



IBM Informix Dynamic Server Guide to SQL: Syntax



IBM Informix Dynamic Server Guide to SQL: Syntax

Note

Before using this information and the product it supports, read the information in "Notices" on page D-1.

This edition replaces SC23-7751-01

This document contains proprietary information of IBM. It is provided under a license agreement and is protected by copyright law. The information contained in this publication does not include any product warranties, and any statements provided in this manual should not be interpreted as such.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1996, 2008.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Introduction	xxiii
In This Introduction	xxiii
About This Publication	xxiii
Types of Users.	xxiii
Software Dependencies	xxiii
Assumptions About Your Locale.	xxiii
Demonstration Databases	xxiv
What's New in SQL Syntax for Dynamic Server, Version 11.50	xxiv
Documentation Conventions	xxx
Typographical Conventions.	xxx
Feature, Product, and Platform Markup	xxxi
Example Code Conventions	xxxi
Additional Documentation	xxxii
Compliance with Industry Standards	xxxii
Syntax Diagrams	xxxii
How to Read a Command-Line Syntax Diagram	xxxiii
Keywords and Punctuation	xxxv
Identifiers and Names	xxxv
How to Provide Documentation Feedback	xxxv
Chapter 1. Overview of SQL Syntax	1-1
How to Enter SQL Statements	1-1
Using Syntax Diagrams and Syntax Tables	1-2
Using Examples	1-2
Using Related Information	1-3
How to Enter SQL Comments	1-3
Examples of SQL Comments	1-4
Non-ASCII Characters in SQL Comments	1-5
Categories of SQL Statements	1-5
Data Definition Language Statements	1-6
Data Manipulation Language Statements	1-8
Data Integrity Statements	1-8
Cursor Manipulation Statements	1-8
Dynamic Management Statements	1-9
Data Access Statements	1-9
Optimization Statements	1-9
Routine Definition Statements	1-10
Auxiliary Statements.	1-10
Client/Server Connection Statements	1-10
Optical Subsystem Statements (IDS)	1-10
ANSI/ISO Compliance and Extensions	1-10
ANSI/ISO-Compliant Statements	1-11
ANSI/ISO-Compliant Statements with Informix Extensions.	1-11
Statements that are Extensions to the ANSI/ISO Standard	1-12
Chapter 2. SQL Statements.	2-1
ALLOCATE COLLECTION	2-1
Examples	2-2
ALLOCATE DESCRIPTOR	2-2
WITH MAX Clause	2-3
ALLOCATE ROW	2-4
ALTER ACCESS_METHOD	2-5
Example	2-6
ALTER FRAGMENT	2-7
Restrictions on the ALTER FRAGMENT Statement	2-8

ALTER FRAGMENT and Transaction Logging	2-9
Determining the Number of Rows in the Fragment	2-9
ATTACH Clause	2-9
DETACH Clause	2-16
INIT Clause	2-18
FRAGMENT BY Clause for Indexes (IDS).	2-21
DROP Clause (IDS)	2-25
MODIFY Clause (IDS)	2-26
ALTER FUNCTION	2-28
Keywords That Introduce Modifications	2-29
ALTER INDEX	2-30
TO CLUSTER Option	2-30
TO NOT CLUSTER Option	2-31
LOCK MODE Options (XPS)	2-31
ALTER PROCEDURE	2-32
ALTER ROUTINE	2-33
Restrictions	2-34
Keywords That Introduce Modifications	2-35
Example of Altering Routine Modifiers	2-35
ALTER SECURITY LABEL COMPONENT	2-36
The ADD ARRAY Clause	2-37
The ADD SET Clause	2-38
The ADD TREE Clause	2-38
ALTER SEQUENCE	2-40
INCREMENT BY Option	2-41
RESTART WITH Option	2-42
MAXVALUE or NOMAXVALUE Option	2-42
MINVALUE or NOMINVALUE Option	2-42
CYCLE or NOCYCLE Option.	2-42
CACHE or NOCACHE Option	2-42
ORDER or NOORDER Option	2-43
Examples	2-43
ALTER TABLE.	2-44
Basic Table Options	2-45
ADD Column Clause	2-47
DEFAULT Clause.	2-49
Single-Column Constraint Format	2-50
DROP Clause	2-55
MODIFY Clause	2-57
Using the MODIFY Clause.	2-58
ADD CONSTRAINT Clause	2-63
Multiple-Column Constraint Format	2-64
DROP CONSTRAINT Clause.	2-66
MODIFY EXTENT SIZE	2-66
MODIFY NEXT SIZE Clause	2-67
LOCK MODE Clause	2-68
Logging TYPE Options	2-69
ADD TYPE Clause (IDS)	2-72
Options Valid on Typed Tables	2-73
BEGIN WORK	2-74
BEGIN WORK and ANSI-Compliant Databases.	2-75
BEGIN WORK WITHOUT REPLICATION (IDS, ESQ/L/C)	2-75
Example of BEGIN WORK	2-75
CLOSE	2-77
Closing a Select or Function Cursor.	2-77
Closing an Insert Cursor	2-78
Closing a Collection Cursor	2-78
Using End of Transaction to Close a Cursor	2-79
Examples	2-79
CLOSE DATABASE	2-80
COMMIT WORK	2-81

Issuing COMMIT WORK in a Database That Is Not ANSI Compliant	2-82
Issuing COMMIT WORK in an ANSI-Compliant Database	2-82
CONNECT	2-83
Privileges for Executing the CONNECT Statement.	2-84
Connection Context	2-84
Database Environment	2-84
Declaring a Connection Name	2-86
USER Validation Clause	2-87
The DEFAULT Connection Specification	2-88
WITH CONCURRENT TRANSACTION Option	2-89
CREATE ACCESS_METHOD	2-90
CREATE AGGREGATE	2-92
Extending the Functionality of Aggregates	2-93
Parallel Execution.	2-94
CREATE CAST	2-95
Source and Target Data Types	2-95
Explicit and Implicit Casts.	2-96
WITH Clause	2-97
CREATE DATABASE	2-97
Logging Options	2-98
Specifying Buffered Logging	2-99
ANSI-Compliant Databases	2-99
CREATE DISTINCT TYPE	2-100
Privileges on Distinct Types	2-101
Support Functions and Casts	2-101
Manipulating Distinct Types.	2-101
CREATE DUPLICATE	2-102
Supported Operations	2-104
CREATE EXTERNAL TABLE (XPS)	2-104
Column Definition	2-105
DATAFILES Clause.	2-108
Table Options.	2-109
Reject Files.	2-111
CREATE FUNCTION	2-113
Privileges Necessary for Using CREATE FUNCTION	2-115
DBA Keyword and Privileges on the Created Function	2-115
The REFERENCING and FOR Clauses	2-116
Overloading the Name of a Function	2-118
DOCUMENT Clause	2-118
WITH LISTING IN Clause	2-119
SPL Functions	2-119
External Procedures (IDS)	2-120
CREATE FUNCTION FROM	2-122
CREATE INDEX	2-124
Index-Type Options.	2-125
UNIQUE or DISTINCT Option	2-126
CLUSTER Option	2-126
Index-Key Specification	2-127
Using a Column or Column List as the Index Key	2-128
Using a Function as an Index Key (IDS)	2-128
Creating Composite Indexes.	2-129
Using the ASC and DESC Sort-Order Options	2-129
Using an Operator Class (IDS)	2-132
USING Access-Method Clause (IDS)	2-132
FILLFACTOR Option	2-133
Providing a Low Percentage Value.	2-134
Providing a High Percentage Value	2-134
Storage Options	2-134
IN Clause	2-135
FRAGMENT BY Clause for Indexes	2-136
Restrictions on Fragmentation Expressions	2-137

Fragmentation of System Indexes	2-138
Fragmentation of Unique Indexes (IDS)	2-138
Fragmentation of Indexes on Temporary Tables	2-138
Index Modes (IDS)	2-138
Specifying Modes for Unique Indexes	2-139
Specifying Modes for Duplicate Indexes	2-140
How the Database Server Treats Disabled Indexes (IDS)	2-140
LOCK MODE Options (XPS)	2-141
Generalized-Key Indexes (XPS)	2-141
The ONLINE Keyword (IDS)	2-143
Automatic Calculation of Distribution Statistics	2-144
CREATE OPAQUE TYPE	2-145
Declaring a Name for an Opaque Type	2-146
INTERNALLENGTH Modifier	2-146
Opaque-Type Modifier	2-147
Defining an Opaque Type	2-147
CREATE OPCLASS	2-149
STRATEGIES Clause	2-151
Strategy Specification	2-151
Indexes on Side-Effect Data	2-152
SUPPORT Clause	2-152
Default Operator Classes	2-153
CREATE PROCEDURE	2-153
Using CREATE PROCEDURE Versus CREATE FUNCTION	2-155
Relationship Between Routines, Functions, and Procedures	2-155
Privileges Necessary for Using CREATE PROCEDURE	2-156
DBA Keyword and Privileges on the Procedure	2-156
The REFERENCING and FOR Clauses	2-157
Procedure Names in Extended Parallel Server	2-159
Procedure Names in Dynamic Server	2-159
DOCUMENT Clause	2-159
Using the WITH LISTING IN Option	2-160
SPL Procedures	2-160
External Procedures (IDS)	2-161
CREATE PROCEDURE FROM	2-163
Default Directory That Holds the File	2-164
CREATE ROLE	2-164
CREATE ROUTINE FROM	2-166
Examples	2-167
CREATE ROW TYPE	2-168
Privileges on Named Row Data Types	2-169
Inheritance and Named ROW Types	2-169
Creating a Subtype	2-169
Type Hierarchies	2-170
Procedure for Creating a Subtype	2-170
Field Definition	2-170
Restrictions on Serial and Simple-Large-Object Data Types	2-171
CREATE SCHEMA	2-172
Creating Database Objects Within CREATE SCHEMA	2-173
CREATE SCRATCH TABLE	2-174
Examples	2-175
CREATE SECURITY LABEL	2-175
Components and Elements of a Security Label	2-176
CREATE SECURITY LABEL COMPONENT	2-177
Types and Elements of Security Components	2-178
ARRAY Components	2-179
SET Components	2-180
TREE Components	2-180
CREATE SECURITY POLICY	2-181
Security Label Components of a Security Policy	2-182
Rules Associated with a Security Policy	2-182

CREATE SEQUENCE	2-184
INCREMENT BY Option	2-185
START WITH Option	2-186
MAXVALUE or NOMAXVALUE Option	2-186
MINVALUE or NOMINVALUE Option	2-186
CYCLE or NOCYCLE Option	2-186
CACHE or NOCACHE Option	2-186
ORDER or NOORDER Option	2-187
Examples	2-187
CREATE SYNONYM	2-187
Synonyms for Remote and External Tables and Views	2-188
PUBLIC and PRIVATE Synonyms	2-189
Synonyms with the Same Name	2-189
Chaining Synonyms	2-190
CREATE TABLE	2-191
Logging Options	2-193
Column Definition	2-193
DEFAULT Clause	2-195
Single-Column Constraint Format	2-197
REFERENCES Clause	2-199
CHECK Clause	2-202
Constraint Definition	2-203
Multiple-Column Constraint Format	2-205
Options Clauses	2-209
Storage Options	2-210
FRAGMENT BY Clause	2-212
RANGE Method Clause (XPS)	2-218
PUT Clause (IDS)	2-222
EXTENT SIZE Options	2-224
USING Access-Method Clause (IDS)	2-224
LOCK MODE Options	2-225
OF TYPE Clause (IDS)	2-227
CREATE TEMP TABLE	2-230
CREATE Temporary TABLE	2-231
Using the TEMP Option	2-232
Using the SCRATCH Option (XPS)	2-232
Naming a Temporary Table	2-232
Using the WITH NO LOG Option	2-232
Column Definition	2-233
Single-Column Constraint Format	2-234
Multiple-Column Constraint Format	2-234
Temporary Table Options	2-235
Storage Options	2-236
Where Temporary Tables are Stored	2-237
Differences between Temporary and Permanent Tables	2-238
Duration of Temporary Tables	2-238
CREATE TRIGGER	2-239
Defining a Trigger Event and Action	2-240
Restrictions on Triggers	2-242
Trigger Modes (IDS)	2-242
Trigger Inheritance in a Table Hierarchy (IDS)	2-243
Triggers and SPL Routines	2-244
Trigger Events	2-244
INSERT Events and DELETE Events	2-246
UPDATE Event	2-246
Defining Multiple Update Triggers	2-247
SELECT Event (IDS)	2-247
Circumstances When a Select Trigger Is Activated	2-248
Stand-alone SELECT Statements	2-249
SELECT Statements Within UDRs in the Select List	2-249
UDRs That EXECUTE PROCEDURE and EXECUTE FUNCTION Call	2-249

Subqueries in the Select List	2-249
Subqueries in the FROM Clause of SELECT	2-250
Subqueries in the WHERE Clause of DELETE or UPDATE	2-250
Select Triggers in Table Hierarchies	2-250
Circumstances When a Select Trigger Is Not Activated	2-251
Action Clause	2-251
Guaranteeing Row-Order Independence	2-253
REFERENCING Clauses	2-254
Correlated Table Action	2-257
Triggered Action.	2-257
Using Correlation Names in Triggered Actions	2-260
Re-Entrancy of Triggers	2-262
Rules for SPL Routines	2-265
Privileges to Execute Trigger Actions	2-266
Cascading Triggers	2-267
External Tables	2-269
Logging and Recovery.	2-269
INSTEAD OF Triggers on Views (IDS)	2-270
CREATE VIEW	2-274
Typed Views	2-275
Subset of SELECT Statements Valid in View Definitions	2-276
Union Views	2-276
Naming View Columns	2-276
Using a View in the SELECT Statement	2-277
WITH CHECK OPTION Keywords	2-277
Updating Through Views.	2-278
CREATE XADATASOURCE	2-279
CREATE XADATASOURCE TYPE	2-280
DATABASE	2-281
SQLCA.SQLWARN Settings Immediately after DATABASE Executes (ESQL/C).	2-282
EXCLUSIVE Keyword	2-282
DEALLOCATE COLLECTION	2-283
DEALLOCATE DESCRIPTOR	2-284
DEALLOCATE ROW	2-285
DECLARE	2-286
Overview of Cursor Types	2-288
Select Cursor or Function Cursor	2-289
Cursor Characteristics	2-293
Associating a Cursor with a Prepared Statement	2-297
Using Cursors with Transactions	2-300
Declaring a Dynamic Cursor in an SPL Routine	2-301
Syntax	2-301
Usage	2-302
DELETE	2-303
Using the ONLY Keyword (IDS)	2-305
Considerations When Tables Have Cascading Deletes	2-305
Using the WHERE Keyword to Specify a Condition	2-306
Subqueries in the WHERE Clause of DELETE	2-306
Using the WHERE CURRENT OF Keywords (ESQL/C, SPL).	2-307
Using the USING or FROM Keyword to Introduce a Join Condition (XPS)	2-308
Deleting Rows That Contain Opaque Data Types (IDS)	2-309
Deleting Rows That Contain Collection Data Types (IDS)	2-309
Data Types in Distributed DELETE Operations (IDS)	2-309
SQLSTATE Values in an ANSI-Compliant Database	2-310
SQLSTATE Values in a Database That Is Not ANSI-Compliant	2-310
DESCRIBE.	2-311
The OUTPUT Keyword	2-312
Describing the Statement Type	2-312
Checking for the Existence of a WHERE Clause	2-312
Describing a Statement with Runtime Parameters.	2-312
Using the SQL DESCRIPTOR Keywords.	2-313

Using the INTO sqllda Pointer Clause	2-314
Describing a Collection Variable (IDS)	2-314
DESCRIBE INPUT	2-315
Describing the Statement Type	2-316
Checking for Existence of a WHERE Clause	2-317
Describing a Statement with Dynamic Runtime Parameters	2-317
Using the SQL DESCRIPTOR Keywords	2-317
Using the INTO sqllda Pointer Clause	2-318
Describing a Collection Variable	2-319
DISCONNECT	2-320
DEFAULT Option	2-321
Specifying the CURRENT Keyword	2-321
When a Transaction is Active	2-321
Disconnecting in a Thread-Safe Environment	2-321
Specifying the ALL Option	2-322
DROP ACCESS_METHOD	2-322
Examples	2-323
DROP AGGREGATE	2-323
DROP CAST	2-324
Examples	2-325
DROP DATABASE	2-325
DROP DUPLICATE	2-326
DROP FUNCTION	2-327
Dropping External Functions	2-328
Examples	2-328
DROP INDEX	2-329
The ONLINE Keyword (IDS)	2-330
Examples	2-331
DROP OPCLASS	2-331
DROP PROCEDURE	2-332
Dropping an External Procedure (IDS)	2-333
DROP ROLE	2-334
DROP ROUTINE	2-335
Restrictions	2-335
Modifying SPL Routines (XPS)	2-336
Dropping an External Routine	2-336
DROP ROW TYPE	2-336
The RESTRICT Keyword	2-337
DROP SECURITY	2-338
Restrictions on Dropping Security Objects	2-339
Examples of Dropping Security Objects	2-339
DROP SEQUENCE	2-340
Examples	2-341
DROP SYNONYM	2-341
DROP TABLE	2-342
Effects of the DROP TABLE Statement	2-342
Specifying CASCADE Mode	2-343
Specifying RESTRICT Mode	2-343
Dropping a Table That Contains Opaque Data Types (IDS)	2-343
Tables That Cannot Be Dropped	2-343
DROP TRIGGER	2-344
DROP TYPE	2-345
DROP VIEW	2-346
DROP XADATASOURCE	2-347
DROP XADATASOURCE TYPE	2-347
EXECUTE	2-348
Scope of Statement Identifiers	2-349
Restrictions with the INTO Clause	2-351
Replacing Placeholders with Parameters	2-351
Saving Values In Host or Program Variables	2-351
Saving Values in a System-Descriptor Area	2-351

Saving Values in an sqlca Structure (ESQL/C)	2-352
The sqlca Record and EXECUTE	2-353
Returned SQLCODE Values with EXECUTE	2-353
Supplying Parameters Through Host or Program Variables	2-354
Supplying Parameters Through a System Descriptor	2-355
Supplying Parameters Through an sqlca Structure (ESQL/C)	2-355
EXECUTE FUNCTION	2-356
Negator Functions and Their Companions	2-357
How the EXECUTE FUNCTION Statement Works	2-357
Data Variables	2-358
INTO Clause with Indicator Variables (ESQL/C)	2-358
INTO Clause with Cursors	2-359
Alternatives to PREPARE ... EXECUTE FUNCTION ... INTO	2-359
Dynamic Routine-Name Specification of SPL Functions	2-360
EXECUTE IMMEDIATE	2-360
EXECUTE IMMEDIATE and Restricted Statements	2-361
Restrictions on Valid Statements	2-362
Handling Exceptions from EXECUTE IMMEDIATE Statements	2-363
Examples of the EXECUTE IMMEDIATE Statement	2-363
EXECUTE PROCEDURE	2-364
Causes of Errors	2-365
Using the INTO Clause	2-365
The WITH TRIGGER REFERENCES Keywords	2-366
Dynamic Routine-Name Specification of SPL Procedures	2-367
FETCH	2-367
FETCH with a Sequential Cursor	2-369
FETCH with a Scroll Cursor	2-369
How the Database Server Implements Scroll Cursors	2-370
Specifying Where Values Go in Memory	2-370
Using the INTO Clause	2-371
Using Indicator Variables	2-371
When the INTO Clause of FETCH is Required	2-371
Using a System-Descriptor Area (X/Open)	2-372
Using sqlca Structures	2-373
Fetching a Row for Update	2-373
Fetching from a Collection Cursor (IDS)	2-374
Checking the Result of FETCH	2-375
FLUSH	2-377
Error Checking FLUSH Statements	2-377
Examples	2-378
FREE	2-379
GET DESCRIPTOR	2-381
Using the COUNT Keyword	2-382
Using the VALUE Clause	2-382
Using LENGTH or ILENGTH	2-384
Describing an Opaque-Type Column (IDS)	2-384
Describing a Distinct-Type Column (IDS)	2-384
Examples	2-385
GET DIAGNOSTICS	2-386
Using the SQLSTATE Error Status Code	2-387
Statement Clause	2-390
EXCEPTION Clause	2-391
The Contents of the SERVER_NAME Field	2-393
The Contents of the CONNECTION_NAME Field	2-394
Using GET DIAGNOSTICS for Error Checking	2-394
GRANT	2-395
Database-Level Privileges	2-397
Table-Level Privileges	2-399
Table Reference	2-402
Type-Level Privileges (IDS)	2-403
Routine-Level Privileges	2-405

Language-Level Privileges (IDS)	2-406
Sequence-Level Privileges (IDS).	2-407
Role Name	2-408
WITH GRANT OPTION Keywords	2-411
AS grantor Clause	2-412
Security Administration Options	2-413
GRANT FRAGMENT	2-423
Fragment-Level Privileges	2-424
Granting Privileges to One User or a List of Users	2-426
Granting One Privilege or a List of Privileges	2-427
WITH GRANT OPTION Clause	2-427
AS grantor Clause	2-427
INFO	2-428
Examples	2-430
INSERT	2-430
Specifying Columns	2-432
Using the AT Clause (ESQL/C, IDS, and SPL).	2-432
Inserting Rows Through a View	2-432
Inserting Rows with a Cursor	2-433
Inserting Rows into a Database Without Transactions	2-433
Inserting Rows into a Database with Transactions	2-434
VALUES Clause	2-434
Execute Routine Clause	2-441
LOAD	2-443
LOAD FROM File	2-444
Loading Simple Large Objects	2-446
Loading Smart Large Objects (IDS)	2-447
Loading Complex Data Types (IDS)	2-447
Loading Opaque-Type Columns (IDS)	2-448
DELIMITER Clause.	2-448
INSERT INTO Clause	2-448
LOCK TABLE	2-449
Concurrent Access to Tables with Shared Locks	2-450
Concurrent Access to Tables with Exclusive Locks	2-450
Databases with Transactions	2-451
Databases Without Transactions	2-451
Locking Granularity	2-452
MERGE	2-452
Target Table Restrictions	2-454
Handling Duplicate Rows	2-454
MOVE TABLE	2-455
Considerations When Using the MOVE TABLE Statement	2-456
Moving Tables with Constraints	2-457
Renaming the Table	2-457
Preserving Table Privileges	2-458
The RESTRICT and CASCADE Options	2-458
Physically Moving Data to Another Database	2-459
OPEN	2-459
Opening a Select Cursor	2-460
Opening an Update Cursor Inside a Transaction	2-461
Opening a Function Cursor	2-461
Reopening a Select or Function Cursor	2-461
Errors Associated with Select and Function Cursors	2-462
Opening an Insert Cursor (ESQL/C)	2-462
Opening a Collection Cursor (ESQL/C)	2-463
USING Clause	2-463
Specifying a System Descriptor Area (ESQL/C)	2-464
Specifying a Pointer to an sqllda Structure (ESQL/C)	2-465
Using the WITH REOPTIMIZATION Option (ESQL/C)	2-465
Relationship Between OPEN and FREE	2-465
DDL Operations on Tables Referenced by Cursors	2-466

OUTPUT	2-466
Sending Query Results to a File	2-467
Displaying Query Results Without Column Headings	2-467
Sending Query Results to Another Program	2-467
PREPARE	2-468
Restrictions	2-469
Declaring a Statement Identifier	2-469
Releasing a Statement Identifier	2-470
Statement Text	2-470
Preparing and Executing User-Defined Routines	2-471
Restricted Statements in Single-Statement Prepares	2-471
Preparing Statements When Parameters Are Known	2-473
Preparing Statements That Receive Parameters	2-473
Preparing Statements with SQL Identifiers	2-474
Preparing Multiple SQL Statements	2-476
Using Prepared Statements for Efficiency	2-477
PUT	2-479
Supplying Inserted Values	2-480
Using the USING Clause	2-482
Inserting into a Collection Cursor (IDS)	2-483
Writing Buffered Rows	2-484
Error Checking	2-484
RELEASE SAVEPOINT	2-485
RENAME COLUMN	2-487
How Views and Check Constraints Are Affected	2-487
How Triggers Are Affected	2-487
Example of RENAME COLUMN	2-488
RENAME DATABASE	2-488
RENAME INDEX	2-489
RENAME SECURITY	2-490
RENAME SEQUENCE	2-491
RENAME TABLE	2-492
REVOKE	2-494
Database-Level Privileges	2-496
Table-Level Privileges	2-497
Effect of Uncommitted Transactions	2-500
Type-Level Privileges (IDS)	2-500
Routine-Level Privileges	2-501
Language-Level Privileges (IDS)	2-503
Sequence-Level Privileges (IDS)	2-503
User List	2-504
Role Name	2-505
Revoking Privileges Granted WITH GRANT OPTION	2-507
The AS Clause	2-507
Controlling the Scope of REVOKE with the RESTRICT Option	2-508
Security Administration Options	2-508
REVOKE FRAGMENT	2-517
Specifying Fragments	2-518
The FROM Clause	2-518
Fragment-Level Privileges	2-518
The AS Clause	2-519
Examples of the REVOKE FRAGMENT Statement	2-519
ROLLBACK WORK	2-520
WORK Keyword	2-521
TO SAVEPOINT Clause	2-521
SAVE EXTERNAL DIRECTIVES	2-522
External Optimizer Directives	2-523
Enabling or Disabling External Directives for a Session	2-523
The directive Specification	2-524
The ACTIVE, INACTIVE, and TEST ONLY Keywords	2-525
The query Specification	2-525

SAVEPOINT	2-526
SELECT	2-528
Projection Clause	2-531
INTO Clause	2-541
FROM Clause	2-544
WHERE Clause of SELECT	2-560
GROUP BY Clause	2-566
HAVING Clause.	2-567
ORDER BY Clause	2-568
FOR UPDATE Clause	2-572
FOR READ ONLY Clause	2-573
INTO Table Clauses	2-574
UNION Operator	2-578
SET ALL_MUTABLES	2-581
SET AUTOFREE	2-582
Globally Affecting Cursors with SET AUTOFREE.	2-583
Using the FOR Clause to Specify a Specific Cursor	2-583
Associated and Detached Statements	2-583
Closing Cursors Implicitly	2-584
SET COLLATION	2-584
Specifying a Collating Order with SET COLLATION	2-585
Restrictions on SET COLLATION	2-585
Collation Performed by Database Objects	2-586
SET CONNECTION	2-587
Making a Dormant Connection the Current Connection	2-587
Making a Current Connection the Dormant Connection	2-588
Dormant Connections in a Single-Threaded Environment	2-588
Dormant Connections in a Thread-Safe Environment	2-588
Identifying the Connection	2-589
DEFAULT Option	2-589
CURRENT Keyword	2-590
When a Transaction is Active	2-590
SET CONSTRAINTS	2-590
SET Database Object Mode	2-591
Privileges Required for Changing Database Object Modes	2-592
Object-List Format	2-593
Table Format	2-593
Modes for Constraints and Unique Indexes.	2-594
Definitions of Database Object Modes	2-595
SET DATASKIP	2-598
Circumstances When a Dbspace Cannot Be Skipped.	2-598
Examples	2-599
SET DEBUG FILE	2-599
Using the WITH APPEND Option.	2-600
Closing the Output File	2-600
Redirecting Trace Output	2-600
Location of the Output File	2-600
SET Default Table Space	2-601
SET Default Table Type	2-602
SET DEFERRED_PREPARE	2-603
Example of SET DEFERRED_PREPARE	2-604
Using Deferred-Prepare with OPTOFC	2-604
SET DESCRIPTOR	2-605
Using the COUNT Clause	2-606
Using the VALUE Clause.	2-606
Item Descriptor	2-606
Modifying Values Set by the DESCRIBE Statement	2-611
SET ENCRYPTION PASSWORD	2-611
Storage Requirements for Encryption	2-612
Specifying a Session Password and Hint.	2-613
Levels of Encryption	2-613

Protecting Passwords	2-614
SET ENVIRONMENT	2-615
BOUND_IMPL_PDQ Environment Option (XPS)	2-616
CLIENT_TZ Environment Option (XPS)	2-617
COMPUTE_QUOTA Environment Option (XPS)	2-618
IMPLICIT_PDQ Environment Option (XPS)	2-618
MAXSCAN Environment Option (XPS)	2-618
TMPSPACE_LIMIT Environment Option (XPS)	2-619
TEMP_TAB_EXT_SIZE and TEMP_TAB_NEXT_SIZE Options (XPS)	2-619
EXTDIRECTIVES Environment Option (IDS)	2-620
IFX_AUTO_REPREPARE Environment Option (IDS)	2-620
OPTCOMPIND Environment Option (IDS)	2-621
USELASTCOMMITTED Environment Option (IDS)	2-622
SET EXPLAIN	2-623
Using the AVOID_EXECUTE Option	2-625
Using the FILE TO Option	2-626
Using the WITH APPEND Option (XPS)	2-626
Default Name and Location of the Output File on UNIX	2-626
Default Name and Location of the Output File on Windows	2-627
SET EXPLAIN Output	2-627
SET INDEX	2-630
SET INDEXES	2-631
SET ISOLATION	2-631
Complete-Connection Level Settings	2-633
Informix Isolation Levels	2-633
Effects of Isolation Levels	2-637
SET LOCK MODE	2-637
WAIT Clause	2-638
SET LOG	2-639
SET OPTIMIZATION	2-641
HIGH and LOW Options	2-641
FIRST_ROWS and ALL_ROWS Options (IDS)	2-642
Optimizing SPL Routines	2-642
Examples	2-642
SET PDQPRIORITY	2-643
Allocating Database Server Resources	2-644
Using a Range of Values	2-645
SET PLOAD FILE	2-645
SET Residency	2-646
Residency and the Changing Status of Fragments	2-647
Examples	2-647
SET ROLE	2-647
Setting the Default Role	2-649
SET SCHEDULE LEVEL	2-649
SET SESSION AUTHORIZATION	2-650
SET SESSION AUTHORIZATION and Transactions	2-652
SET STATEMENT CACHE	2-652
Precedence and Default Behavior	2-652
Turning the Cache ON	2-653
Turning the Cache OFF	2-654
Statement Qualification	2-654
SET TABLE	2-655
SET TRANSACTION	2-656
Comparing SET TRANSACTION with SET ISOLATION	2-657
Informix Isolation Levels	2-658
Default Isolation Levels	2-659
Access Modes	2-660
Effects of Isolation Levels	2-660
SET Transaction Mode	2-661
Statement-Level Checking	2-661
Transaction-Level Checking	2-661

Duration of Transaction Modes	2-662
Specifying All Constraints or a List of Constraints	2-662
Specifying Remote Constraints	2-662
Examples of Setting the Transaction Mode for Constraints	2-662
SET TRIGGERS	2-663
START VIOLATIONS TABLE	2-663
Relationship to SET Database Object Mode Statement	2-665
Effect on Concurrent Transactions	2-665
Stopping the Violations and Diagnostics Tables	2-666
USING Clause	2-666
Using the MAX ROWS Clause (IDS)	2-666
Using the MAX VIOLATIONS Clause (XPS)	2-666
Specifying the Maximum Number of Rows in the Diagnostics Table (IDS)	2-667
Specifying the Maximum Number of Rows in the Violations Table (XPS)	2-667
Privileges Required for Starting Violations Tables	2-667
Structure of the Violations Table	2-667
Examples of START VIOLATIONS TABLE Statements	2-668
Relationships Among the Target, Violations, and Diagnostics Tables	2-668
Initial Privileges on the Violations Table	2-669
Example of Privileges on the Violations Table	2-670
Using the Violations Table	2-671
Example of a Violations Table	2-672
Structure of the Diagnostics Table (IDS)	2-672
Initial Privileges on the Diagnostics Table	2-673
Using the Diagnostics Table	2-675
STOP VIOLATIONS TABLE	2-676
Example of Stopping the Violations and Diagnostics Tables	2-677
Example of Dropping the Violations and Diagnostics Tables	2-677
Privileges Required for Stopping a Violations Table	2-677
TRUNCATE (IDS)	2-677
The TABLE Keyword	2-679
The Table Specification	2-679
The Storage Specification	2-679
The AM_TRUNCATE Purpose Function	2-680
Performance Advantages of TRUNCATE	2-680
Restrictions	2-681
TRUNCATE (XPS)	2-682
Restrictions	2-682
Using the ONLY and TABLE Keywords	2-683
After the TRUNCATE Statement Executes	2-683
When You Might Use the TRUNCATE Statement	2-683
UNLOAD	2-683
UNLOAD TO File	2-684
DELIMITER Clause	2-687
UNLOCK TABLE	2-688
UPDATE	2-689
Using the ONLY Keyword (IDS)	2-691
Updating Rows Through a View	2-691
Updating Rows in a Database Without Transactions	2-692
Updating Rows in a Database with Transactions	2-692
Locking Considerations	2-692
SET Clause	2-693
UPDATE STATISTICS	2-705
Scope of UPDATE STATISTICS	2-706
Updating Statistics for Tables	2-706
Updating Statistics for Columns of User-Defined Types (IDS)	2-708
Using the LOW Mode Option	2-709
Using the MEDIUM Mode Option	2-710
Using the HIGH Mode Option	2-711
Resolution Clause	2-711
Routine Statistics	2-714

Updating Statistics When You Upgrade the Database Server (IDS)	2-717
Performance	2-717
WHENEVER	2-718
The Scope of WHENEVER	2-719
SQLERROR Keyword	2-720
ERROR Keyword	2-720
SQLWARNING Keyword	2-720
NOT FOUND Keywords	2-721
CONTINUE Keyword	2-721
STOP Keyword	2-721
GOTO Keyword	2-721
CALL Clause	2-722

Chapter 3. SPL Statements 3-1

<< Label >>	3-1
Examples of Labels	3-2
CALL	3-3
Receiving Input from the Called UDR	3-4
Receiving Input from the Called UDR	3-4
CASE	3-5
How the Database Server Executes a CASE Statement	3-5
Computation of the Value Expression in CASE	3-6
Example of CASE Statement	3-6
CONTINUE	3-7
DEFINE	3-8
Referencing TEXT and BYTE Variables	3-9
Redeclaration or Redefinition	3-9
Declaring Global Variables	3-9
Declaring Local Variables	3-12
EXIT	3-17
EXIT From FOREACH Statements	3-17
FOR	3-19
Using the TO Keyword to Define a Range	3-20
Using an Expression List as the Range	3-21
Mixing Range and Expression Lists in the Same FOR Statement	3-21
Specifying a Labelled FOR Loop	3-22
FOREACH	3-22
Using a SELECT ... INTO Statement	3-24
Using the ORDER BY Clause of the SELECT Statement	3-25
Using Hold Cursors	3-25
Updating or Deleting Rows Identified by Cursor Name	3-25
Using Collection Variables (IDS)	3-25
Using Select Cursors with FOREACH	3-26
Calling a UDR in the FOREACH Loop	3-27
GOTO	3-27
IF	3-28
ELIF Clause	3-29
ELSE Clause	3-29
Conditions in an IF Statement	3-30
Subset of SPL Statements Allowed in the IF Statement List	3-30
SQL Statements Not Valid in an IF Statement	3-31
LET	3-32
Using a SELECT Statement in a LET Statement	3-33
Calling a Function in a LET Statement	3-33
LOOP	3-34
Simple LOOP Statements	3-35
FOR LOOP Statements	3-36
WHILE LOOP Statements	3-36
Labeled LOOP Statements	3-36
ON EXCEPTION	3-38
Placement of the ON EXCEPTION Statement	3-39

Using the IN Clause to Trap Specific Exceptions	3-40
Receiving Error Information in the SET Clause	3-40
Forcing Continuation of the Routine	3-40
RAISE EXCEPTION	3-41
Special Error Number -746.	3-42
RETURN	3-42
WITH RESUME Keyword	3-43
SYSTEM	3-44
Executing the SYSTEM Statement on UNIX	3-45
Executing the SYSTEM Statement on Windows	3-46
Setting Environment Variables in SYSTEM Commands	3-46
TRACE	3-46
TRACE ON.	3-47
TRACE OFF	3-47
TRACE PROCEDURE	3-47
Displaying Expressions	3-48
Example Showing Different Forms of TRACE	3-48
Looking at the Traced Output	3-48
WHILE	3-49
Example of WHILE Loops in an SPL Routine	3-49
Labeled WHILE Loops	3-50
Chapter 4. Data Types and Expressions	4-1
Scope of Segment Descriptions.	4-1
Use of Segment Descriptions	4-1
Segments in This Chapter	4-2
Collection Subquery	4-3
Table Expressions in the FROM Clause	4-5
Condition	4-5
Comparison Conditions (Boolean Expressions)	4-6
Column Name	4-8
Quotation Marks in Conditions	4-8
Relational-Operator Condition	4-9
BETWEEN Condition	4-9
IN Condition	4-10
IS NULL and IS NOT NULL Conditions	4-12
Trigger-Type Boolean Operator	4-12
LIKE and MATCHES Condition	4-13
Stand-Alone Condition	4-15
Condition with Subquery	4-15
NOT Operator.	4-19
Conditions with AND or OR	4-20
Data Type	4-21
Built-In Data Types	4-21
User-Defined Data Type (IDS)	4-32
Complex Data Type (IDS)	4-34
DATETIME Field Qualifier.	4-38
Expression	4-40
Syntax of SQL Expressions.	4-40
Usage.	4-42
List of Expressions	4-42
Arithmetic Operators	4-51
Bitwise Logical Functions	4-52
Concatenation Operator	4-56
Cast Expressions (IDS)	4-58
Column Expressions	4-59
Conditional Expressions	4-65
Constant Expressions	4-70
Constructor Expressions (IDS)	4-82
NULL Keyword	4-85
Function Expressions	4-87

Statement-Local Variable Expressions (IDS)	4-164
Aggregate Expressions	4-165
INTERVAL Field Qualifier	4-176
Literal Collection	4-178
Element Literal Value	4-179
Nested Quotation Marks	4-180
Literal DATETIME	4-181
Casting Numeric Date and Time Strings to DATE Data Types	4-182
Literal INTERVAL	4-183
Literal Number	4-185
Integer Literals	4-185
Fixed-Point Decimal Literals	4-186
Floating-Point Decimal Literals	4-186
Literal Numbers and the MONEY Data Type	4-186
Literal Row	4-186
Literals of an Unnamed Row Type	4-188
Literals of a Named Row Type	4-188
Literals for Nested Rows	4-189
Quoted String	4-189
Restrictions on Specifying Characters in Quoted Strings	4-190
The DELIMIDENT Environment Variable	4-190
Newline Characters in Quoted Strings	4-191
Using Quotation Marks in Strings	4-192
DATETIME and INTERVAL Values as Strings	4-192
LIKE and MATCHES in a Condition	4-192
Inserting Values as Quoted Strings	4-192
Numeric Operations on Character Columns	4-193
Relational Operator	4-193
Using Operator Functions in Place of Relational Operators	4-194
Collating Order for U.S. English Data	4-195
Support for ASCII Characters in Nondefault Code Sets (GLS)	4-196
Literal Numbers as Operands	4-196
Chapter 5. Other Syntax Segments	5-1
Arguments	5-1
Comparing Arguments to the Parameter List	5-2
Subset of Expressions Valid as an Argument	5-3
Arguments to UDRs in Remote Databases	5-3
Collection-Derived Table	5-4
Accessing a Collection Through a Virtual Table	5-5
Table Expressions in the FROM Clause	5-6
Restrictions with the Collection-Expression Format	5-6
Row Type of the Resulting Collection-Derived Table	5-7
Accessing a Collection Through a Collection Variable	5-10
Using a Collection Variable to Manipulate Collection Elements	5-10
Accessing a Nested Collection	5-14
Accessing a Row Variable	5-14
Database Name	5-15
Using Keywords as Table Names	5-16
Database Object Name	5-16
Specifying a Database Object in an External Database	5-17
Routine Overloading and Routine Signatures (IDS)	5-19
Owners of Objects Created by UDRs	5-19
External Routine Reference	5-19
VARIANT or NOT VARIANT Option	5-20
Example of a C User-Defined Function	5-21
Examples	5-21
Identifier	5-21
Use of Uppercase Characters	5-22
Use of Keywords as Identifiers	5-23
Support for Non-ASCII Characters in Identifiers	5-23

Delimited Identifiers	5-23
Enabling Delimited Identifiers	5-24
Potential Ambiguities and Syntax Errors	5-26
Using the Names of Built-In Functions as Column Names	5-26
Using Keywords as Column Names	5-27
Using ALL, DISTINCT, or UNIQUE as a Column Name	5-27
Using INTERVAL or DATETIME as a Column Name	5-27
Using rowid as a Column Name (IDS).	5-28
Using Keywords as Table Names	5-28
Jar Name	5-34
Optimizer Directives	5-35
Optimizer Directives as Comments	5-36
Restrictions on Optimizer Directives	5-37
Access-Method Directives	5-37
Join-Order Directive	5-40
Join-Method Directives	5-40
Optimization-Goal Directives (IDS)	5-42
Explain-Mode Directives	5-43
Rewrite Method Directive (XPS)	5-44
External Directives	5-45
Owner Name	5-45
Using Quotation Marks.	5-46
Referencing Tables Owned by User informix.	5-47
ANSI-Compliant Database Restrictions and Case Sensitivity	5-48
Setting ANSIOWNER for an ANSI-Compliant Database	5-49
Default Owner Names	5-49
Summary of Lettercase Rules for Owner Names	5-49
Purpose Options	5-50
Purpose Options for Access Methods	5-51
Purpose Functions, Methods, Flags, and Values.	5-52
Purpose Options for XA Data Source Types	5-54
Return Clause	5-55
Limits on Returned Values.	5-56
Subset of SQL Data Types	5-56
Using the REFERENCES Clause to Point to a Simple Large Object	5-57
Returning a Value from Another Database (IDS)	5-57
Named Return Parameters (IDS).	5-58
Cursor and Noncursor Functions	5-59
Routine Modifier	5-59
Adding or Modifying a Routine Modifier	5-60
Modifier Descriptions	5-61
CLASS	5-61
COSTFUNC	5-62
HANDLESNULLS	5-62
INTERNAL.	5-62
ITERATOR	5-63
NEGATOR	5-63
PARALLELIZABLE	5-63
PERCALL_COST (C)	5-64
SELCONST (C)	5-65
SELFUNC (C)	5-65
STACK (C)	5-66
VARIANT and NOT VARIANT	5-66
Examples	5-66
Routine Parameter List	5-67
Subset of SQL Data Types	5-68
Using the LIKE Clause	5-68
Using the REFERENCES Clause	5-68
Using the DEFAULT Clause	5-69
Specifying OUT Parameters for User-Defined Routines	5-69
Specifying INOUT Parameters for a User-Defined Routine (IDS)	5-70

Shared-Object Filename	5-70
C Shared-Object File	5-71
Java Shared-Object File	5-72
Specific Name	5-73
Restrictions on the Owner Name	5-74
Restrictions on the Specific Name	5-74
Statement Block	5-74
Subset of SPL Statements Valid in the Statement Block	5-75
SQL Statements Valid in SPL Statement Blocks	5-76
Nested Statement Blocks	5-77
Restrictions on SPL Routines in Data-Manipulation Statements	5-78
Transactions in SPL Routines	5-79
Support for Roles and User Identity	5-79

Chapter 6. Built-In Routines 6-1

SQL Admin API Functions	6-1
ADMIN and TASK Arguments and Size Specifications	6-3
Mnemonic Initial Commands	6-6
The ONMODE Initial Command	6-27
Calling Syntax of the "ONMODE" Commands	6-29
onmode a: Add a shared-memory segment	6-30
onmode c: Force a checkpoint	6-30
onmode C: Control the B-tree scanner	6-31
onmode d: Set data-replication types	6-32
onmode D: Set PDQ	6-33
onmode e: Change usage of the SQL statement cache	6-33
onmode F: Free unused memory segments	6-34
onmode j: Switch database server to administration mode	6-34
onmode l: Switch to next logical log	6-35
onmode m: Switch to multi-user mode	6-35
onmode M: Change decision-support memory	6-35
onmode n: Unlock resident memory	6-36
onmode O: Override a space down that blocks a checkpoint	6-36
onmode p: Add or remove virtual processors	6-36
onmode Q: Set maximum number for decision-support queries	6-39
onmode r: Lock memory resident	6-39
onmode S: Specify maximum number of decision support scans	6-39
onmode W: Reset statement cache attributes	6-40
onmode {wf wm}: Dynamically update a configuration parameter	6-40
onmode wm: Dynamically change LRU tuning status	6-41
onmode Y: Dynamically change SET EXPLAIN for a user session	6-43
onmode z: Terminate a user session	6-44
onmode Z: Terminate a distributed transaction ungracefully	6-44
High Availability Initial Commands	6-44
Enterprise Replication Initial Commands	6-48
Restricting the Size of the command_history Table	6-49
Session Configuration Procedures	6-50
Using SYSDBOPEN and SYSDBCLOSE Procedures	6-51
The EXPLAIN_SQL Routine	6-53
UDR Definition Routines	6-54
IFX_REPLACE_MODULE Function	6-54
IFX_UNLOAD_MODULE Function	6-55
jvpcontrol Function	6-56
Using the MEMORY Keyword	6-57
Using the THREADS Keyword	6-57
SQLJ Driver Built-In Procedures	6-57
sqlj.install_jar	6-58
sqlj.replace_jar	6-59
sqlj.remove_jar	6-59
sqlj.alter_java_path	6-60
sqlj.setUDTextName	6-61

sqlj.unsetUDTextName	6-62
DRDA Support Functions	6-62
Metadata Function	6-62
sysibm.SQLCAMessage Function	6-64
Appendix A. Reserved Words for IBM Informix Dynamic Server	A-1
Appendix B. Reserved Words for IBM Informix Extended Parallel Server	B-1
Appendix C. Accessibility	C-1
Accessibility features for IBM Informix Dynamic Server	C-1
Accessibility Features.	C-1
Keyboard Navigation.	C-1
Related Accessibility Information.	C-1
IBM and Accessibility	C-1
Dotted Decimal Syntax Diagrams	C-1
Notices	D-1
Trademarks	D-3
Index	X-1

Introduction

In This Introduction

This introduction provides an overview of the information in this publication and describes the documentation conventions that it uses.

About This Publication

This publication describes the syntax of the Structured Query Language (SQL) and of the Stored Procedure Language (SPL) for Version 11.50 of IBM® Informix® Dynamic Server and Version 8.50 of IBM Informix Extended Parallel Server.

This publication is a companion volume to the *IBM Informix Guide to SQL: Reference*, the *IBM Informix Guide to SQL: Tutorial*, and the *IBM Informix Database Design and Implementation Guide*. The *IBM Informix Guide to SQL: Reference* provides reference information about the system catalog, the built-in SQL data types, and environment variables that can affect SQL statements. The *IBM Informix Guide to SQL: Tutorial* shows how to use basic and advanced SQL and SPL routines to access and manipulate the data in your databases. The *IBM Informix Database Design and Implementation Guide* shows how to use SQL to implement and manage relational databases.

Types of Users

This publication is written for the following users:

- Database users
- Database administrators
- Database-application programmers

This publication assumes that you have the following background:

- A working knowledge of your computer, your operating system, and the utilities that your operating system provides
- Some experience working with relational databases or exposure to database concepts
- Some experience with computer programming

If you have limited experience with relational databases, SQL, or your operating system, refer to the *IBM Informix Dynamic Server Getting Started Guide* for your database server for a list of supplementary titles.

Software Dependencies

This publication assumes that you are using one of the following database servers:

- IBM Informix Extended Parallel Server, Version 8.50
- IBM Informix Dynamic Server, Version 11.50

Assumptions About Your Locale

IBM Informix products can support many languages, cultures, and code sets. All culture-specific information is brought together in a single environment, called a Global Language Support (GLS) locale.

This publication assumes that you use the U.S. 8859-1 English locale as the default locale. The default is **en_us.8859-1** (ISO 8859-1) on UNIX[®] platforms or **en_us.1252** (Microsoft[®] 1252) for Windows[®] environments. These locales support U.S. English format conventions for dates, times, and currency, and also support the ISO 8859-1 or Microsoft 1252 code set, which includes the ASCII code set plus many 8-bit characters such as è, é, and ñ.

If you plan to use non-ASCII characters in your data or in SQL identifiers, or if you want to conform to localized collation rules of character data, you need to specify an appropriate nondefault locale.

For instructions on how to specify a nondefault locale, additional syntax, and other considerations related to GLS locales, see the *IBM Informix GLS User's Guide*.

Demonstration Databases

The DB–Access utility, which is provided with your IBM Informix database server products, includes one or more of the following demonstration databases:

- The **stores_demo** database illustrates a relational schema with information about a fictitious wholesale sporting-goods distributor. Many examples in IBM Informix publications are based on the **stores_demo** database.
- For Extended Parallel Server only, the **sales_demo** database illustrates a dimensional schema for data-warehousing applications. For conceptual information about dimensional data modeling, see the *IBM Informix Database Design and Implementation Guide*.
- For Dynamic Server only, the **superstores_demo** database illustrates an object-relational schema that contains examples of extended data types, data-type inheritance and table inheritance, and user-defined routines.

For information about how to create and populate the demonstration databases, see the *IBM Informix DB–Access User's Guide*. For descriptions of the databases and their contents, see the *IBM Informix Guide to SQL: Reference*.

The scripts that you use to install the demonstration databases reside in the **\$INFORMIXDIR/bin** directory on UNIX platforms and in the **%INFORMIXDIR%\bin** directory in Windows environments.

What's New in SQL Syntax for Dynamic Server, Version 11.50

For a comprehensive list of new features for this release, see the *IBM Informix Dynamic Server Getting Started Guide*. The following changes and enhancements are relevant to this publication.

| *Table 1. What's New in IBM Informix Guide to SQL: Syntax for Version 11.50.xC3.*

Overview	Reference
<p> Rolling Back SQL Transactions to a Savepoint</p> <p> You can now declare or reference savepoint objects in SQL statements. A savepoint identifies an arbitrary location within the statements of an SQL transaction. Within its transaction, the savepoint resembles a statement label to which the ROLLBACK statement of SQL can cancel any changes to the database that the transaction produced between the savepoint and the ROLLBACK statement. A client application that declares one or more savepoints within a transaction can implement error-handling logic that rolls back only the portion of the transaction that follows the specified savepoint, rather than cancelling the entire transaction if an exception occurs.</p> <p> The following new SQL statements are implemented for savepoints</p> <ul style="list-style-type: none"> • The SAVEPOINT statement creates a savepoint within the current SQL transaction. • The RELEASE SAVEPOINT statement destroys a specified savepoint, as well as any other savepoints that exist between the RELEASE statement and the savepoint that it references. • The ROLLBACK WORK TO SAVEPOINT statement discards changes to the schema of the database or to its data values by statements that follow the savepoint, but the effects of DDL and DML statements that preceded the savepoint persist. <p> This release of Dynamic Server directly supports the SQL statements that declare or reference savepoint objects in DB-Access and in UDRs written in the SPL, C, or Java languages. Methods that support savepoints are also supported in JDBC applications. Applications written in ESQL/C that use the savepoint statements of SQL are also supported with this release of Client SDK.</p>	<p> "TO SAVEPOINT Clause" on page 2-521</p>
<p> Changing the Size of the First Extent of a Table</p> <p> When you create a table, you specify a first extent size based on the eventual estimated size of the table. If the table becomes larger or smaller than that estimate, you might want to change the size of the first extent to avoid either having too many extents or creating extents that are larger than necessary. You can now change the size of the first extent of a table in a dbspace. When you change the size of the first extent, Dynamic Server records the change (in the system catalog and on the partition page), but only makes the actual change when the table is rebuilt or a new partition or fragment is created.</p>	<p> "MODIFY EXTENT SIZE" on page 2-66</p>

| *Table 1. What's New in IBM Informix Guide to SQL: Syntax for Version 11.50.xC3. (continued)*

Overview	Reference
Using SQL Admin API to Dynamically Update Configuration Parameters You can dynamically set configuration parameters by using the new SQL Admin API SET ONCONFIG commands. These new commands provide the same function as the onmode -wf or onmode -wm commands. Use the SET ONCONFIG or SET ONCONFIG PERMANENT commands to change configuration parameters in the ONCONFIG file. Use the SET ONCONFIG MEMORY command to configure parameters only for the current session.	"SET ONCONFIG SQL Admin API Commands" on page 6-23
Improved SQL Tracing with the SQL Admin API You can use these new commands to manage SQL tracing by databases, sessions, and users: SET SQL TRACING DATABASE, SET SQL TRACING SESSION, and SET SQL TRACING USER. Previously, you could only trace SQL for all databases at the server. Now you can control which databases to include in the SQL trace. You can also turn tracing on or off for a specific session, and specify whether you want to trace SQL statements run by specific users. You can also suspend and resume all tracing at the server, without deallocating any resources, by using the SET SQL TRACING SUSPEND and RESUME commands.	"SET SQL TRACING SQL Admin API Command" on page 6-24, "SET SQL TRACING DATABASE SQL Admin API Commands" on page 6-25, "SET SQL TRACING SESSION SQL Admin API Command" on page 6-26, and "SET SQL TRACING USER SQL Admin API Commands" on page 6-27

| *Table 2. What's New in IBM Informix Guide to SQL: Syntax for Version 11.50.xC2.*

Overview	Reference
Subquery Support in UPDATE and DELETE Statements The FROM clause of a subquery in the WHERE clause of the DELETE or UPDATE statement can specify as a data source the same table or view that the FROM clause of the DELETE or UPDATE statement specifies. In this version, Dynamic Server supports DELETE and UPDATE operations with subqueries that reference the same table object if all of the following conditions are true: <ul style="list-style-type: none"> • The subquery either returns only a single row, or else has no references to columns in any table that is not listed in its FROM clause. • The subquery is specified in the WHERE clause of the outer DELETE or UPDATE statement, using the Condition with Subquery syntax. If you call SPL routines within the same subquery, these cannot reference the table that is being modified.	"Subqueries in the WHERE Clause of UPDATE" on page 2-700 and "Subqueries in the WHERE Clause of DELETE" on page 2-306

Table 2. What's New in IBM Informix Guide to SQL: Syntax for Version 11.50.xC2. (continued)

Overview	Reference
<p>Longer Return Strings from String Manipulation Functions</p> <p>The following SQL string-manipulation functions now support operations where the returned string can be up to 32 KB in length:</p> <ul style="list-style-type: none"> • LPAD • RPAD • REPLACE • SUBSTR • SUBSTRING • TRIM • LTRIM • RTRIM • ENCRYPT_AES • ENCRYPT_TDES • DECRYPT_CHAR • DECRYPT_BINARY • CONCAT (and the operator) <p>In addition, the CONCAT function and the operator now provide native support for the LVARCHAR data type.</p>	<p>"Return Types from CONCAT and String Functions" on page 4-142</p>
<p>Controlling I/O of B-Tree Indexes with Compression Levels</p> <p>B-tree scanners can now compress indexes by merging two partially used index pages if the data on those pages totals a set level (low, medium, or high). You can specify the index compression level by modifying the value of the compression field of the BTSCANNER configuration parameter option, by running an onmode -C compression value command, or by running an SQL API function with a SET INDEX COMPRESSION command. command.</p>	<p>"Mnemonic Initial Commands" on page 6-6 and "onmode C: Control the B-tree scanner" on page 6-31</p>

Table 3. What's New in IBM Informix Guide to SQL: Syntax for Version 11.50.xC1.

Overview	Reference
<p>Data Types BIGINT and BIGSERIAL</p> <p>This release introduces two ANSI/ISO-standard exact numeric data types, BIGINT and BIGSERIAL. These data types have the same ranges as the existing INT8 and SERIAL8 data types, but have storage and computational advantages.</p>	<p>"Exact Numeric Data Types" on page 4-26</p>

Table 3. What's New in IBM Informix Guide to SQL: Syntax for Version 11.50.xC1. (continued)

Overview	Reference
<p>SQL Expressions with the IS [NOT] NULL Predicate</p> <p>Now you can use SQL expressions as operands in Boolean conditions that use the IS NULL or IS NOT NULL predicate. Previously, only column names were valid as operands of IS NULL and IS NOT NULL, which return a TRUE or FALSE result. The use of IS NULL or IS NOT NULL with expressions enables you to provide a value for entries that otherwise are not computable (because NULL is not a valid numeric value).</p>	<p>"IS NULL and IS NOT NULL Conditions" on page 4-12</p>
<p>SQL Admin API Commands to Configure High-Availability Clusters</p> <p>New commands for the SQL Admin API enable user informix (or members of the DBSA group who hold Connect privilege on the sysadmin database and on the new sysha database) to configure high-availability data replication cluster environments remotely.</p> <p>These SQL function calls emulate the effects of various onmode -d and onmode -wk command-line options.</p>	<p>"High Availability Initial Commands" on page 6-44</p> <p><i>IBM Informix Dynamic Server Administrator's Guide</i></p> <p><i>IBM Informix Administrator's Reference</i></p>
<p>Control External Directives for the Current Session</p> <p>You now can use the new EXTDIRECTIVES session environment option of the SET ENVIRONMENT statement to control whether external directives are enabled, disabled, or have default behavior during a session. You can specify the default directives behavior with the EXT_DIRECTIVES configuration parameter and the client-side IFX_EXTDIRECTIVES environment variable.</p>	<p>"SET ENVIRONMENT " on page 2-615</p> <p>"EXTDIRECTIVES Environment Option (IDS)" on page 2-620</p>

Table 3. What's New in IBM Informix Guide to SQL: Syntax for Version 11.50.xC1. (continued)

Overview	Reference
<p>Dynamic SQL in SPL Routines</p> <p>You can now use the following dynamic SQL statements in SPL routines:</p> <ul style="list-style-type: none"> • EXECUTE IMMEDIATE, to prepare, execute and free an SQL statement whose text is specified by a quoted string or by a character expression. • PREPARE, to define a prepared object from the text of an SQL query or function call. • DECLARE, to associate the prepared object with a Select cursor or with a Function cursor. • OPEN, to activate a cursor and to replace any question mark (?) placeholders in the text of the prepared object with values in the USING clause of the OPEN statement. • FETCH, to sequentially process the rows of the active set of a cursor. • CLOSE, to deactivate a cursor after all of its rows have been processed. • FREE, to deallocate the resources of a cursor that is no longer needed. <p>For some of these dynamic SQL statements, their syntax in SPL routines is a subset of their Informix ESQL/C syntax, but the PREPARE and EXECUTE IMMEDIATE statements now support character expressions in SPL routines.</p> <p>In addition, the new SQLCODE built-in function returns the value of sqlca.sqlcode for the most recently executed SQL statement. This function expression, which can be invoked only from SPL routines, is useful in error handling and in program logic to exit from a loop after the last row of the active set of a cursor has been processed.</p>	<p>"CLOSE " on page 2-77</p> <p>"Declaring a Dynamic Cursor in an SPL Routine" on page 2-301</p> <p>"EXECUTE IMMEDIATE" on page 2-360</p> <p>"Fetching from Dynamic Cursors in SPL Routines" on page 2-376</p> <p>"FREE " on page 2-379</p> <p>"OPEN " on page 2-459</p> <p>"PREPARE " on page 2-468</p> <p>"SQLCODE Function (SPL)" on page 4-98</p>

Table 3. What's New in IBM Informix Guide to SQL: Syntax for Version 11.50.xC1. (continued)

Overview	Reference
<p data-bbox="423 285 927 342">Automated Table Statistics Maintenance System</p> <p data-bbox="423 369 927 541">Auto Update Statistics (AUS) is an UPDATE STATISTICS maintenance system that can automate the potentially complex DBA task of updating table statistics. This AUS system is implemented as two tasks of the Scheduler:</p> <ol data-bbox="423 552 927 1108" style="list-style-type: none"> <li data-bbox="423 552 927 867">1. The first AUS task, Auto Update Statistics Evaluation, analyzes database logs to evaluate whether or not each table needs its statistics updated. Built-in policies specify criteria for identifying the tables whose column distributions should be recalculated, but you can reset AUS parameter values to modify these default policies. Output from this task is a prioritized set of UPDATE STATISTICS statements that AUS generates. <li data-bbox="423 877 927 1108">2. The second AUS task, Auto Update Statistics Refresh, recalculates table statistics for use by the query optimizer. Statements from the Auto Update Statistics Evaluation task are executed in their order of priority during a recurring time interval that the DBA specifies for this task. <p data-bbox="423 1129 927 1360">The AUS system can simplify the work of the DBA in maintaining column distributions in the system catalog for tables that DML statements have modified. This feature can improve the performance of the database server by regularly providing the query optimizer with updated table statistics as the basis for efficient query plans.</p>	<p data-bbox="927 285 1427 342">"Automated Table Statistics Maintenance" on page 2-706</p>

Documentation Conventions

This section describes the following conventions, which are used in the product documentation for IBM Informix Dynamic Server:

- Typographical conventions
- Feature, product, and platform conventions
- Syntax diagrams
- Command-line conventions
- Example code conventions

Typographical Conventions

This publication uses the following conventions to introduce new terms, illustrate screen displays, describe command syntax, and so forth.

Convention	Meaning
KEYWORD	Keywords of SQL, SPL, and some other programming languages appear in uppercase letters in a serif font.
<i>italics</i>	Within text, new terms and emphasized words appear in italics. Within syntax and code examples, variable values that you are to specify appear in italics.
boldface	Names of program entities (such as classes, events, and tables), environment variables, file names, path names, and interface elements (such as icons, menu items, and buttons) appear in boldface.
monospace	Information that the product displays and information that you enter appear in a monospace typeface.
KEYSTROKE	Keys that you are to press appear in uppercase letters in a sans serif font.
>	This symbol indicates a menu item. For example, “Choose Tools > Options ” means choose the Options item from the Tools menu.

Technical changes to the text are indicated by special characters depending on the format of the documentation:

HTML documentation

New or changed information is surrounded by blue \gg and \ll characters.

PDF documentation

A plus sign (+) is shown to the left of the current changes. A vertical bar (|) is shown to the left of changes made in earlier shipments.

Feature, Product, and Platform Markup

Feature, product, and platform markup identifies paragraphs that contain feature-specific, product-specific, or platform-specific information. Some examples of this markup follow:

Dynamic Server only: Identifies information that is specific to the Windows operating system

Windows only: Identifies information that is specific to the Windows operating system

This markup can apply to one or more paragraphs within a section. When an entire section applies to a particular product or platform, this is noted as part of the heading text, for example:

Table Sorting (Windows)

Example Code Conventions

Examples of SQL code occur throughout this publication. Except as noted, the code is not specific to any single IBM Informix application development tool.

If only SQL statements are listed in the example, they are not delimited by semicolons. For instance, you might see the code in the following example:

```
CONNECT TO stores_demo
...

DELETE FROM customer
WHERE customer_num = 121
```

...

```
COMMIT WORK  
DISCONNECT CURRENT
```

To use this SQL code for a specific product, you must apply the syntax rules for that product. For example, if you are using an SQL API, you must use EXEC SQL at the start of each statement and a semicolon (or other appropriate delimiter) at the end of the statement. If you are using DB–Access, you must delimit multiple statements with semicolons.

Tip: Ellipsis points in a code example indicate that more code would be added in a full application, but it is not necessary to show it to describe the concept being discussed.

For detailed directions on using SQL statements for a particular application development tool or SQL API, see the documentation for your product.

Additional Documentation

You can view, search, and print all of the product documentation from the IBM Informix Dynamic Server information center on the Web at <http://publib.boulder.ibm.com/infocenter/idshelp/v115/index.jsp>.

For additional documentation about IBM Informix Dynamic Server and related products, including release notes, machine notes, and documentation notes, go to the online product library page at <http://www.ibm.com/software/data/informix/pubs/library/>. Alternatively, you can access or install the product documentation from the Quick Start CD that is shipped with the product.

Compliance with Industry Standards

The American National Standards Institute (ANSI) and the International Organization of Standardization (ISO) have jointly established a set of industry standards for the Structured Query Language (SQL). IBM Informix SQL-based products are fully compliant with SQL-92 Entry Level (published as ANSI X3.135-1992), which is identical to ISO 9075:1992. In addition, many features of IBM Informix database servers comply with the SQL-92 Intermediate and Full Level and X/Open SQL Common Applications Environment (CAE) standards.

Syntax Diagrams

This guide uses syntax diagrams built with the following components to describe the syntax for statements and all commands other than system-level commands.

Table 4. Syntax Diagram Components

Component represented in PDF	Component represented in HTML	Meaning
	<code>>>-----</code>	Statement begins.
	<code>-----></code>	Statement continues on next line.
	<code>>-----</code>	Statement continues from previous line.

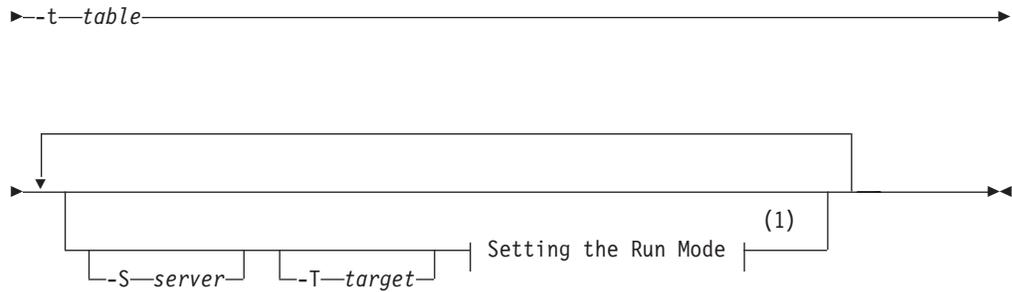
Table 4. Syntax Diagram Components (continued)

Component represented in PDF	Component represented in HTML	Meaning
	----->>	Statement ends.
	-----SELECT-----	Required item.
	---+-----+--- '-----LOCAL-----'	Optional item.
	---+-----+--- +--DISTINCT-----+ '---UNIQUE-----'	Required item with choice. One and only one item must be present.
	---+-----+--- +--FOR UPDATE-----+ '--FOR READ ONLY--'	Optional items with choice are shown below the main line, one of which you might specify.
	.---NEXT-----. ---+-----+--- +---PRIOR-----+ '---PREVIOUS-----'	The values below the main line are optional, one of which you might specify. If you do not specify an item, the value above the line will be used as the default.
	-----,----- v ---+-----+--- +---index_name---+ '---table_name---	Optional items. Several items are allowed; a comma must precede each repetition.
	>> Table Reference -><	Reference to a syntax segment.
Table Reference 	Table Reference ---+-----+--- +-----table-----+ '-----synonym-----'	Syntax segment.

How to Read a Command-Line Syntax Diagram

The following command-line syntax diagram uses some of the elements listed in the table in Syntax Diagrams.

Creating a No-Conversion Job

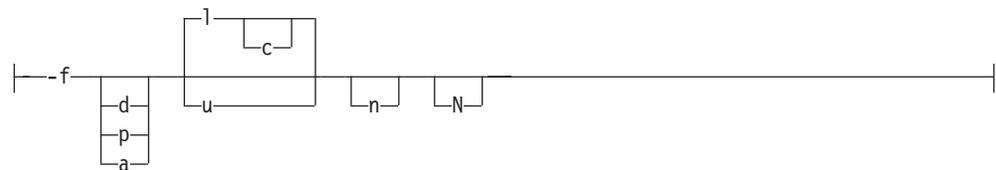


Notes:

1 See page Z-1

The second line in this diagram has a segment named “Setting the Run Mode,” which according to the diagram footnote, is on page Z-1. If this was an actual cross-reference, you would find this segment in on the first page of Appendix Z. Instead, this segment is shown in the following segment diagram. Notice that the diagram uses segment start and end components.

Setting the Run Mode:



To see how to construct a command correctly, start at the top left of the main diagram. Follow the diagram to the right, including the elements that you want. The elements in this diagram are case sensitive because they illustrate utility syntax. Other types of syntax, such as SQL, are not case sensitive.

The Creating a No-Conversion Job diagram illustrates the following steps:

1. Type **onpladm create job** and then the name of the job.
2. Optionally, type **-p** and then the name of the project.
3. Type the following required elements:
 - **-n**
 - **-d** and the name of the device
 - **-D** and the name of the database
 - **-t** and the name of the table
4. Optionally, you can choose one or more of the following elements and repeat them an arbitrary number of times:
 - **-S** and the server name
 - **-T** and the target server name
 - The run mode. To set the run mode, follow the Setting the Run Mode segment diagram to type **-f**, optionally type **d**, **p**, or **a**, and then optionally type **l** or **u**.
5. Follow the diagram to the terminator.

Keywords and Punctuation

Keywords are words reserved for statements and all commands except system-level commands. When a keyword appears in a syntax diagram, it is shown in uppercase letters. When you use a keyword in a command, you can write it in uppercase or lowercase letters, but you must spell the keyword exactly as it appears in the syntax diagram.

You must also use any punctuation in your statements and commands exactly as shown in the syntax diagrams.

Identifiers and Names

Variables serve as placeholders for identifiers and names in the syntax diagrams and examples. You can replace a variable with an arbitrary name, identifier, or literal, depending on the context. Variables are also used to represent complex syntax elements that are expanded in additional syntax diagrams. When a variable appears in a syntax diagram, an example, or text, it is shown in *lowercase italic*.

The following syntax diagram uses variables to illustrate the general form of a simple SELECT statement.

►—SELECT—*column_name*—FROM—*table_name*—◄◄

When you write a SELECT statement of this form, you replace the variables *column_name* and *table_name* with the name of a specific column and table.

How to Provide Documentation Feedback

You are encouraged to send your comments about IBM Informix user documentation by using one of the following methods:

- Send e-mail to docinf@us.ibm.com.
- Go to the Information Center at <http://publib.boulder.ibm.com/infocenter/idshelp/v115/index.jsp> and open the topic that you want to comment on. Click the feedback link at the bottom of the page, fill out the form, and submit your feedback.

Feedback from both methods is monitored by those who maintain the user documentation of Dynamic Server. The feedback methods are reserved for reporting errors and omissions in our documentation. For immediate help with a technical problem, contact IBM Technical Support. For instructions, see the IBM Informix Technical Support Web site at <http://www.ibm.com/planetwide/>.

We appreciate your suggestions.

Chapter 1. Overview of SQL Syntax

In This Chapter

This chapter provides information about how to use the SQL statements, SPL statements, and syntax segments that subsequent chapters of this book discuss. The chapter is organized into the following sections.

Section	Scope
"How to Enter SQL Statements"	How to use syntax diagrams and descriptions to enter SQL statements correctly
"How to Enter SQL Comments" on page 1-3	How to enter comments in SQL statements
"Categories of SQL Statements" on page 1-5	The SQL statements, listed by functional category
"ANSI/ISO Compliance and Extensions" on page 1-10	The SQL statements, listed by degree of ANSI/ISO compliance

How to Enter SQL Statements

SQL is a free-form language, like C or PASCAL, that generally ignores white space characters like TAB, LINEFEED, and extra blank spaces between statements or statement elements. At least one blank character or other delimiter, however, must separate keywords and identifiers from other syntax tokens.

SQL is lettercase insensitive, except within quoted strings; see also "Identifier " on page 5-21. In an ANSI-compliant database, if you do not delimit the *owner* of an object by double (") quotation marks, and the **ANSIOWNER** environment variable was not set to 1 when the database server was initialized, the database server stores the *owner* name in uppercase letters.

Statement descriptions are provided in this publication to help you to enter SQL statements successfully. A statement description includes this information:

- A brief introduction that explains what the statement does
- A syntax diagram that shows how to enter the statement correctly
- A syntax table that explains each input parameter in the syntax diagram
- Rules of usage, typically with examples that illustrate these rules

For some statements, this information is provided for individual clauses.

Most statement descriptions conclude with references to related information in this publication and in other publications.

Chapter 2, "SQL Statements," on page 2-1 provides descriptions of each SQL statement, arranged in alphabetical order. Chapter 3, "SPL Statements," on page 3-1 describes each of the SPL statements, using the same format.

The major aids for entering SQL statements include:

- The combination of the syntax diagram and syntax table
- The examples of syntax that appear in the rules of usage

- The references to related information

Using Syntax Diagrams and Syntax Tables

Before you try to use the syntax diagrams in this chapter, it is helpful to read the syntax diagram section of the Introduction. This section is the key to understanding the syntax diagrams and explains the elements that can appear in a syntax diagram and the paths that connect the elements to each other. This section also includes an example that illustrates the elements of typical syntax diagrams. The narrative that follows the example diagram shows how to read the diagram in order to enter the statement successfully.

Notes to the diagram can reference other syntax segments or can specify various restrictions. If you are using an application programming interface, such as ESQL/C or 4GL, only the SQL syntax rules that both your client application and the database server support are valid.

When a syntax diagram includes input specifications that are not keywords, syntax segments, nor punctuation symbols, such as identifiers, expressions, filenames, or host variables, the syntax diagram is followed by a table that describes how to enter the term without generating errors. Each syntax table includes four columns:

- The **Element** column lists each variable term in the syntax diagram.
- The **Description** column briefly describes the term and identifies the default value, if the term has one.
- The **Restrictions** column summarizes the restrictions on the term, such as acceptable ranges of values. (For some diagrams, restrictions that cannot be tersely summarized appear in the **Usage** notes, rather than in this column.)
- The **Syntax** column points to the SQL segment that gives the detailed syntax for the term. For a few terms, such as the names of host variables, pathnames, or literal characters, no page reference is provided.

The diagrams generally provide an intuitive notation for what is valid in a given SQL statement, but for some statements, dependencies or restrictions among syntax elements are identified only in the text of the Usage section.

Using Examples

To understand the main syntax diagram and subdiagrams for a statement, study the examples of syntax that appear in the rules of usage for each statement. These examples have two purposes:

- To show how to accomplish specific tasks with the statement or its clauses
- To show how to use syntax of the statement or its clauses in a concrete way

Tip: An efficient way to understand a syntax diagram is to find an example of the syntax and compare it with the keywords and parameters in the syntax diagram. By mapping the concrete elements of the example to the abstract elements of the syntax diagram, you can understand the syntax diagram and use it more effectively.

For an explanation of the conventions used in the examples in this publication, see the syntax diagram section of the Introduction.

These code examples are program fragments to illustrate valid syntax, rather than complete SQL programs. In some code examples, ellipsis (. . .) symbols indicate

that additional code has been omitted. To save space, however, ellipses are not shown at the beginning or end of the program fragments.

Using Related Information

For help in understanding concepts and terms in the SQL statement description, check the “Related Information” section at the end of each statement.

This section points to related information in this publication and other publications to help you understand the statement in question. The section provides some or all of the following information:

- The names of related statements that might contain a fuller discussion of topics in this statement
- The titles of other publications that provide extended discussions of topics in this statement

Tip: If you do not have extensive knowledge and experience with SQL, the *IBM Informix Guide to SQL: Tutorial* gives you the basic SQL knowledge that you need to understand and use the statement descriptions in this document.

How to Enter SQL Comments

You can add comments to clarify the purpose or effect of particular SQL statements. You can also use comment symbols during program development to disable individual statements without deleting them from your source code.

Your comments can help you or others to understand the role of the statement within a program, SPL routine, or command file. The code examples in this document sometimes include comments that clarify the role of an SQL statement within the code, but your own SQL programs will be easier to read and to maintain if you document them with frequent comments.

The following table shows the SQL comment indicators that you can enter in your code. Here a Y in a column signifies that you can use the symbol with the product or with the type of database identified in the column heading. An N in a column signifies that you cannot use the symbol with the indicated product or with a database of the indicated ANSI-compliance status.

Comment Symbol	ESQL/C	SPL Routine	DB-Access	ANSI-Compliant Databases	Databases Not ANSI Compliant	Description
double hyphen (--)	Y	Y	Y	Y	Y	The double hyphen precedes a comment within a single line. To comment more than one line, put double hyphen symbols at the beginning of each comment line.
braces ({ . . . })	N	Y	Y	Y	Y	Braces enclose the comment. The { precedes the comment, and the } follows it. Braces can delimit single- or multiple-line comments, but comments cannot be nested.

Comment Symbol	ESQL/C	SPL Routine	DB–Access	ANSI-Compliant Databases	Databases Not ANSI Compliant	Description
slash and asterisk /* . . . */	Y	Y	Y	Y	Y	C-language style slash and asterisk (/* */) paired delimiters enclose the comment. The /* precedes the comment, and the */ follows it. These can delimit single-line or multiple-line comments, but comments cannot be nested.

Characters within the comment are ignored by the database server.

The section “ Optimizer Directives ” on page 5-35 describes a context where information within comments can influence query plans of Dynamic Server.

If the product that you use supports all of these comment symbols, your choice of a comment symbol depends on requirements for ANSI/ISO compliance:

- Double hyphen (--) complies with the ANSI/ISO standard for SQL.
- Braces ({ }) are an Informix extension to the ANSI/ISO standard.
- C-style slash-and-asterisk (/* . . . */) comply with the SQL-99 standard.

If ANSI/ISO compliance is not an issue, your choice of comment symbols is a matter of personal preference.

In DB–Access, you can use any of these comment symbols when you enter SQL statements with the SQL editor and when you create SQL command files with the SQL editor or with a system editor.

An SQL command file is an operating-system file that contains one or more SQL statements. Command files are also known as command scripts. For more information about command files, see the discussion of command scripts in the *IBM Informix Guide to SQL: Tutorial*. For information on how to create and modify command files with the SQL editor or a system editor in DB–Access, see the *IBM Informix DB–Access User’s Guide*.

You can use any of these comment symbols in any line of an SPL routine. See the discussion of how to comment and document an SPL routine in the *IBM Informix Guide to SQL: Tutorial*.

In Informix ESQL/C, the double hyphen (--) can begin a comment that extends to the end of the same line. For information on language-specific comment symbols in Informix ESQL/C programs, see the *IBM Informix ESQL/C Programmer’s Manual*.

Examples of SQL Comments

These examples illustrate different ways to use the SQL comment indicators.

The following examples use each style of comment indicator, including the double hyphen (--), braces ({ }), and C-style (/* . . . */) comment delimiters to include a comment after an SQL statement. The comment appears on the same line as the statement.

```
SELECT * FROM customer; -- Selects all columns and rows
```

```
SELECT * FROM customer; {Selects all columns and rows}
```

```
SELECT * FROM customer; /*Selects all columns and rows*/
```

The next three examples use the same SQL statement and the same comments as in the preceding examples, but place the comment on a separate line:

```
SELECT * FROM customer;
  -- Selects all columns and rows
```

```
SELECT * FROM customer;
  {Selects all columns and rows}
```

```
SELECT * FROM customer;
  /*Selects all columns and rows*/
```

In the following examples, the user enters the same SQL statement as in the preceding example but now a multiple-line comment (or for the double-hyphen indicator, two comments) follows each statement:

```
SELECT * FROM customer;
  -- Selects all columns and rows
  -- from the customer table
```

```
SELECT * FROM customer;
  {Selects all columns and rows
  from the customer table}
```

```
SELECT * FROM customer;
  /*Selects all columns and rows
  from the customer table*/
```

Comments in any of these styles can also appear within an SQL statement:

```
SELECT *          -- Selects all columns and rows
  FROM customer;  -- from the customer table
```

```
SELECT *          {Selects all columns and rows}
  FROM customer;  {from the customer table}
```

```
SELECT *          /*Selects all columns and rows*/
  FROM customer;  /*from the customer table*/
```

If you use braces or C-style comments that are delimited by paired opening and closing indicators, the closing comment indicator must be in the same style as the opening comment indicator.

Non-ASCII Characters in SQL Comments

You can enter non-ASCII characters (including multibyte characters) in SQL comments if the database locale supports the non-ASCII characters. For further information on the GLS aspects of SQL comments, see the *IBM Informix GLS User's Guide*.

Categories of SQL Statements

SQL statements are traditionally divided into the following logical categories:

- **Data definition statements.** These *data definition language* (DDL) statements can declare, rename, modify, or destroy objects in the local database.
- **Data manipulation statements.** These *data manipulation language* (DML) statements can retrieve, insert, delete, or modify data values.
- **Cursor manipulation statements.** These statements can declare, open, and close *cursors*, which are data structures for operations on multiple rows of data.
- **Dynamic management statements.** These statements support memory management and allow users to specify at runtime the details of DML operations.
- **Data access statements.** These statements specify discretionary access privileges and support concurrent access to the database by multiple users.
- **Data integrity statements.** These implement transaction logging and support the referential integrity of the database.
- **Optimization statements.** These can be used to improve the performance of operations on the database.
- **Routine definition statements.** These can declare, define, modify, execute, or destroy user-defined routines that the database stores.
- **Client/server connection statements.** These can open or close a connection between a database and a client application.
- **Auxiliary statements.** These can provide information about the database. (This is also a residual category for statements that are not closely related to the other statement categories.)
- **Optical subsystem statements.** These statements are separately documented in *IBM Informix Optical Subsystem Guide*.

The SQL statements of each category are listed in the pages that follow.

As Chapter 2, "SQL Statements," on page 2-1 indicates, some statements (and some statement options, as noted in the syntax diagrams) are specific to Dynamic Server (sometimes abbreviated as "IDS") or are specific to Extended Parallel Server (sometimes abbreviated as "XPS").

Data Definition Language Statements

- ALTER ACCESS_METHOD
- ALTER FRAGMENT
- ALTER FUNCTION
- ALTER INDEX
- ALTER PROCEDURE
- ALTER ROUTINE
- ALTER SEQUENCE
- ALTER SECURITY LABEL COMPONENT
- ALTER TABLE
- CLOSE DATABASE
- CREATE ACCESS_METHOD
- CREATE AGGREGATE
- CREATE CAST
- CREATE DATABASE
- CREATE DISTINCT TYPE
- CREATE DUPLICATE

- CREATE EXTERNAL TABLE
- CREATE FUNCTION
- CREATE FUNCTION FROM
- CREATE INDEX
- CREATE OPAQUE TYPE
- CREATE OPCLASS
- CREATE PROCEDURE
- CREATE PROCEDURE FROM
- CREATE ROLE
- CREATE ROUTINE FROM
- CREATE ROW TYPE
- CREATE SCHEMA
- CREATE SECURITY LABEL
- CREATE SECURITY LABEL COMPONENT
- CREATE SECURITY POLICY
- CREATE SEQUENCE
- CREATE SYNONYM
- CREATE TABLE
- CREATE Temporary TABLE
- CREATE TRIGGER
- CREATE VIEW
- CREATE XADATASOURCE
- CREATE XADATASOURCE TYPE
- DROP ACCESS_METHOD
- DROP AGGREGATE
- DROP CAST
- DROP DATABASE
- DROP DUPLICATE
- DROP FUNCTION
- DROP INDEX
- DROP OPCLASS
- DROP PROCEDURE
- DROP ROLE
- DROP ROUTINE
- DROP ROW TYPE
- DROP SECURITY
- DROP SEQUENCE
- DROP SYNONYM
- DROP TABLE
- DROP TRIGGER
- DROP TYPE
- DROP VIEW
- DROP XADATASOURCE
- DROP XADATASOURCE TYPE
- MOVE

- RENAME COLUMN
- RENAME DATABASE
- RENAME INDEX
- RENAME SECURITY
- RENAME SEQUENCE
- RENAME TABLE
- TRUNCATE
- UPDATE STATISTICS

Data Manipulation Language Statements

- DELETE
- INSERT
- LOAD
- MERGE
- SELECT
- UNLOAD
- UPDATE

Note: The DELETE, INSERT, SELECT, and UPDATE statements are DML statements in the ANSI/ISO standard for SQL. Although LOAD, UNLOAD, and MERGE resemble DML in their functionality, these are out-of-scope for most references in this document to “DML statements.”

Data Integrity Statements

- BEGIN WORK
- COMMIT WORK
- SAVEPOINT
- RELEASE SAVEPOINT
- ROLLBACK WORK
- SET Database Object Mode
- SET LOG
- SET PLOAD FILE
- SET Transaction Mode
- START VIOLATIONS TABLE
- STOP VIOLATIONS TABLE

Cursor Manipulation Statements

- CLOSE
- DECLARE
- FETCH
- FLUSH
- FREE
- OPEN
- PUT
- SET AUTOFREE

Dynamic Management Statements

- ALLOCATE COLLECTION
- ALLOCATE DESCRIPTOR
- ALLOCATE ROW
- DEALLOCATE COLLECTION
- DEALLOCATE DESCRIPTOR
- DEALLOCATE ROW
- DESCRIBE
- DESCRIBE INPUT
- EXECUTE
- EXECUTE IMMEDIATE
- FREE
- GET DESCRIPTOR
- INFO
- PREPARE
- SET DEFERRED_PREPARE
- SET DESCRIPTOR

Data Access Statements

- GRANT
- GRANT FRAGMENT
- LOCK TABLE
- REVOKE
- REVOKE FRAGMENT
- SET ISOLATION
- SET LOCK MODE
- SET ROLE
- SET SESSION AUTHORIZATION
- SET TRANSACTION
- SET Transaction Mode
- UNLOCK TABLE

Optimization Statements

- SAVE EXTERNAL DIRECTIVES
- SET ALL_MUTABLES
- SET Default Table Space
- SET Default Table Type
- SET ENVIRONMENT
- SET EXPLAIN
- SET OPTIMIZATION
- SET PDQPRIORITY
- SET Residency
- SET SCHEDULE LEVEL
- SET STATEMENT CACHE

Routine Definition Statements

- ALTER FUNCTION
- ALTER PROCEDURE
- ALTER ROUTINE
- CREATE FUNCTION
- CREATE FUNCTION FROM
- CREATE PROCEDURE
- CREATE PROCEDURE FROM
- CREATE ROUTINE FROM
- DROP FUNCTION
- DROP PROCEDURE
- DROP ROUTINE
- EXECUTE FUNCTION
- EXECUTE PROCEDURE
- SET DEBUG FILE TO

Auxiliary Statements

- GET DIAGNOSTICS
- INFO
- OUTPUT
- SET COLLATION
- SET DATASKIP
- SET ENCRYPTION PASSWORD
- WHENEVER

Client/Server Connection Statements

- CONNECT
- DATABASE
- DISCONNECT
- SET CONNECTION

Optical Subsystem Statements (IDS)

- ALTER OPTICAL CLUSTER
- CREATE OPTICAL CLUSTER
- DROP OPTICAL CLUSTER
- RELEASE
- RESERVE
- SET MOUNTING TIMEOUT

Important: See the *IBM Informix Optical Subsystem Guide* for more information.

ANSI/ISO Compliance and Extensions

Lists that follow show statements that match the ANSI SQL-92 standard at the entry level, statements that are ANSI compliant but include Informix extensions, and statements that are Informix extensions to the ANSI/ISO standard.

ANSI/ISO-Compliant Statements

- CLOSE
- COMMIT WORK
- RELEASE SAVEPOINT
- SET SESSION AUTHORIZATION
- SET TRANSACTION

ANSI/ISO-Compliant Statements with Informix Extensions

- ALLOCATE DESCRIPTOR
- ALTER TABLE
- CONNECT
- CREATE FUNCTION
- CREATE PROCEDURE
- CREATE TRIGGER
- CREATE SCHEMA
- CREATE TABLE
- CREATE Temporary TABLE
- CREATE VIEW
- DEALLOCATE DESCRIPTOR
- DECLARE
- DELETE
- DESCRIBE
- DESCRIBE INPUT
- DISCONNECT
- EXECUTE
- EXECUTE IMMEDIATE
- FETCH
- GET DESCRIPTOR
- GET DIAGNOSTICS
- GRANT
- INSERT
- MERGE (XPS only)
- OPEN
- PREPARE
- REVOKE
- ROLLBACK WORK
- SAVEPOINT
- SELECT
- SET CONNECTION
- SET CONSTRAINTS
- SET DESCRIPTOR
- SET Transaction Mode
- UPDATE STATISTICS
- WHENEVER

Statements that are Extensions to the ANSI/ISO Standard

- ALLOCATE COLLECTION
- ALLOCATE ROW
- ALTER ACCESS_METHOD
- ALTER FRAGMENT
- ALTER FUNCTION
- ALTER INDEX
- ALTER OPTICAL CLUSTER
- ALTER PROCEDURE
- ALTER ROUTINE
- ALTER SECURITY LABEL COMPONENT
- ALTER SEQUENCE
- BEGIN WORK
- CLOSE DATABASE
- CREATE ACCESS_METHOD
- CREATE AGGREGATE
- CREATE CAST
- CREATE DATABASE
- CREATE DISTINCT TYPE
- CREATE DUPLICATE (XPS only)
- CREATE EXTERNAL TABLE
- CREATE FUNCTION FROM
- CREATE INDEX
- CREATE OPAQUE TYPE
- CREATE OPCLASS
- CREATE OPTICAL CLUSTER
- CREATE PROCEDURE FROM
- CREATE ROLE
- CREATE ROUTINE FROM
- CREATE ROW TYPE
- CREATE SECURITY LABEL
- CREATE SECURITY LABEL COMPONENT
- CREATE SECURITY POLICY
- CREATE SEQUENCE
- CREATE SYNONYM
- CREATE XADATASOURCE
- CREATE XADATASOURCE TYPE
- DATABASE
- DEALLOCATE COLLECTION
- DEALLOCATE ROW
- DROP ACCESS_METHOD
- DROP AGGREGATE
- DROP CAST
- DROP DATABASE
- DROP DUPLICATE

- DROP FUNCTION
- DROP INDEX
- DROP OPCLASS
- DROP OPTICAL CLUSTER
- DROP PROCEDURE
- DROP ROLE
- DROP ROUTINE
- DROP ROW TYPE
- DROP SECURITY LABEL
- DROP SECURITY LABEL COMPONENT
- DROP SECURITY POLICY
- DROP SEQUENCE
- DROP SYNONYM
- DROP TABLE
- DROP TRIGGER
- DROP TYPE
- DROP VIEW
- DROP XADATASOURCE
- DROP XADATASOURCE TYPE
- EXECUTE FUNCTION
- EXECUTE PROCEDURE
- FLUSH
- FREE
- GRANT FRAGMENT
- LOAD
- LOCK TABLE
- MOVE (XPS only)
- OUTPUT
- PUT
- RELEASE
- RENAME COLUMN
- RENAME DATABASE
- RENAME INDEX
- RENAME SECURITY LABEL
- RENAME SECURITY LABEL COMPONENT
- RENAME SECURITY POLICY
- RENAME SEQUENCE
- RENAME TABLE
- RESERVE
- REVOKE FRAGMENT
- SAVE EXTERNAL DIRECTIVES
- SET ALL_MUTABLES (XPS only)
- SET AUTOFREE
- SET COLLATION
- SET Database Object Mode

- SET DATASKIP
- SET DEBUG FILE TO
- SET Default Table Space (XPS only)
- SET Default Table Type (XPS only)
- SET DEFERRED_PREPARE
- SET ENCRYPTION PASSWORD
- SET ENVIRONMENT
- SET EXPLAIN
- SET ISOLATION
- SET LOCK MODE
- SET LOG
- SET MOUNTING TIMEOUT
- SET OPTIMIZATION
- SET PDQPRIORITY
- SET PLOAD FILE (XPS only)
- SET Residency (XPS only)
- SET ROLE
- SET SCHEDULE LEVEL (XPS only)
- SET STATEMENT CACHE
- START VIOLATIONS TABLE
- STOP VIOLATIONS TABLE
- TRUNCATE
- UNLOAD
- UNLOCK TABLE
- UPDATE STATISTICS

Chapter 2. SQL Statements

In This Chapter

This chapter describes the syntax and semantics of SQL statements that are recognized by Dynamic Server or by Extended Parallel Server. Statements (and statement segments, and notes describing usage) that are not explicitly restricted to one of these database servers are valid for both.

The statement descriptions appear in alphabetical order. For some statements, important details of the semantics appear in other volumes of this documentation set, as indicated by cross-references.

For many statements, the syntax diagram, or the table of terms immediately following the diagram, or both, includes references to syntax segments in Chapter 4, "Data Types and Expressions," on page 4-1 or in Chapter 5, "Other Syntax Segments," on page 5-1

When the name of an SQL statement includes lowercase characters, such as "CREATE Temporary TABLE," it means that the first mixed-lettercase string in the statement name is not an SQL keyword, but that two or more different SQL keywords can follow the preceding uppercase keyword.

For an explanation of the structure of statement descriptions, see Chapter 1, "Overview of SQL Syntax," on page 1-1

ALLOCATE COLLECTION

Use the ALLOCATE COLLECTION statement to allocate memory for a variable of a collection data type (such as LIST, MULTISSET, or SET) or an untyped collection variable.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL. Use this statement with ESQL/C.

Syntax

►►—ALLOCATE COLLECTION—*variable*—►►

Element	Description	Restrictions	Syntax
<i>variable</i>	Name of the typed or untyped collection variable to allocate	Must be an unallocated Informix ESQL/C collection-type host variable	Language-specific rules for names

Usage

The ALLOCATE COLLECTION statement allocates memory for an ESQL/C variable that can store the value of a collection data type.

To create a collection variable for an Informix ESQL/C program

1. Declare the collection variable as a client collection variable in an Informix ESQL/C program.
The collection variable can be a typed or untyped collection variable.
2. Allocate memory for the collection variable with the ALLOCATE COLLECTION statement.

The ALLOCATE COLLECTION statement sets **SQLCODE** (that is, **sqlca.sqlcode**) to zero (0) if the memory allocation was successful and to a negative error code if the allocation failed.

When you no longer need the collection variable, you must explicitly release the memory that it occupies with the DEALLOCATE COLLECTION statement. After the DEALLOCATE COLLECTION statement executes successfully, you can reuse the collection variable.

Tip: The ALLOCATE COLLECTION statement allocates memory for an Informix ESQL/C collection variable only. To allocate memory for an Informix ESQL/C row variable, use the ALLOCATE ROW statement.

Examples

The following example shows how to allocate resources with the ALLOCATE COLLECTION statement for the untyped collection variable, **a_set**:

```
EXEC SQL BEGIN DECLARE SECTION;
      client collection a_set;
EXEC SQL END DECLARE SECTION;
. . .
EXEC SQL allocate collection :a_set;
```

The following example uses ALLOCATE COLLECTION to allocate resources for a typed collection variable, **a_typed_set**:

```
EXEC SQL BEGIN DECLARE SECTION;
      client collection set(integer not null) a_typed_set;
EXEC SQL END DECLARE SECTION;
. . .
EXEC SQL allocate collection :a_typed_set;
```

Related Information

Related examples: Refer to the collection-variable example in “PUT ” on page 2-479.

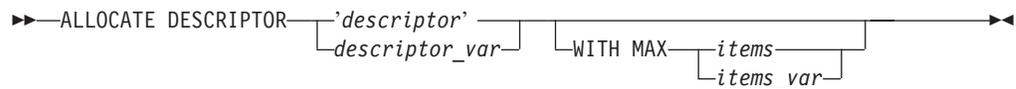
Related statements: “ALLOCATE ROW ” on page 2-4 and “DEALLOCATE COLLECTION ” on page 2-283

For a discussion of collection data types in ESQL/C programs, see the *IBM Informix ESQL/C Programmer's Manual*.

ALLOCATE DESCRIPTOR

Use the ALLOCATE DESCRIPTOR statement to allocate memory for a system-descriptor area (SDA). Use this statement with ESQL/C.

Syntax



Element	Description	Restrictions	Syntax
<i>descriptor</i>	Name of an unallocated system-descriptor area	Enclose in single (') quotation marks. Must be unique among SDA names	"Quoted String" on page 4-189.
<i>descriptor_var</i>	Host variable that stores the name of a system-descriptor area	Must contain name of unallocated system-descriptor area	Language specific
<i>items</i>	Number of item descriptors in <i>descriptor</i> . Default value is 100.	Must be an unsigned INTEGER greater than zero	"Literal Number " on page 4-185
<i>items_var</i>	Host variable that contains the number of items	Data type must be INTEGER or SMALLINT	Language specific

Usage

The ALLOCATE DESCRIPTOR statement creates a *system-descriptor area*, which is a location in memory that holds information that the DESCRIBE statement can display, or that holds information about the WHERE clause of a query.

A system-descriptor area (SDA) contains one or more fields called *item descriptors*. Each item descriptor holds a data value that the database server can receive or send. The item descriptors also contain information about the data, such as data type, length, scale, precision, and support for NULL values.

A system-descriptor area holds information that a DESCRIBE ... USING SQL DESCRIPTOR statement obtains or that holds information about the WHERE clause of a dynamically executed statement.

If the name that you declare for a system-descriptor area matches the name of an existing system-descriptor area, the database server returns an error. After you free the specified descriptor with the DEALLOCATE DESCRIPTOR statement, however, you can reuse the same descriptor name.

WITH MAX Clause

You can use the WITH MAX clause to indicate the maximum number of item descriptors you need. When you use this clause, the COUNT field is set to the number of *items that* you specify. If you do not specify the WITH MAX clause, the default value of the COUNT field is 100. You can change the value of the COUNT field with the SET DESCRIPTOR statement.

The following examples show valid ALLOCATE DESCRIPTOR statements that include the WITH MAX clause. This example uses embedded variable names to identify the system-descriptor area and to specify the desired number of item descriptors:

```
EXEC SQL allocate descriptor :desname with max :occ;
```

The next example uses a quoted string to identify the system-descriptor area and an unsigned integer to specify the desired number of item descriptors:

```
EXEC SQL allocate descriptor 'desc1' with max 3;
```

Related Information

Related statements: "DEALLOCATE DESCRIPTOR " on page 2-284, "DECLARE " on page 2-286, "DESCRIBE" on page 2-311, "EXECUTE " on page 2-348, "FETCH " on page 2-367, "GET DESCRIPTOR " on page 2-381, "OPEN " on page 2-459, "PREPARE " on page 2-468, "PUT " on page 2-479, and "SET DESCRIPTOR " on page 2-605

For more information on system-descriptor areas, refer to the *IBM Informix ESQL/C Programmer's Manual*.

ALLOCATE ROW

Use the ALLOCATE ROW statement to allocate memory for a **row** variable. Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL. Use this statement with ESQL/C.

Syntax

▶▶—ALLOCATE ROW—*variable*—————▶▶

Element	Description	Restrictions	Syntax
<i>variable</i>	Name of a typed or untyped row variable to allocate	Must be an unallocated Informix ESQL/C row-type host variable	Language specific

Usage

The ALLOCATE ROW statement allocates memory for a host variable that stores row-type data. To create a **row** variable, an ESQL/C program must do the following:

1. Declare the **row** variable. The row variable can be a typed or untyped **row** variable.
2. Allocate memory for the **row** variable with the ALLOCATE ROW statement.

The following example shows how to allocate resources with the ALLOCATE ROW statement for the typed **row** variable, **a_row**:

```
EXEC SQL BEGIN DECLARE SECTION;  
    row (a int, b int) a_row;  
EXEC SQL END DECLARE SECTION;  
.  
.  
EXEC SQL allocate row :a_row;
```

The ALLOCATE ROW statement sets **SQLCODE** (the contents of **sqlca.sqlcode**) to zero (0) if the memory allocation operation was successful, or to a negative error code if the allocation failed.

You must explicitly release memory with the DEALLOCATE ROW statement. Once you free the **row** variable with the DEALLOCATE ROW statement, you can reuse the **row** variable.

Tip: The ALLOCATE ROW statement allocates memory for an Informix ESQL/C **row** variable only. To allocate memory for an Informix ESQL/C **collection** variable, use the ALLOCATE COLLECTION statement.

When you use the same **row** variable in multiple function calls without deallocating it, a memory leak on the client computer results. Because there is no way to determine if a pointer is valid when it is passed, Informix ESQL/C assumes that the pointer is not valid and assigns it to a new memory location.

Related Information

Related statements: "ALLOCATE COLLECTION " on page 2-1 and "DEALLOCATE ROW " on page 2-285

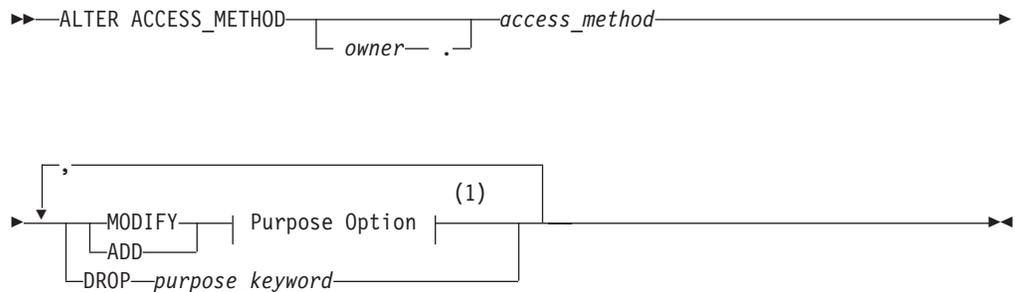
For a discussion of complex data types in ESQL/C programs, see the *IBM Informix ESQL/C Programmer's Manual*.

ALTER ACCESS_METHOD

The ALTER ACCESS_METHOD statement changes the attributes of a user-defined access method in the **sysams** system catalog table.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 See "Purpose Options" on page 5-50

Element	Description	Restrictions	Syntax
<i>access_method</i>	Name of the access method to alter	Access method must be registered in the sysams system catalog table by a previous CREATE ACCESS_METHOD statement	"Identifier " on page 5-21
<i>owner</i>	Name of the owner of the access method	Must own the access method	" Owner Name " on page 5-45
<i>purpose_keyword</i>	A keyword that indicates which feature to change	Keyword must be associated with the access method by a previous CREATE or ALTER ACCESS_METHOD statement	" Purpose Functions, Methods, Flags, and Values" on page 5-52

Usage

Use ALTER ACCESS_METHOD to modify the definition of a user-defined access method. You cannot modify a built-in access method.

You must own the access method or hold DBA privileges to alter a built-in access method. In an ANSI-compliant database, the DBA must qualify the name of the access method with the *owner* name if another user is the owner of the access method.

When you alter an access method, you change the purpose-option specifications (purpose functions, purpose methods, purpose flags, or purpose values) that define the access method. For example, you might alter an access method to assign a new user-defined function or method name or to provide a multiplier for the scan cost on a table.

If a transaction is in progress, the database server waits to alter the access method until after the transaction is committed or rolled back. No other users can execute the access method until the transaction has completed.

Example

The following statement alters the **remote** user-defined access method:

```
ALTER ACCESS_METHOD remote
  ADD am_scancost = FS_scancost,
  ADD am_rowids,
  DROP am_getbyid,
  MODIFY am_costfactor = 0.9;
```

The preceding example makes the following changes to the access method:

- Adds a user-defined function or method named **FS_scancost()**, which is associated in the **sysams** table with the **am_scancost** keyword
- Sets (adds) the **am_rowids** flag
- Drops the user-defined function or method associated with the **am_getbyid** keyword
- Modifies the **am_costfactor** value

Related Information

Related statements: “CREATE ACCESS_METHOD ” on page 2-90 and “DROP ACCESS_METHOD” on page 2-322

For information about how to set purpose-option specifications, see “Purpose Options” on page 5-50.

For the schema of the **sysams** system catalog table, see the *IBM Informix Guide to SQL: Reference*.

For more information on primary-access methods, see the *IBM Informix Virtual-Table Interface Programmer’s Guide*.

For more information on secondary-access methods, see the *IBM Informix Virtual-Index Interface Programmer’s Guide* and the *IBM Informix User-Defined Routines and Data Types Developer’s Guide*.

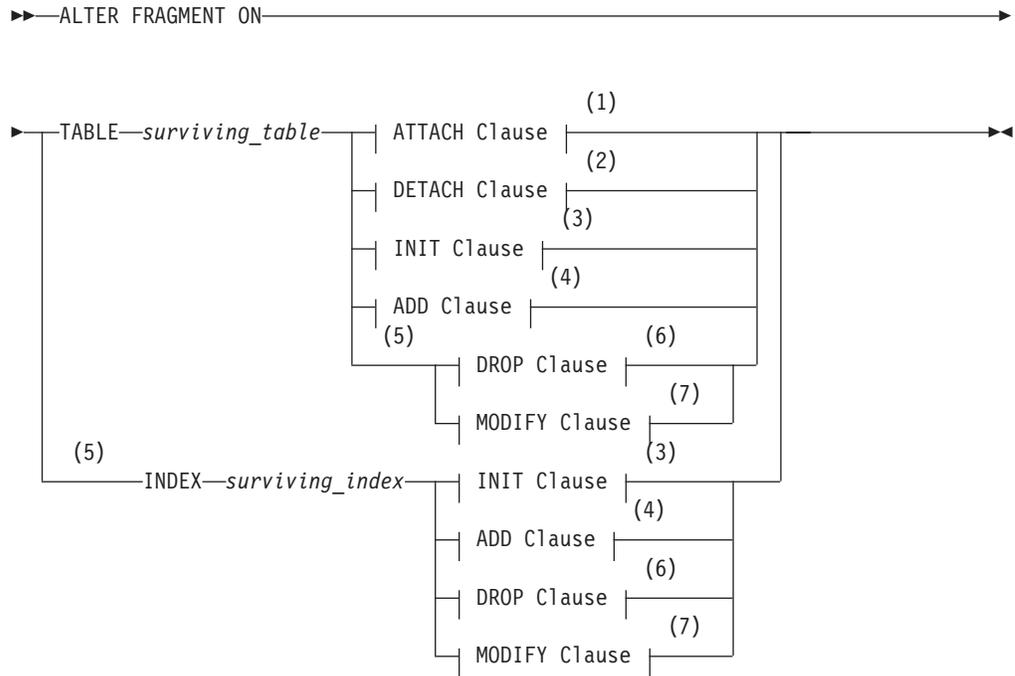
For a discussion of privileges, see the “GRANT ” on page 2-395 statement in this chapter, or the *IBM Informix Database Design and Implementation Guide*.

ALTER FRAGMENT

Use the ALTER FRAGMENT statement to change the distribution strategy or the storage location of an existing table or index.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 See "ATTACH Clause" on page 2-9
- 2 See "DETACH Clause" on page 2-16
- 3 See "INIT Clause" on page 2-18
- 4 See "ADD Clause" on page 2-23
- 5 Dynamic Server only
- 6 See "DROP Clause (IDS)" on page 2-25
- 7 See "MODIFY Clause (IDS)" on page 2-26

Element	Description	Restrictions	Syntax
<i>surviving_index</i>	Index on which to modify the distribution or storage	Must exist when the statement executes	"Identifier " on page 5-21
<i>surviving_table</i>	Table on which to modify the distribution or storage	Must exist. See "Restrictions on the ALTER FRAGMENT Statement" on page 2-8.	"Identifier " on page 5-21

Usage

The ALTER FRAGMENT statement applies only to table fragments or index fragments that are located at the current site (or the current cluster, for Extended Parallel Server). No remote information is accessed or updated.

You must have the Alter or the DBA privilege to change the fragmentation strategy of a table. You must have the Index or the DBA privilege to alter the fragmentation strategy of an index.

Attention: This statement can cause indexes to be dropped and rebuilt. Before undertaking alter operations, check corresponding sections in your *IBM Informix Performance Guide* to review effects and strategies.

Clauses of the ALTER FRAGMENT statement support the following tasks.

Clause Effect

ATTACH

Combines two or more tables that have the same schema into a single fragmented table

DETACH

Detaches a table fragment (or in Extended Parallel Server, a slice) from a fragmentation strategy, and places it in a new table

INIT Provides the following options:

- Defines and initializes a fragmentation strategy on a table
- Creates a fragmentation strategy for tables
- Changes the order of evaluation of fragment expressions
- Alters the fragmentation strategy of a table or index
- Changes the storage location of an existing table

ADD Adds an additional fragment to an existing fragmentation list

DROP Drops an existing fragment from a fragmentation list

MODIFY

Changes an existing fragmentation expression

Use the CREATE TABLE statement or the INIT clause of the ALTER FRAGMENT statement to create fragmented tables.

After a dbspace has been renamed successfully by the **onspaces** utility, only the new name can reference the renamed dbspace. Existing fragmentation strategies for tables or indexes are automatically updated, however, by the database server to replace the old dbspace name with the new name. You do not need to take any additional action to update a distribution strategy or storage location that was defined using the old dbspace name, but you must use the new name if you reference the dbspace in an ALTER FRAGMENT or ALTER TABLE statement.

Restrictions on the ALTER FRAGMENT Statement

You cannot use the ALTER FRAGMENT statement on a temporary table, an external table, or on a view. If your table or index is not already fragmented, the only clauses available to you are INIT and ATTACH.

In Extended Parallel Server, you cannot use ALTER FRAGMENT on a generalized-key (GK) index. If the *surviving_table* has hash fragmentation, the only

clauses available are ATTACH and INIT. You cannot use ALTER FRAGMENT on any table that has a dependent GK index defined on it. In addition, you cannot use this statement on a table that has range fragmentation.

In Dynamic Server, you cannot use ALTER FRAGMENT on a typed table that is part of a table hierarchy.

ALTER FRAGMENT and Transaction Logging

If your database supports transaction logging, ALTER FRAGMENT is executed within a single transaction. If the fragmentation strategy uses large numbers of records, you might run out of log space or disk space. (To alter a fragmentation strategy, the database server requires extra disk space that it later frees.)

If you run out of log space or disk space, try one of the following procedures to reduce your log-space or disk-space requirements:

- Turn off logging and turn it back on again at the end of the operation. This procedure indirectly requires a backup of the **root** dbspace.
- Split the operations into multiple ALTER FRAGMENT statements, moving a smaller portion of records each time.

For information about log-space requirements and disk-space requirements, see your *IBM Informix Administrator's Guide*. That guide also contains detailed instructions about how to turn off logging. For information about backups, refer to your *IBM Informix Backup and Restore Guide*.

Determining the Number of Rows in the Fragment

You can place as many rows into a fragment as the available space in the partition, dbspace, or dbslice allows.

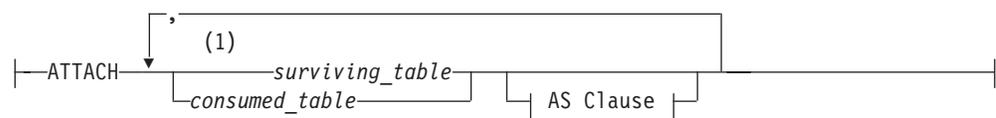
To find out how many rows are in a fragment

1. Run the UPDATE STATISTICS statement on the table. This step fills the **sysfragments** system catalog table with the current table information.
2. Query the **sysfragments** system catalog table to examine the **npused** and **nrows** values. The **npused** column shows the number of data pages used in the fragment, and the **nrows** column shows the number of rows in the fragment.

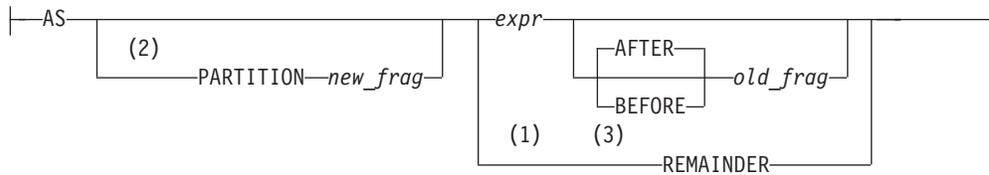
ATTACH Clause

Use the ATTACH clause of the ALTER FRAGMENT ON TABLE statement to combine tables that have identical structures into a fragmentation strategy.

ATTACH Clause:



AS Clause:



Notes:

- 1 Use path no more than once
- 2 Dynamic Server only; required if another *surviving_table* fragment has same name as dbspace
- 3 Required for fragmentation by expression; optional for round-robin fragmentation

Element	Description	Restrictions	Syntax
<i>consumed_table</i>	Table that loses its identity to be merged with <i>surviving_table</i>	Must exist. Cannot include serial columns nor unique, referential, or primary key constraints. See also "General Restrictions for the ATTACH Clause."	"Identifier " on page 5-21
<i>expr</i>	Expression defining which rows are stored in a fragment of a fragmented-by-expression table	Can include only columns from the current table and only data values from a single row. See also "General Restrictions for the ATTACH Clause."	"Condition " on page 4-5; "Expression" on page 4-40
<i>new_frag</i>	Name declared here for a <i>consumed_table</i> fragment. Default is the dbspace name.	Must be unique among the names of partitions of <i>surviving_table</i>	"Identifier " on page 5-21
<i>old_frag</i>	Partition or dbspace where a <i>surviving_table</i> fragment exists	Must exist. See also "Altering Hybrid-Fragmented Tables (XPS)" on page 2-12.	"Identifier " on page 5-21
<i>surviving_table</i>	Table on which to modify the distribution or storage location	Must exist. Cannot have any constraints. See also "Restrictions on the ALTER FRAGMENT Statement" on page 2-8.	"Identifier " on page 5-21

To use this clause, you must have the DBA privilege or else be the owner of the specified tables. The ATTACH clause supports the following tasks:

- Creates a single fragmented table by combining two or more identically-structured, nonfragmented tables
(See "Combining Nonfragmented Tables to Create a Fragmented Table" on page 2-12.)
- Attaches one or more tables to a fragmented table
(See "Attaching a Table to a Fragmented Table" on page 2-12.)

General Restrictions for the ATTACH Clause

This clause is not valid in ALTER FRAGMENT ON INDEX statements.

Any tables that you attach must have been created previously in separate partitions. You cannot attach the same table more than once.

All consumed tables listed in the ATTACH clause must have the same structure as the surviving table. The number, names, data types, and relative position of the columns must be identical.

The *expression* cannot include aggregates, subqueries, or variant functions.

Additional Restrictions on the ATTACH Clause in Dynamic Server:

User-defined routines and references to fields of a ROW-type column are not valid. You cannot attach a fragmented table to another fragmented table.

All of the dbspaces that store the fragments must have the same page size.

For fragmented tables that are protected by a security policy, attaching a fragment to the table fails if any of the following conditions are not satisfied:

- The source table and the target table are both protected by the same security policy;
- Both tables have the same protection granularity (either row-level, or column-level, or both row-level and column-level);
- In both tables, the same set of protected columns is protected by the same security labels. If there is more than one protected column, there can be more than one security label in each table, but the same label must protect the same column in both tables.

If the ATTACH operation fails because one or more of these conditions are not satisfied, you can use the ALTER TABLE statement to make the schemas of the two tables identical, and then repeat the ALTER FRAGMENT ATTACH statement on the modified tables

Only a user who holds the DBSECADM role can reference a protected table in the ALTER FRAGMENT statement.

Additional Restrictions on the ATTACH Clause in XPS: In addition to the general restrictions, every consumed table must be of the same usage type as the surviving table. For information about how to specify the usage type of a table, refer to “Logging Options” on page 2-193.

The ATTACH clause is not valid under either of the following conditions:

- If the consumed tables contain data that belongs in some existing fragment of the surviving table
- If existing data in the surviving table would belong in a new fragment

Thus, you cannot use the ATTACH clause for data movement among fragments. To perform this task, see the “INIT Clause” on page 2-18.

Using the BEFORE, AFTER, and REMAINDER Options

The BEFORE and AFTER options allow you to place a new fragment either before or after an existing dbspace. You cannot use the BEFORE and AFTER options when the distribution scheme is round-robin.

When you attach a new fragment without an explicit BEFORE or AFTER option, the database server places the added fragment at the end of the fragmentation list, unless a remainder fragment exists. If a remainder fragment exists, the new fragment is placed just before the remainder fragment. You cannot attach a new fragment after the remainder fragment.

In Extended Parallel Server, when you create or append to a hybrid-fragmented table, the positioning specification (BEFORE or AFTER) applies to an entire dbslice. You can use any dbspace in a dbslice to identify the dbslice for the BEFORE or AFTER position.

If you omit the AS PARTITION *fragment* specification, the name of the fragment is the name of the dbspace where it is stored.

Combining Nonfragmented Tables to Create a Fragmented Table

When you transform tables with identical table structures into fragments in a single table, you allow the database server to manage the fragmentation instead of allowing the application to manage the fragmentation. The distribution scheme can be round-robin or expression based.

To make a single, fragmented table from two or more identically-structured, nonfragmented tables, the ATTACH clause must contain the surviving table in the *attach list*. The attach list is the list of tables in the ATTACH clause.

In Dynamic Server, to include a **rowid** column in the newly-created single, fragmented table, attach all tables first and then add the **rowid** with the ALTER TABLE statement.

Attaching a Table to a Fragmented Table

To attach a nonfragmented table to an already fragmented table, the nonfragmented table must have been created in a separate dbspace and must have the same table structure as the fragmented table. In the following example, a round-robin distribution scheme fragments the table **cur_acct**, and the table **old_acct** is a nonfragmented table that resides in a separate dbspace. The following example shows how to attach (as the consumed table) **old_acct** to **cur_acct** (as the surviving table):

```
ALTER FRAGMENT ON TABLE cur_acct ATTACH old_acct;
```

In Dynamic Server, when you attach one or more tables to a fragmented table, a *consumed_table* must be nonfragmented.

In Extended Parallel Server, when you attach one or more tables to a fragmented table, a *consumed_table* can be nonfragmented or have hash fragmentation. If you specify a *consumed_table* that has hash fragmentation, the hash column specification must match that of the *surviving_table* and of any other *consumed_table* involved in the attach operation.

Altering Hybrid-Fragmented Tables (XPS)

When you alter a hybrid-fragmented table with an ATTACH or DETACH clause, you need specify only one dbspace to identify the entire set of dbspaces associated with a given expression in the base fragmentation strategy of the table. The set of dbspaces associated with an expression in the base fragmentation strategy of the table might have been defined as one or more dbslices or dbspaces. For more information, see “Fragmenting by HYBRID (XPS)” on page 2-217.

If you know the dbslice name, but not the names of any of the dbspaces that make it up, you can name the first dbspace in the dbslice by appending .1 to the dbslice name. For example, if the dbslice were named **dbsl1**, you could specify **dbsl1.1**.

Effect of the ATTACH Clause

After an ATTACH operation, all consumed tables no longer exist. Any CHECK constraints or NOT NULL constraints on the consumed tables also no longer exist. You must reference the records that were in the consumed tables through the surviving table.

What Happens to Indexes?: A detached index on the surviving table retains its same fragmentation strategy. That is, a detached index does not automatically adjust to accommodate the new fragmentation of the surviving table. For more information on what happens to indexes, see the discussion about altering table fragments in your *IBM Informix Performance Guide*.

In a database that supports transaction logging, an ATTACH operation extends any attached index on the surviving table according to the new fragmentation strategy of the surviving table. All rows in the consumed table are subject to these automatically adjusted indexes. For information on whether the database server completely rebuilds the index on the surviving table or reuses an index that was on the consumed table, see your *IBM Informix Performance Guide*.

In a nonlogging database of Dynamic Server, an ATTACH operation does not extend indexes on the surviving table according to the new fragmentation strategy of the surviving table. To extend the fragmentation strategy of an attached index according to the new fragmentation strategy of the surviving table, you must drop the index and re-create it on the surviving table.

Some ALTER FRAGMENT ... ATTACH operations to attach a fragment can cause the database server to update the index structure. When an index is rebuilt in those cases, the associated column distribution is automatically recalculated, and is available to the query optimizer when it designs query plans for the table on which the fragment was attached.

- For an indexed column (or a set of columns) on which ALTER FRAGMENT ... ATTACH automatically rebuilds a B-tree index, the recalculated column distribution statistics are equivalent to distributions created by the UPDATE STATISTICS statement in HIGH mode.
- If the rebuilt index is not a B-tree index, the automatically recalculated statistics correspond to distributions created by the UPDATE STATISTICS statement in LOW mode.

See also the section “Automatic Calculation of Distribution Statistics” on page 2-144 in the description of the CREATE INDEX statement for additional information about statistical distributions that are produced automatically when an index or constraint is created on an existing table.

What Happens to BYTE and TEXT Columns?: When an ATTACH occurs, BYTE and TEXT fragments of the consumed table become part of the surviving table and continue to be associated with the same rows and data fragments as they were before the ATTACH operation.

In Dynamic Server, each BYTE and TEXT column in every table that is specified in the ATTACH clause must have the same storage type, either blobspace or tblspace. If the BYTE or TEXT column is stored in a blobspace, the same column in all tables must be in the same blobspace. If the BYTE or TEXT column is stored in a tblspace, the same column must be stored in a tblspace in all tables.

In Extended Parallel Server, BYTE and TEXT columns are stored in separate, additional fragments that the database server creates for that purpose in the same dbspace as each regular table fragment where a given row resides.

What Happens to Triggers and Views?: When you attach tables, triggers on the surviving table survive the ATTACH, but triggers on the consumed table are automatically dropped. No triggers are activated by the ATTACH clause, but subsequent data-manipulation operations on the new rows can activate triggers.

Views on the surviving table survive the ATTACH operation, but views on the consumed table are automatically dropped.

What Happens with the Distribution Scheme?: You can attach a nonfragmented table to a table with any type of supported distribution scheme. In general, the resulting table has the same fragmentation strategy as the prior fragmentation strategy of the surviving_table.

When you attach two or more nonfragmented tables, however, the distribution scheme can either be based on expression or round-robin.

For Dynamic Server, only the following distribution schemes can result from combining the distribution schemes of the tables in the ATTACH clause.

Prior Distribution Scheme of Surviving Table	Prior Distribution Scheme of Consumed Table	Resulting Distribution Scheme
None	None	Round-robin or expression
Round-robin	None	Round-robin
Expression	None	Expression

For Extended Parallel Server, this table shows the distribution schemes that can result from different distribution schemes of the tables in the ATTACH clause.

Prior Distribution Scheme of Surviving Table	Prior Distribution Scheme of Consumed Table	Resulting Distribution Scheme
None	None	Round-robin or expression
None	Hash	Hybrid
Round-robin	None	Round-robin
Expression	None	Expression
Hash	None	Hybrid
Hash	Hash	Hybrid
Hybrid	None	Hybrid
Hybrid	Hash	Hybrid

In Extended Parallel Server, when you attach a nonfragmented table to a table that has hash fragmentation, the resulting table has hybrid fragmentation. You can attach a table with a hash distribution scheme to a table that currently has no fragmentation, or hash fragmentation, or hybrid fragmentation. In any of these situations, the resulting table has a hybrid distribution scheme.

The following examples illustrate the use of the ATTACH clause to create fragmented tables with different distribution schemes.

Round-Robin Distribution Scheme: The following example combines nonfragmented tables **pen_types** and **pen_makers** into a single, fragmented table, **pen_types**. Table **pen_types** resides in dbspace **dbsp1**, and table **pen_makers** resides in dbspace **dbsp2**. Table structures are identical in each table.

```
ALTER FRAGMENT ON TABLE pen_types ATTACH pen_types, pen_makers;
```

After you execute the ATTACH clause, the database server fragments the table **pen_types** using a round-robin scheme into two dbspaces: the dbspace that

contained **pen_types** and the dbspace that contained **pen_makers**. Table **pen_makers** is consumed and no longer exists; all rows that were in table **pen_makers** are now in table **pen_types**.

Expression Distribution Scheme: Consider the following example that combines tables **cur_acct** and **new_acct** and uses an expression-based distribution scheme. Table **cur_acct** was originally created as a fragmented table and has fragments in dbspaces **dbsp1** and **dbsp2**. The first statement of the example shows that table **cur_acct** was created with an expression-based distribution scheme. The second statement of the example creates table **new_acct** in **dbsp3** without a fragmentation strategy. The third statement combines the tables **cur_acct** and **new_acct**. Table structures (columns) are identical in each table.

```
CREATE TABLE cur_acct (a int) FRAGMENT BY EXPRESSION
    a < 5 in dbsp1, a >= 5 and a < 10 in dbsp2;
CREATE TABLE new_acct (a int) IN dbsp3;
ALTER FRAGMENT ON TABLE cur_acct ATTACH new_acct AS a>=10;
```

When you examine the **sysfragments** system catalog table after you alter the fragment, you see that table **cur_acct** is fragmented by expression into three dbspaces. For additional information about the **sysfragments** system catalog table, see the *IBM Informix Guide to SQL: Reference*.

In addition to simple range rules, you can use the **ATTACH** clause to fragment by expression with hash or arbitrary rules. For a discussion of all types of expressions that you can use in an expression-based distribution scheme, see “Fragmenting by **EXPRESSION**” on page 2-214.

Warning: When you specify a date value as the default value for a parameter, make sure to specify 4 digits instead of 2 digits for the year. If you specify a 2-digit year, the setting of the **DBCENTURY** environment variable can affect how the database server interprets the date value. For more information, see the *IBM Informix Guide to SQL: Reference*.

Hybrid Fragmentation Distribution Scheme (XPS): Consider a case where monthly sales data is added to the **sales_info** table defined below. Due to the large amount of data, the table is distributed evenly across multiple coservers with a system-defined hash function. To manage monthly additions of data to the table, it is also fragmented by a date expression. The combined hybrid fragmentation is declared in the following **CREATE TABLE** statement:

```
CREATE TABLE sales_info (order_num INT, sale_date DATE, ...)
    FRAGMENT BY HYBRID (order_num) EXPRESSION
        sale_date >= '01/01/2006' AND sale_date < '02/01/2006'
            IN sales_slice_0601,
        sale_date >= '02/01/2006' AND sale_date < '03/01/2006'
            IN sales_slice_0602,
        . . .
        sale_date >= '12/01/2006' AND sale_date < '01/01/2007'
            IN sales_slice_0612;
```

Values for a new month are originally loaded from an external source. The data values are distributed evenly across the same coservers on which the **sales_info** table is defined, using a system-defined hash function on the same column:

```
CREATE TABLE jan_07 (order_num INT, sale_date DATE, ...)
    FRAGMENT BY HASH (order_num) IN sales_slice_0701
INSERT INTO jan_07 SELECT (...) FROM ... ;
```

After data values are loaded, you can attach the new table to **sales_info**. You can issue the following **ALTER FRAGMENT** statement to attach the new table:

```
ALTER FRAGMENT ON TABLE sales_info ATTACH jan_07
AS sale_date >= '01/01/2007' AND sale_date < '02/01/2007';
```

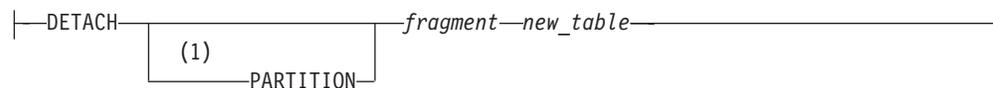
DETACH Clause

Use the DETACH clause of the ALTER FRAGMENT ON TABLE statement to detach a table fragment from a distribution scheme and place the contents into a new nonfragmented table. This clause is not valid in ALTER FRAGMENT ON INDEX statements.

In Extended Parallel Server, the new table can also be a table with hash fragmentation.

For an explanation of distribution schemes, see “FRAGMENT BY Clause” on page 2-212.

DETACH Clause:



Notes:

- 1 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>fragment</i>	Partition (IDS only), dbspace, or dbslice (XPS only) that contains the table fragment to be detached. With hybrid-fragmentation, <i>dbslice</i> identifies a set of dbspaces. See “Altering Hybrid-Fragmented Tables (XPS)” on page 2-12.	Must exist at the time of execution	“Identifier ” on page 5-21
<i>new_table</i>	Nonfragmented table that results after you execute the ALTER FRAGMENT statement. (In XPS, this can also be a hash-fragmented table.)	Must not exist before execution	“Identifier ” on page 5-21

The new table that results from executing the DETACH clause does not inherit any indexes nor constraints from the original table. Only data values remain.

Similarly, the new table does not inherit any privileges from the original table. Instead, the new table has the default privileges of any new table. For further information on default table-level privileges, see the GRANT statement on “Table-Level Privileges” on page 2-399.

The DETACH clause cannot be applied to a table if that table is the parent of a referential constraint or if a **rowid** column is defined on the table.

In Dynamic Server, if you omit the PARTITION keyword, the name of the fragment must be the name of the dbspace where the fragment is stored.

Some ALTER FRAGMENT ... DETACH operations to attach a fragment can cause the database server to update the index structure. When an index is rebuilt in those cases, the database server also recalculates the associated column distributions, and these statistics are available to the query optimizer when it designs query plans for the table from which the fragment was detached:

- For an indexed column (or for a set of columns) on which ALTER FRAGMENT ... DETACH automatically rebuilds a B-tree index, the recalculated column distribution statistics are equivalent to distributions created by the UPDATE STATISTICS statement in HIGH mode.
- If the rebuilt index is not a B-tree index, the automatically recalculated statistics correspond to distributions created by the UPDATE STATISTICS statement in LOW mode.

See also the section “Automatic Calculation of Distribution Statistics” on page 2-144 in the description of the CREATE INDEX statement for additional information about statistical distributions that are produced automatically when an index or constraint is created on an existing table.

In Extended Parallel Server, you cannot use the DETACH clause if the table has a dependent generalized key (GK) index defined on it.

Detach with BYTE and TEXT Columns

If the DETACH clause specifies the first fragment of a table that includes simple large objects of data type BYTE or TEXT, the database server applies locks on the blobspaces of every fragment of the table. To detach any other fragment of the table locks only the blobspaces of the specified fragment, rather than the blobspaces of all fragments, so fewer locks are required if the fragment is not the first.

Detach from a Protected Table

If a successfully executed DETACH clause specifies a table that is protected by a security policy, the database server creates a new table that is protected by the same security policy, and has the same IDSSECURITYLABEL column for row security labels and the same set of protected columns as the original table. The IDSSECURITYLABEL column has a NOT NULL constraint. Only a user who holds the DBSECADM role can reference a protected table in the ALTER FRAGMENT statement.

Detach That Results in a Nonfragmented Table

The following example uses the table **cur_acct** fragmented into two dbspaces, **dbsp1** and **dbsp2**:

```
ALTER FRAGMENT ON TABLE cur_acct DETACH dbsp2 accounts;
```

This example detaches **dbsp2** from the distribution scheme for **cur_acct** and places the rows in a new table, **accounts**. Table **accounts** now has the same structure (column names, number of columns, data types, and so on) as table **cur_acct**, but the table **accounts** does not contain any indexes or constraints from the table **cur_acct**. Both tables are now nonfragmented. The following example shows a table that contains three fragments:

```
ALTER FRAGMENT ON TABLE bus_acct DETACH dbsp3 cli_acct;
```

This statement detaches **dbsp3** from the distribution scheme for **bus_acct** and places the rows in a new table, **cli_acct**. Table **cli_acct** now has the same structure (column names, number of columns, data types, and so on) as **bus_acct**, but the table **cli_acct** does not contain any indexes or constraints from the table **bus_acct**. Table **cli_acct** is a nonfragmented table, but table **bus_acct** remains fragmented.

Detach That Results in a Table with Hash Fragmentation (XPS)

The new table is a hash-fragmented table if the *surviving_table* had hybrid fragmentation and the detached dbslice has more than one fragment. In a hybrid-fragmented table, the dbslice is detached if you specify any dbspace in that

slice. For example, see the **sales_info** table discussed in the “Hybrid Fragmentation Distribution Scheme (XPS)” on page 2-15. Once the January 2007 data is available in **sales_info**, you might archive year-old **sales_info** data.

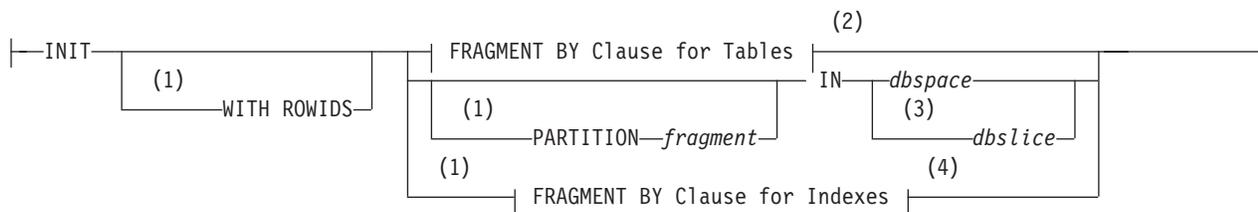
```
ALTER FRAGMENT ON TABLE sales_info
  DETACH sales_slice_0601.1 jan_06;
```

In this example, data from January 2006 is detached from the **sales_info** table and placed in a new table called **jan_06**.

INIT Clause

The INIT clause of the ALTER FRAGMENT statement can define or modify the fragmentation strategy or the storage location of an existing table or (for Dynamic Server) an existing index. The INIT clause has the following syntax.

INIT Clause:



Notes:

- 1 Dynamic Server only
- 2 See “FRAGMENT BY Clause for Tables” on page 2-20
- 3 Extended Parallel Server only
- 4 See “FRAGMENT BY Clause for Indexes (IDS)” on page 2-21

Element	Description	Restrictions	Syntax
<i>dbslice</i>	Dbslice storing fragmented data	Must exist at time of execution	“Identifier ” on page 5-21
<i>dbspace</i>	Dbspace storing fragmented data	Must exist at time of execution	“Identifier ” on page 5-21
<i>fragment</i>	Partition within <i>dbspace</i>	No more than 2048 for same table	“Identifier ” on page 5-21

The INIT clause can accomplish tasks that include the following:

- Move a nonfragmented table from one partition (IDS only) or dbslice (XPS only) or dbspace to another partition or to another dbslice or dbspace.
- Convert a fragmented table to a nonfragmented table.
- Fragment an existing non-fragmented table without redefining it.
- Convert a fragmentation strategy to another fragmentation strategy.
- Fragment an existing index that is not fragmented without redefining the index (IDS only).
- Convert a fragmented index to a nonfragmented index (IDS only).
- Add a new **rowid** column to the table definition (IDS only).

With Extended Parallel Server, you cannot use the INIT clause to change the fragmentation strategy of a table that has a GK index.

When you use the INIT clause to modify a table, the **tabid** value in the system catalog tables changes for the affected table. The **constrid** value of all unique and referential constraints on the table also changes.

For more information about the storage spaces in which you can store a table, see “Using the IN Clause” on page 2-211.

Attention: When you execute the ALTER FRAGMENT statement with this clause, it results in data movement if the table contains any data. If data values move, the potential exists for significant logging, for the transaction being aborted as a long transaction, and for a relatively long exclusive lock being held on the affected tables. Use this statement when it does not interfere with day-to-day operations.

WITH ROWIDS Option

Nonfragmented tables contain a hidden column called **rowid**. Its integer value defines the physical location of a row.

To include a **rowid** column in a fragmented table, you must explicitly request it by using the WITH ROWIDS clause in CREATE TABLE (or ADD ROWIDS in ALTER TABLE, or WITH ROWIDS in ALTER FRAGMENT INIT). The **rowid** in a row of a fragmented table does not identify a physical location for the row in the same way that a **rowid** in a non-fragmented table does.

When you use the WITH ROWIDS option to add a new **rowid** column to a fragmented table, the database server assigns a unique **rowid** number to each row and creates an index that it can use to find the physical location of the row. Performance using this access method is comparable to using a SERIAL, BIGSERIAL, or SERIAL8 column. The **rowid** value of a row cannot be updated, but remains stable during the existence of the row. Each row requires an additional four bytes to store the **rowid** column after you specify the WITH ROWIDS option.

Recommendation: When creating new applications, use primary keys as a method of row identification instead of using **rowid** values.

Converting a Fragmented Table to a Nonfragmented Table

You might decide that you no longer want a table to be fragmented. You can use the INIT clause to convert a fragmented table to a nonfragmented table. The following example shows the original fragmentation definition, as well as how to use the INIT clause of the ALTER FRAGMENT statement to convert the table to a nonfragmented table:

```
CREATE TABLE checks (col1 INT, col2 INT)
  FRAGMENT BY ROUND ROBIN IN dbsp1, dbsp2, dbsp3;

ALTER FRAGMENT ON TABLE checks INIT IN dbsp1;
```

You must use the IN *dbspace* clause to place the table in a dbspace explicitly.

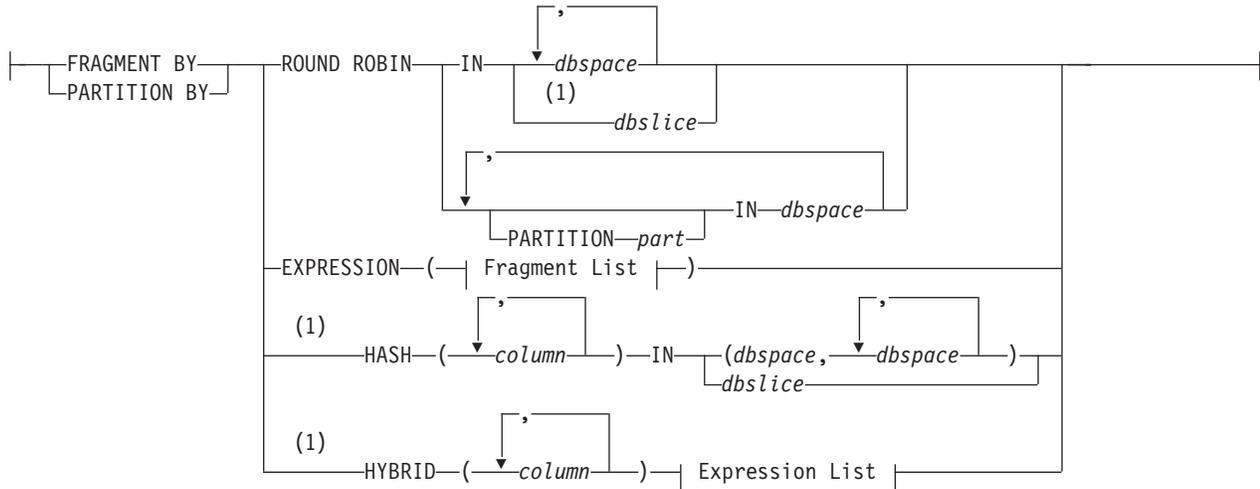
When you use the INIT clause to change a fragmented table to a nonfragmented table, all attached indexes become nonfragmented indexes. In addition, constraints that do not use existing user-defined indexes (detached indexes) become nonfragmented indexes. All newly nonfragmented indexes exist in the same dbspace as the new nonfragmented table.

Using the INIT clause to change a fragmented table to a nonfragmented table has no effect on the fragmentation strategy of detached indexes, nor of constraints that use detached indexes.

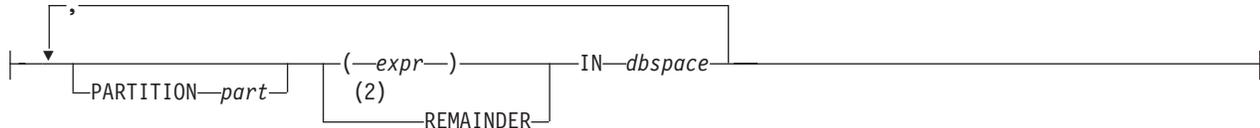
FRAGMENT BY Clause for Tables

Use the FRAGMENT BY (or PARTITION BY) option of the INIT clause to fragment an existing non-fragmented table, or to convert one fragmentation strategy to another. The PARTITION BY keywords are a synonym for the FRAGMENT BY keywords in this context.

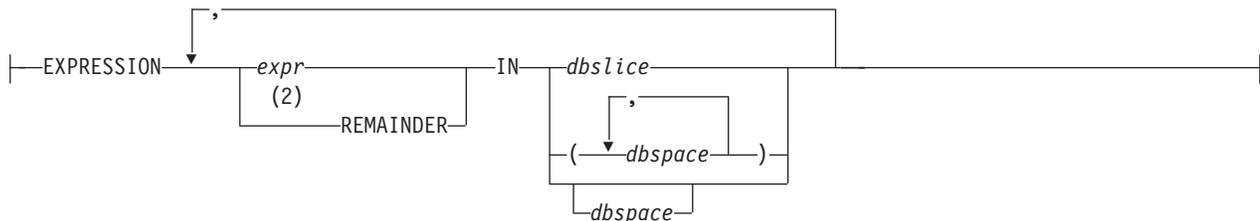
FRAGMENT BY Clause for Tables:



Fragment List:



Expression List:



Notes:

- 1 Extended Parallel Server only
- 2 If included, must be the last item in fragment list

Element	Description	Restrictions	Syntax
<i>column</i>	Column to which the strategy applies	Must exist in the table	"Identifier " on page 5-21
<i>dbslice</i>	Dbslice that contains the table fragment	Must be defined	"Identifier " on page 5-21

Element	Description	Restrictions	Syntax
<i>dbspace</i>	Dbospace that contains the table fragment	Must specify at least 2 but no more than 2,048 dbspaces	"Identifier " on page 5-21
<i>expr</i>	Expression that defines a table fragment	Must evaluate to a Boolean value (t or f)	"Expression" on page 4-40
<i>part</i>	Name declared here for a partition of <i>dbspace</i> that contains the table fragment	Must be unique among names of partitions of <i>table</i>	"Identifier " on page 5-21

In the HYBRID option of Extended Parallel Server, *column* identifies a column on which to apply the hash portion of the hybrid table fragmentation strategy. The expression can reference columns only from the current table and data values only from a single row. Subqueries, aggregates, and the built-in **CURRENT**, **DATE**, **SYSDATE**, and **TODAY** functions are not valid in the expression.

For more information on the available fragmentation strategies for tables, see the "FRAGMENT BY Clause" on page 2-212.

Changing an Existing Fragmentation Strategy on a Table

You can redefine a fragmentation strategy on a table if you decide that your initial strategy does not fulfill your needs. When you alter a fragmentation strategy, the database server discards the existing fragmentation strategy and moves records to fragments as defined in the new fragmentation strategy.

The following example shows the statement that originally defined the fragmentation strategy on the table **account** and then shows an ALTER FRAGMENT statement that redefines the fragmentation strategy:

```
CREATE TABLE account (col1 INT, col2 INT)
  FRAGMENT BY ROUND ROBIN IN dbosp1, dbosp2;
ALTER FRAGMENT ON TABLE account
  INIT FRAGMENT BY EXPRESSION
  col1 < 0 IN dbosp1,
  col2 >= 0 IN dbosp2;
```

If an existing dbospace is full when you redefine a fragmentation strategy, you must not use it in the new fragmentation strategy.

Defining a Fragmentation Strategy on a Nonfragmented Table

The INIT clause can define a fragmentation strategy on a nonfragmented table, regardless of whether the table was created with a storage option.

```
CREATE TABLE balances (col1 INT, col2 INT) IN dbosp1;
ALTER FRAGMENT ON TABLE balances INIT
  FRAGMENT BY EXPRESSION col1 <= 500 IN dbosp1,
  col1 > 500 AND col1 <=1000 IN dbosp2, REMAINDER IN dbosp3;
```

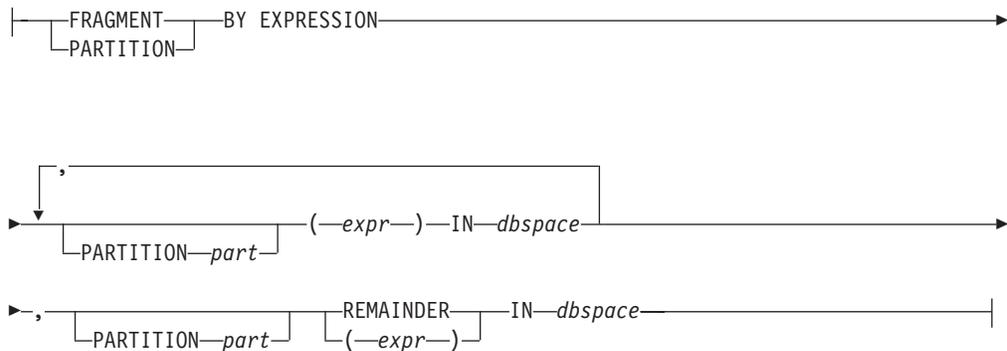
In Dynamic Server, when you use the INIT clause to fragment an existing nonfragmented table, all indexes on the table become fragmented in the same way as the table.

In Extended Parallel Server, when the INIT clause fragments an existing nonfragmented table, any indexes retain their existing fragmentation strategy.

FRAGMENT BY Clause for Indexes (IDS)

The INIT FRAGMENT BY clause for indexes of the ALTER FRAGMENT statement can fragment an existing nonfragmented index by an expression-based distribution scheme without redefining the index.

FRAGMENT BY Clause for Indexes:



Element	Description	Restrictions	Syntax
<i>dbspace</i>	DbSPACE that contains the fragmented information	Must specify at least two, but no more than 2,048 dbSPACES	"Identifier " on page 5-21
<i>expr</i>	Expression defining an index fragment	Must return a Boolean value	"Condition " on page 4-5; "Expression" on page 4-40
<i>part</i>	Name that you declare here for a partition of the specified <i>dbspace</i>	Required for any partition in the same dbSPACE as another partition of the same index. Name must be unique among partitions of the same table.	"Identifier " on page 5-21

The keywords **FRAGMENT BY** and **PARTITION BY** are synonyms in this context. You can convert an existing fragmentation strategy to another expression-based fragmentation strategy. Dynamic Server discards the existing fragmentation strategy and moves the data records to fragments that you define in the new fragmentation strategy. (To convert an existing fragmented index to a nonfragmented index, you can use the **INIT** clause to specify **IN *dbspace*** or else **PARTITION *partition* IN *dbspace*** as the only storage specification for a previously fragmented index.)

The *expression* can contain only columns from the current table and data values from only a single row. No subqueries nor aggregates are allowed. The built-in **CURRENT**, **DATE**, **SYSDATE**, and **TODAY** functions are not valid here.

Fragmenting Unique and System Indexes

You can fragment unique indexes on a table that uses a round-robin or expression-based distribution scheme, but any columns referenced in the fragment expression must be indexed columns. If your index fragmentation strategy violates this restriction, the **ALTER FRAGMENT INIT** statement fails, and work is rolled back.

You might have an attached unique index on a table fragmented by **Column A**. If you attempt to use **ALTER FRAGMENT INIT** to change the table fragmentation to **Column B**, the statement fails because the unique index is defined on **Column A**. To resolve this issue, use the **INIT** clause on the index to detach it from the table fragmentation strategy and fragment it separately.

System indexes (such as those used in referential constraints and unique constraints) use user-defined indexes if the indexes exist. If no user-defined indexes can be used, system indexes remain nonfragmented and are moved to the *dbspace* where the database was created. To fragment a system index, create the fragmented index on the constraint columns and then use the ALTER TABLE statement to add the constraint.

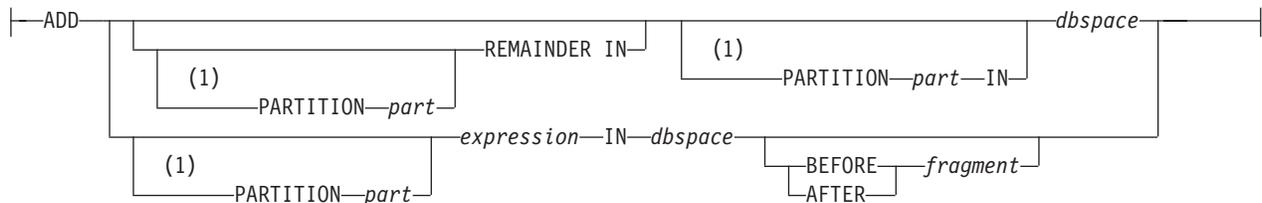
Detaching an Index from a Table-Fragmentation Strategy

You can detach an index from a table-fragmentation strategy with the INIT clause of the ALTER TABLE FOR INDEX statement, so that an attached index becomes a detached index. This breaks any dependency of the index on the table fragmentation strategy. If the INIT clause specifies only IN *dbspace* or PARTITION *fragment* IN *dbspace* for a previously fragmented index, or specifies an index fragmentation strategy that differs from the storage option of the table, then the index becomes a detached index.

ADD Clause

Use the ADD clause to add another fragment to an existing fragment list for a table or (for Dynamic Server only) for an index.

ADD Clause:



Notes:

- 1 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>fragment</i>	Name of an existing fragment in the fragmentation list	Must exist at the time when you execute the statement	"Identifier " on page 5-21
<i>expression</i>	Expression that defines the new fragment that is to be added	Must return a Boolean value (t or f)	"Condition " on page 4-5; "Expression" on page 4-40
<i>dbspace</i>	Name of <i>dbspace</i> to be added to the fragmentation scheme	Must exist at the time when you execute the statement	"Identifier " on page 5-21
<i>part</i>	Name that you declare here for the fragment. The default name is the name of the <i>dbspace</i>	Required for any partition in the same <i>dbspace</i> as another partition of the same index	"Identifier " on page 5-21

The *expression* can contain *column* names only from the current table and data values only from a single row. No subqueries or aggregates are allowed. In addition, the built-in current, date, and time functions are not valid here.

Adding a New Dbspace to a Round-Robin Distribution Scheme:

You can add more *dbspaces* to a round-robin distribution scheme. The following example shows the original round-robin definition:

```
CREATE TABLE book (col1 INT, col2 INT)
  FRAGMENT BY ROUND ROBIN IN dbsp1, dbsp4;
```

To add another dbspace, use the ADD clause, as in this example:

```
ALTER FRAGMENT ON TABLE book ADD dbsp3;
```

Adding a New Partition to a Round-Robin Distribution Scheme (IDS):

In Dynamic Server, you can add a partition of a dbspace to an existing round-robin distribution scheme. The name must be unique within the distribution among partitions of the same dbspace. The following example specifies the same original round-robin fragmentation definition as in the previous section:

```
CREATE TABLE book (col1 INT, col2 INT)
  FRAGMENT BY ROUND ROBIN IN dbsp1, dbsp4;
```

To add a new fragment that is stored in a partition within a dbspace, you can use the ADD clause, as in the following example:

```
ALTER FRAGMENT ON TABLE book
  ADD PARTITION chapter3 IN dbsp1;
```

The new distribution uses **dbsp1**, **dbsp4**, and **chapter3** as storage locations for a 3-part round-robin fragmentation scheme. The records in the fragment **chapter3** are stored in a separate partition of **dbsp1** from the records in the first fragment, which are stored in a partition whose default name is **dbsp1**. (The database server issues an error if you attempt to declare the same name for different fragments of the same fragmented table or index.)

Adding Fragment Expressions:

Adding a fragment expression to the fragmentation list in an expression-based distribution scheme can relocate records from existing fragments into the new fragment. When you add a new fragment into the middle of the fragmentation list, the database server reevaluates all the data existing in fragments after the new fragment. The next example shows the original expression definition:

```
FRAGMENT BY EXPRESSION c1 < 100 IN dbsp1, c1 >= 100 AND c1 < 200 IN dbsp2,
  REMAINDER IN dbsp3;
```

To add another fragment in a new partition of **dbspace dbsp2** to hold rows for **c1** values between 200 and 299, use the following ALTER FRAGMENT statement:

```
ALTER FRAGMENT ON TABLE news
  ADD PARTITION century3 (c1 >= 200 AND c1 < 300) IN dbsp2;
```

Any rows that were formerly in the remainder fragment but that fit the criteria (**c1 >= 200 AND c1 < 300**) move to the new **century3** partition in dbspace **dbsp2**.

Using the BEFORE and AFTER Options:

The BEFORE and AFTER options can position the new fragment either before or after an existing fragment within the fragmentation list. The name of a fragment is the name of the partition where it is stored (or the name of the dbspace, if the dbspace has only one partition). You cannot use the BEFORE and AFTER options if the distribution scheme is round-robin.

When you attach a new fragment without an explicit BEFORE or AFTER option, the database server places the added fragment at the end of the fragmentation list, unless a remainder fragment exists. If a remainder fragment exists, the new

fragment is placed immediately before the remainder fragment. You cannot attach a new fragment after the remainder fragment.

Using the REMAINDER Option:

You cannot add a remainder fragment if one already exists. If you add a new fragment when a remainder exists, the database server retrieves and reevaluates all records in the remainder fragment; some records might move to the new fragment. The remainder fragment is always the last item in the fragment list.

DROP Clause (IDS)

Use the DROP clause to remove an existing fragment from the fragmentation list of a fragmented table or index.

DROP Clause:



Element	Description	Restrictions	Syntax
<i>fragment</i>	Name of a partition of the dbspace that stores the dropped fragment	Must exist when you execute the statement	"Identifier " on page 5-21

If the table is fragmented by expression, you cannot drop a fragment containing data that cannot be moved to another fragment. (If the distribution scheme has a REMAINDER option, or if the expressions overlap, you can drop a fragment that contains data.) You cannot drop a fragment if the table has only two fragments.

When you want to make a fragmented table nonfragmented, use either the INIT clause or the DETACH clause of the ALTER FRAGMENT statement.

If the *fragment* was not given a name when it was created or added then the name of the dbspace is also the name of the fragment.

When you drop a fragment, the underlying partition or dbspace is not affected. Only the fragment data values within that partition or dbspace are affected.

When you drop a fragment, the database server attempts to move all records in the dropped fragment to another fragment. In this case, the destination fragment might not have enough space for the additional records. If this happens, follow one of the procedures that "ALTER FRAGMENT and Transaction Logging" on page 2-9 describes to increase your available space, and retry the procedure.

The following examples show how to drop a fragment from a fragmentation list. The first line shows how to drop an index fragment, and the second line shows how to drop a table fragment.

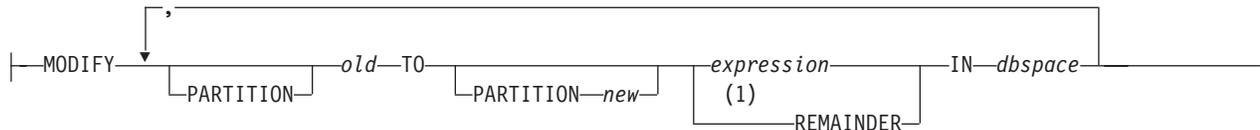
```
ALTER FRAGMENT ON INDEX cust_indx DROP dbsp2;
```

```
ALTER FRAGMENT ON TABLE customer DROP dbsp1;
```

MODIFY Clause (IDS)

Use the MODIFY clause to change an existing fragment expression on an existing partition in the fragmentation list of a table or of an index. You can also use the MODIFY clause to relocate a fragment corresponding to an expression from one dbspace to a different dbspace.

MODIFY Clause:



Notes:

- 1 Use this path no more than once

Element	Description	Restrictions	Syntax
<i>dbspace</i>	Dbspace that stores the modified information	Must exist at time of execution	"Identifier " on page 5-21
<i>expression</i>	Modified expression	Can specify columns in current table only and data from only a single row	"Condition " on page 4-5; "Expression" on page 4-40
<i>new</i>	Name that you declare here for a partition of <i>dbspace</i>	Must be unique in fragmentation list among names of partitions of <i>dbspace</i>	"Identifier " on page 5-21
<i>old</i>	Name of an existing fragment	Must exist at time of execution	"Identifier " on page 5-21

Here *dbspace* and *old* (or *old* and *new*) can be identical, if you are not changing the storage location.

You must declare a *new* partition (using the PARTITION keyword) if more than one fragment of the same table or index is named the same as the *dbspace*.

The *expression* must evaluate to a Boolean value (true or false).

No subqueries or aggregates are allowed in the *expression*. In addition, the built-in CURRENT, DATE, SYSDATE, and TODAY functions are not valid.

When you use the MODIFY clause, the underlying dbspaces are not affected. Only the fragment data values within the partitions or dbspaces are affected.

You cannot change a REMAINDER fragment into a nonremainder fragment if records within the REMAINDER fragment do not satisfy the *new expression*.

When you use the MODIFY clause to change an expression without changing the storage location for the expression, you must use the same name for the *old* and the *new* fragment (or else the same name for *old* and for *dbspace*, if the *dbspace* consists of only a single partition, as in the following example):

```

ALTER FRAGMENT ON TABLE cust_acct
  MODIFY dbsp1 TO acct_num < 65 IN dbsp1;
    
```

When you use the `MODIFY` clause to move an expression from one dbspace to another, *old* is the name of the dbspace where the expression was previously located, and *dbspace* is the new location for the expression:

```
ALTER FRAGMENT ON TABLE cust_acct
  MODIFY dbbsp1 TO acct_num < 35 IN dbbsp2;
```

Here the distribution scheme for the `cust_acct` table is modified so that all row items in column `acct_num` that are less than 35 are now contained in the dbspace `dbbsp2`. These items were formerly contained in the dbspace `dbbsp1`.

When you use the `MODIFY` clause to move an expression from one partition of a dbspace to another partition, *old* is the name of the partition where information fragmented by the expression was previously located, and *new* is the name of the partition that is the new location for the expression, as in the following example:

```
ALTER FRAGMENT ON TABLE cust_acct
  MODIFY PARTITION part1 TO PARTITION part2 (acct_num < 35) IN dbbsp2;
```

Here the distribution scheme for the `cust_acct` table is modified so that all row items in column `acct_num` that have a value less than 35 are now contained in the partition `part2` of dbspace `dbbsp2`. These items were formerly contained in the partition `part1`.

To use the `MODIFY` clause both to change the *expression* and to move its corresponding fragment to a new storage location, you must change the *expression* and you must also specify the name of a different *dbspace* or partition.

If the indexes on a table are attached indexes, and you modify the table, the index fragmentation strategy is also modified.

Examples

The following series of examples illustrate the use of `ALTER FRAGMENT ON INDEX` with the `INIT`, `ADD`, `DROP`, and `MODIFY` options. This first example creates an index stored in `dbbsp1`:

```
CREATE INDEX item_idx ON items (stock_num) IN dbbsp1;
```

The following statements modify the index to add fragmentation. Values up to 50 are stored in `dbbsp1`, values between 51 and 80 in `dbbsp2`, and the remainder in `dbbsp3`:

```
ALTER FRAGMENT ON INDEX item_idx INIT
  FRAGMENT BY EXPRESSION
    stock_num <= 50 IN dbbsp1,
    stock_num > 50 AND stock_num <= 80 IN dbbsp2,
    REMAINDER IN dbbsp3;
```

The following statements add a new fragment to the index:

```
ALTER FRAGMENT ON INDEX item_idx
  ADD stock_num > 80 AND stock_num <= 120 IN dbbsp4;
```

The following statements change the first fragment of the index:

```
ALTER FRAGMENT ON INDEX item_idx
  MODIFY dbbsp1 TO stock_num <= 40 IN dbbsp1;
```

The following statements drop the fragment in `dbbsp4` from the index:

```
ALTER FRAGMENT ON INDEX item_idx
  DROP dbbsp4;
```

The following statements create a fragmented index with fragments in partitions of the dbspaces **dbsp1** and **dbsp2**:

```
ALTER FRAGMENT ON INDEX item_idx INIT
PARTITION BY EXPRESSION
PARTITION part1 stock_num <= 10 IN dbsp1,
PARTITION part2 stock_num > 20 AND stock_num <= 30 IN dbsp1,
PARTITION part3 REMAINDER IN dbsp2;
```

The following statements add a new partition:

```
ALTER FRAGMENT ON INDEX item_idx ADD
PARTITION part4 stock_num > 30 AND stock_num <= 40 IN dbsp2
BEFORE part3;
```

Related Information

Related statements: “CREATE TABLE ” on page 2-191, “CREATE INDEX ” on page 2-124, and “ALTER TABLE” on page 2-44

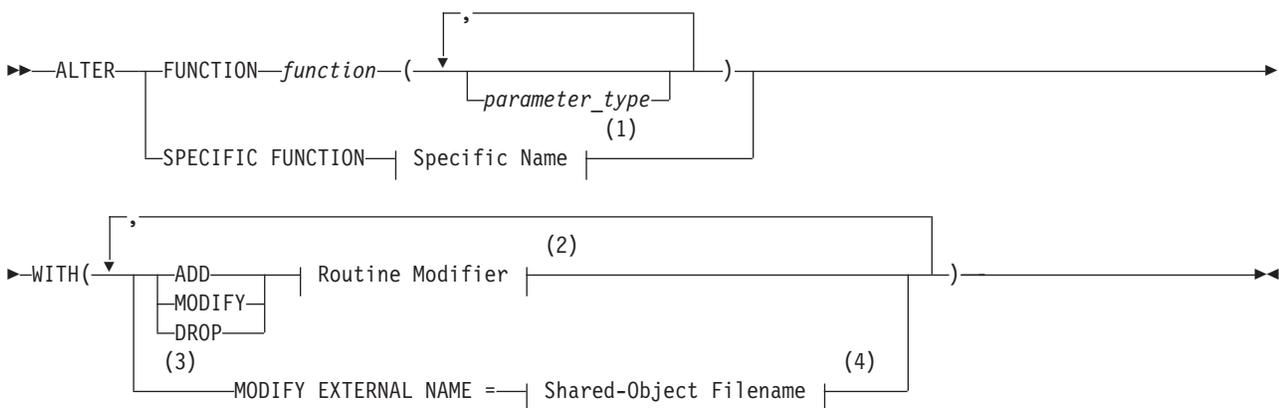
For a discussion of fragmentation strategy, refer to the *IBM Informix Database Design and Implementation Guide*.

For information on how to maximize performance when you make fragment modifications, see your *IBM Informix Performance Guide*.

ALTER FUNCTION

Use the ALTER FUNCTION statement to change the routine modifiers or pathname of a user-defined function. Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 See “Specific Name ” on page 5-73
- 2 See “Routine Modifier” on page 5-59
- 3 External routines only
- 4 See “ Shared-Object Filename ” on page 5-70

Element	Description	Restrictions	Syntax
<i>function</i>	User-defined function to be modified	Must be registered in the database. If the name does not uniquely identify a function, you must enter one or more appropriate values for <i>parameter_type</i> .	"Identifier " on page 5-21
<i>parameter_type</i>	Data type of a parameter	Must be the same data types (and specified in the same order) as in the definition of <i>function</i>	"Data Type " on page 4-21

Usage

The ALTER FUNCTION statement can modify a user-defined function to tune its performance by modifying characteristics that control how the function executes. You can also add or replace related user-defined routines (UDRs) that provide alternatives for the query optimizer, which can improve performance.

All modifications take effect on the next invocation of the function.

Only the UDR owner or the DBA can use the ALTER FUNCTION statement.

Keywords That Introduce Modifications

Use the following keywords to introduce what you modify in the UDR.

Keyword	Effect on Specified Routine Modifier
ADD	Add a new routine modifier to the UDR
MODIFY	Change an attribute of the routine modifier
DROP	Delete the routine modifier from the UDR
MODIFY EXTERNAL NAME (for external functions only)	Replace the file specification of the executable file. When the IFX_EXTEND_ROLE configuration parameter = ON, this option is valid only for users to whom the DBSA has granted the EXTEND role. With IFX_EXTEND_ROLE = OFF (or not set), the UDR owner or the DBA can use this option.
WITH	Introduces all modifications

If the routine modifier is a BOOLEAN value, MODIFY sets the value to be t (equivalent of using the keyword ADD to add the routine modifier). For example, both of the following statements alter the **func1** function so that it can be executed in parallel in the context of a parallelizable data query:

```
ALTER FUNCTION func1 WITH (MODIFY PARALLELIZABLE);
ALTER FUNCTION func1 WITH (ADD PARALLELIZABLE);
```

See also "Example of Altering Routine Modifiers" on page 2-35.

Related Information

Related Statements: "ALTER PROCEDURE" on page 2-32, "ALTER ROUTINE" on page 2-33, "CREATE FUNCTION " on page 2-113, and "CREATE PROCEDURE " on page 2-153

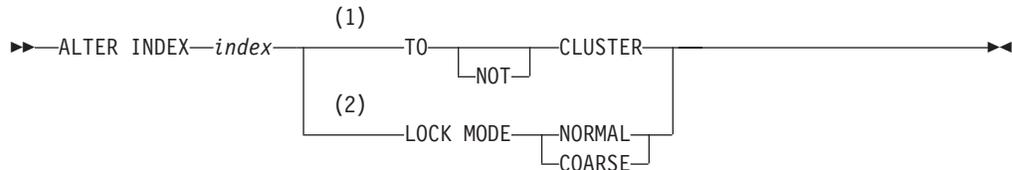
For a discussion of how to create and use SPL routines, see the *IBM Informix Guide to SQL: Tutorial*. For a discussion of how to create and use external routines, see *IBM Informix User-Defined Routines and Data Types Developer's Guide* and the *IBM*

ALTER INDEX

Use the ALTER INDEX statement to change the clustering attribute or the locking mode of an existing index.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 Dynamic Server only
- 2 Extended Parallel Server only

Element	Description	Restrictions	Syntax
<i>index</i>	Name of the index to be altered	Must exist	"Identifier " on page 5-21

Usage

ALTER INDEX is valid only for indexes that the CREATE INDEX statement created explicitly. ALTER INDEX cannot modify an index on a temporary table, nor an index that the database server created implicitly to support a constraint.

You cannot change the collating order of an existing index. If you use ALTER INDEX to modify an index after the SET COLLATION statement of Dynamic Server has specified a non-default collating order, the SET COLLATION statement has no effect on the index.

TO CLUSTER Option

The TO CLUSTER option causes the database server to reorder the rows of the physical table according to the indexed order.

The next example shows how you can use the ALTER INDEX TO CLUSTER statement to order the rows in the **orders** table physically. The CREATE INDEX statement creates an index on the **customer_num** column of the table. Then the ALTER INDEX statement causes the physical ordering of the rows.

```

CREATE INDEX ix_cust ON orders (customer_num);
ALTER INDEX ix_cust TO CLUSTER;
  
```

For an ascending index, TO CLUSTER puts rows in lowest-to-highest order. For a descending index, the rows are reordered in highest-to-lowest order.

When you reorder, the entire file is rewritten. This process can take a long time, and it requires sufficient disk space to maintain two copies of the table.

While a table is clustering, it is locked IN EXCLUSIVE MODE. When another process is using the table to which the index name belongs, the database server cannot execute the ALTER INDEX statement with the TO CLUSTER option; it returns an error unless lock mode is set to WAIT. (When lock mode is set to WAIT, the database server retries the ALTER INDEX statement.)

Over time, if you modify the table, you can expect the benefit of an earlier cluster to disappear because rows are added in space-available order, not sequentially. You can recluster the table to regain performance by issuing another ALTER INDEX TO CLUSTER statement on the clustered index. You do not need to drop a clustered index before you issue another ALTER INDEX TO CLUSTER statement on a currently clustered index.

Extended Parallel Server cannot use the CLUSTER option on STANDARD tables.

TO NOT CLUSTER Option

The TO NOT CLUSTER option drops the cluster attribute on the index name without affecting the physical table. Because no more than one clustered index can exist on a given table, you must use the TO NOT CLUSTER option to release the cluster attribute from one index before you assign it to another index on the same table. The following statements illustrate how to remove clustering from one index and how a second index physically reclusters the table:

```
CREATE UNIQUE INDEX ix_ord ON orders (order_num);  
  
CREATE CLUSTER INDEX ix_cust ON orders (customer_num);  
ALTER INDEX ix_cust TO NOT CLUSTER;  
  
ALTER INDEX ix_ord TO CLUSTER;
```

The first two statements create indexes for the **orders** table and cluster the physical table in ascending order on the **customer_num** column. The last two statements recluster the physical table in ascending order on the **order_num** column.

LOCK MODE Options (XPS)

Use the LOCK MODE clause to specify the locking granularity of the index.

When you use the COARSE mode, index-level locks are acquired on the index instead of item-level or page-level locks. This mode reduces the number of lock calls on an index. The COARSE mode offers performance advantages when you know that the index is unlikely to change; for example, when read-only operations are performed on the index.

Use the NORMAL mode to place item-level or page-level locks on the index as necessary. Use this mode when the index gets updated frequently.

When the database server executes the LOCK MODE COARSE option, it acquires an exclusive lock on the table for the duration of the ALTER INDEX statement. Any transactions currently using a lock of finer granularity must complete before the database server switches to the COARSE lock mode.

Related Information

Related statements: "CREATE INDEX " on page 2-124 and "CREATE TABLE " on page 2-191

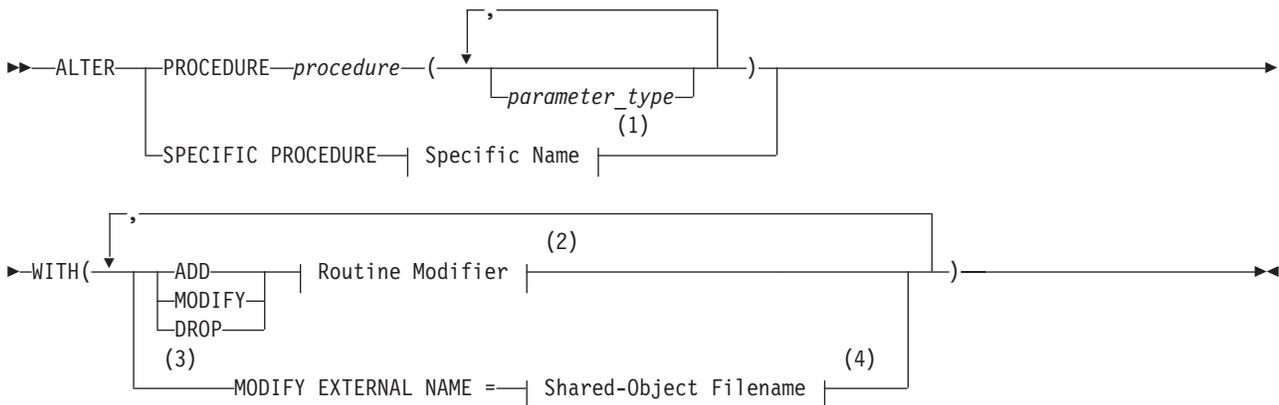
For a discussion of the performance implications of clustered indexes, see your *IBM Informix Performance Guide*.

ALTER PROCEDURE

Use the ALTER PROCEDURE statement to change the routine modifiers or pathname of a previously defined external procedure.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 See "Specific Name " on page 5-73
- 2 See "Routine Modifier" on page 5-59
- 3 External routines only
- 4 See " Shared-Object Filename " on page 5-70

Element	Description	Restrictions	Syntax
<i>procedure</i>	User-defined procedure to modify	Must be registered in the database. If the name does not uniquely identify a function, you must enter one or more appropriate values for <i>parameter_type</i> .	"Identifier " on page 5-21
<i>parameter_type</i>	Data type of a parameter	Must be the same data types (and specified in the same order) as in the definition of <i>procedure</i>	"Data Type " on page 4-21

Usage

The ALTER PROCEDURE statement enables you to modify an external procedure to tune its performance by modifying characteristics that control how it executes. You can also add or replace related UDRs that provide alternatives for the optimizer, which can improve performance. All modifications take effect on the next invocation of the procedure.

Only the UDR owner or the DBA can use the ALTER PROCEDURE statement.

If the procedure name is not unique among routines registered in the database, you must enter one or more appropriate values for *parameter_type*.

The following keywords introduce what you want to modify in *procedure*.

Keyword	Effect
ADD	Add a new routine modifier to the UDR
MODIFY	Change an attribute of a routine modifier
DROP	Delete a routine modifier from the UDR
MODIFY EXTERNAL NAME (for external procedures only)	Replace the file specification of the executable file. When the IFX_EXTEND_ROLE configuration parameter = ON, this option is valid only for users to whom the DBSA has granted the EXTEND role. With IFX_EXTEND_ROLE = OFF (or not set), the UDR owner or the DBA can use this option.
MODIFY EXTERNAL NAME (for external procedures only)	Replace the file specification of the executable file. (Valid only for users who have the EXTEND role)
WITH	Introduces all modifications

If the routine modifier is a BOOLEAN value, MODIFY sets the value to be T (equivalent to using the keyword ADD to add the routine modifier). For example, both of the following statements alter the **proc1** procedure so that it can be executed in parallel in the context of a parallelizable data query:

```
ALTER PROCEDURE proc1 WITH (MODIFY PARALLELIZABLE);
ALTER PROCEDURE proc1 WITH (ADD PARALLELIZABLE);
```

See also “Example of Altering Routine Modifiers” on page 2-35.

Related Information

Related statements: “ALTER ROUTINE,” “CREATE PROCEDURE ” on page 2-153, “DROP PROCEDURE ” on page 2-332, and “DROP ROUTINE ” on page 2-335.

For a discussion of SPL routines, see the *IBM Informix Guide to SQL: Tutorial*.

For a discussion of how to create and use external routines, see *IBM Informix User-Defined Routines and Data Types Developer’s Guide*.

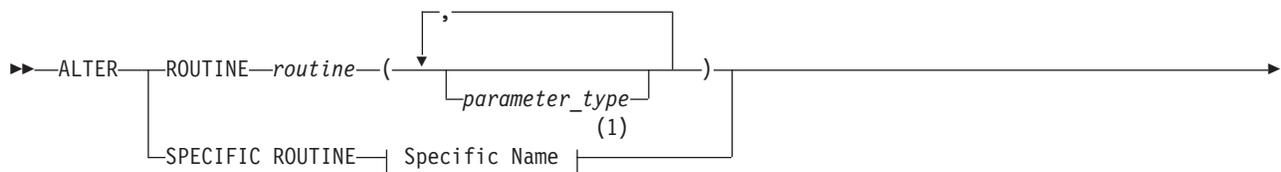
For information about how to create C UDRs, see the *IBM Informix Datablade API Programmer’s Guide*.

ALTER ROUTINE

Use the ALTER ROUTINE statement to change the routine modifiers or pathname of a previously defined user-defined routine (UDR).

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax





Notes:

- 1 See "Specific Name " on page 5-73
- 2 See "Routine Modifier" on page 5-59
- 3 External routines only
- 4 See " Shared-Object Filename " on page 5-70

Element	Description	Restrictions	Syntax
<i>routine</i>	User-defined routine to modify	Must be registered in the database. If the name does not uniquely identify a routine, you must enter one or more appropriate values for <i>parameter_type</i> .	"Identifier " on page 5-21
<i>parameter_type</i>	Data type of a parameter	Must be the same data types (and specified in the same order) as in the definition of <i>routine</i>	"Data Type " on page 4-21

Usage

The ALTER ROUTINE statement allows you to modify a previously defined UDR to tune its performance by modifying characteristics that control how the UDR executes. You can also add or replace related UDRs that provide alternatives for the optimizer, which can improve performance.

This statement is useful when you do not know whether a UDR is a user-defined function or a user-defined procedure. When you use this statement, the database server alters the appropriate user-defined procedure or user-defined function.

All modifications take effect on the next invocation of the UDR.

Only the UDR owner or the DBA can use the ALTER ROUTINE statement.

Restrictions

If the name does not uniquely identify a UDR, you must enter one or more appropriate values for *parameter_type*.

When you use this statement, the type of UDR cannot be ambiguous. The UDR that you specify must refer to either a user-defined function or a user-defined procedure. If either of the following conditions exist, the database server returns an error:

- The name (and parameters) that you specify applies to both a user-defined procedure and a user-defined function.
- The specific name that you specify applies to both a user-defined function and a user-defined procedure.

Keywords That Introduce Modifications

Use these keywords to introduce the items in the UDR that you want to modify:

Keyword	Effect
ADD	Add a routine modifier to the UDR
DROP	Delete a routine modifier from the UDR
MODIFY	Change an attribute of the routine modifier
MODIFY EXTERNAL NAME (for external routines only)	Replace the file specification of the executable file. When the IFX_EXTEND_ROLE configuration parameter = ON, this option is valid only for users to whom the DBSA has granted the EXTEND role. With IFX_EXTEND_ROLE = OFF (or not set), the UDR owner or the DBA can use this option.
WITH	Introduces all modifications

If the routine modifier is a BOOLEAN value, MODIFY sets the value to be T (equivalent to using the keyword ADD to add the routine modifier).

For example, both of the following statements alter the **func1** UDR so that it can be executed in parallel in the context of a parallelizable data query statement:

```
ALTER ROUTINE func1 WITH (MODIFY PARALLELIZABLE);
ALTER ROUTINE func1 WITH (ADD PARALLELIZABLE);
```

Example of Altering Routine Modifiers

Suppose you have an external function **func1** that is set to handle NULL values and has a cost per invocation set to 40. The following ALTER ROUTINE statement adjusts the settings of the function by dropping the ability to handle NULL values, tunes the **func1** by changing the cost per invocation to 20, and indicates that the function can execute in parallel:

```
ALTER ROUTINE func1(CHAR, INT, BOOLEAN)
  WITH (
    DROP HANDLESNULLS,
    MODIFY PERCALL_COST = 20,
    ADD PARALLELIZABLE
  );
```

Because the name **func1** is not unique to the database, the data type parameters are specified so that the routine signature is unique. If this function had a Specific Name, for example, **raise_sal**, specified when it was created, you could identify the function with the following first line:

```
ALTER SPECIFIC ROUTINE raise_sal;
```

Related Information

Related Statements: “ALTER FUNCTION ” on page 2-28, “ALTER PROCEDURE” on page 2-32, “CREATE FUNCTION ” on page 2-113, “CREATE PROCEDURE ” on page 2-153, “DROP FUNCTION ” on page 2-327, “DROP PROCEDURE ” on page 2-332, and “DROP ROUTINE ” on page 2-335.

For a discussion of how to create and use SPL routines, see the *IBM Informix Guide to SQL: Tutorial*.

For a discussion of how to create and use external routines, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

For information about how to create C UDRs, see the *IBM Informix Datablade API Programmer's Guide*.

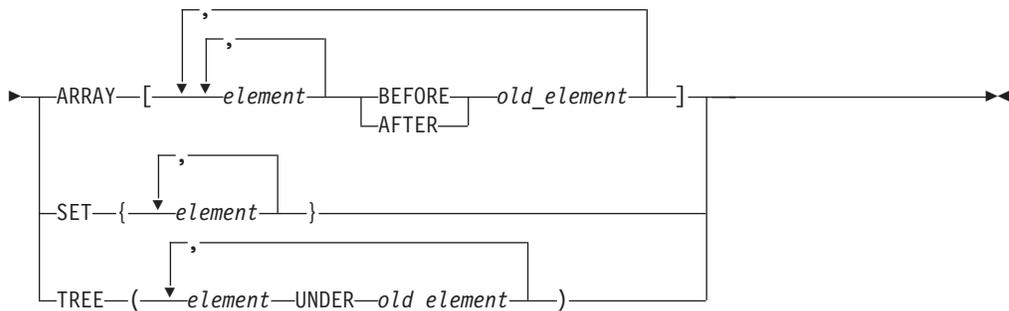
ALTER SECURITY LABEL COMPONENT

Use the ALTER SECURITY LABEL COMPONENT statement to add one or more new elements to an existing security label component in the current database.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

►► ALTER SECURITY LABEL COMPONENT— *component*—ADD—►►



Element	Description	Restrictions	Syntax
<i>component</i>	Component to which <i>element</i> is added	Must already exist in the database.	"Identifier" on page 5-21
<i>element</i>	New element of <i>component</i>	Must be unique among elements of this <i>component</i> , and no longer than 32 bytes. The left (() and right ()) parentheses, comma (,), and colon (:) symbols are not valid characters.	"Quoted String" on page 4-189
<i>old_element</i>	Existing element of <i>component</i>	Must be an element of <i>component</i> .	"Quoted String" on page 4-189

Usage

Only DBSECADM can issue the ALTER SECURITY LABEL COMPONENT statement, which defines new elements of an existing security label component. The new elements become part of any security policy defined in a CREATE SECURITY POLICY statement that references the specified component.

A security label component consists of a set of no more than 64 *elements* that the CREATE SECURITY LABEL COMPONENT statement defines as string constants. Each string constant can have no more than 32 bytes, and each must be unique

among elements of this component. The declaration of each element, which is a valid value that the component can have, defines a category of data sensitivity. By adding new elements to an existing component, the ALTER SECURITY LABEL COMPONENT statement expands the set of values that a component can have within a security policy that includes the component, or within a security label that supports the security policy.

When the ALTER SECURITY LABEL COMPONENT statement executes successfully, Dynamic Server updates the following tables of the system catalog of the current database:

- The **sysseclabelcomponentelements** table, to add new rows for each new elements of the component,
- the **sysseclabelcomponents** table, to show the new cardinality of the security elements that comprise the modified security component.

This statement can define new elements of a security label component, but it cannot modify or drop an existing element. If the security design changes so that different elements are required, DBSECADM can add the new elements, if the total number of elements remains within the size and cardinality limits, and not use any obsolete elements in defining labels that include the component.

Alternatively, DBSECADM can use the DROP SECURITY LABEL COMPONENT statement to drop the component, and then use the CREATE SECURITY LABEL COMPONENT statement to redefine a new component that has only the required elements. You cannot, however, drop a security component if it is part of an existing security policy. See “DROP SECURITY ” on page 2-338 for information about restrictions on dropping security label components and other security objects of Dynamic Server.

The security label component to which new elements are added must be one of three component types. The ARRAY, SET, or TREE keyword that immediately follows the *component* name must specify the same component type that the CREATE SECURITY LABEL COMPONENT statement specified when the component was originally defined. The syntax for specifying the new list of elements depends on whether the specified component is of type ARRAY, SET, or TREE, which are the three types of security component that Dynamic Server supports.

The ADD ARRAY Clause

A security label component of type ARRAY is an ordered set of no more than 64 elements. The order in which array elements are declared is significant, because it defines a descending order of data sensitivity, with each successive element ranking lower in data sensitivity than the preceding element. The set of label elements of the array and their comma (,) separators must be enclosed between a pair of bracket ([...]) symbols. The same new *element* cannot be declared more than once in the same ADD ARRAY clause.

In the ADD ARRAY clause, the BEFORE or AFTER keyword must follow the new element (or a comma-separated list of new elements) to specify the position of the new element within the descending order of data sensitivity. Within restrictions on the size and the number of elements, this syntax enables DBSECADM to insert a new element anywhere in the array, including the highest or the lowest position, or between consecutive existing elements. The ALTER SECURITY LABEL COMPONENT statement fails with an error, however, if the BEFORE or AFTER keyword of the ADD ARRAY clause specifies an array element that was not

previously defined, either when the array component was created, or else in a previous ALTER SECURITY LABEL COMPONENT statement that modified the same array component.

If multiple ALTER SECURITY LABEL COMPONENT operations are performed to add elements to the same component of type ARRAY, DBSECADM might not be able to reach the maximum of 64 array elements because of how array elements are encoded. For information on how security elements are encoded, see the Informix Dynamic Server Security Guide.

The following example defines a security label component of type ARRAY called **aquilae** that is an ordered set of five elements, with **imperator** highest in data sensitivity and **asinus** lowest. The subsequent ALTER SECURITY LABEL COMPONENT statement adds two new elements:

- a new element called **legatus** that ranks between **imperator** and **tribunus**
- a new element called **cunctator** that ranks below **asinus** as the new low in data sensitivity.

```
CREATE SECURITY LABEL COMPONENT aquilae
  ARRAY [ "imperator", "tribunus", "centurio", "miles", "asinus" ];

ALTER SECURITY LABEL COMPONENT aquilae
  ADD ARRAY [ "legatus" BEFORE "tribunus", "cunctator" AFTER "asinus" ];
```

Successful execution of this ALTER SECURITY LABEL COMPONENT ... ADD ARRAY statement modifies the **aquilae** security label component array, so that the new descending order of component elements becomes this: **imperator, legatus, tribunus, centurio, miles, asinus, cunctator**.

The ADD SET Clause

A security label component of type SET is an unordered set of no more than 64 elements. The order in which the elements of a SET component are declared is not significant. The set of elements of the array and their comma separators must be enclosed between a pair of braces ({ ... }) symbols. The same new *element* cannot be declared more than once in the same ADD SET clause.

The following example defines a type SET security label component called **departments** that is an unordered set of three elements, called **Marketing**, **HR**, and **Finance**, which the ALTER SECURITY LABEL COMPONENT statement modifies by adding three new elements called **Training**, **QA**, and **Security**:

```
CREATE SECURITY LABEL COMPONENT departments
  SET { 'Marketing', 'HR', 'Finance' };

ALTER SECURITY LABEL COMPONENT departments
  ADD SET { 'Training', 'QA', 'Security' };
```

Unlike ADD ARRAY or ADD TREE specifications, ADD SET operations of ALTER SECURITY LABEL COMPONENT do not create "greater than" or "less than" data-sensitivity relationships among the new and existing elements of the redefined component, because the elements of a type SET component have no implicit order of data sensitivity.

The ADD TREE Clause

A security label component of type TREE has the logical topology of a simple graph with no loops. Each TREE component has a single root node and no more than 63 additional nodes. Any new elements that the ALTER SECURITY LABEL

COMPONENT statement adds to this hierarchy must be inserted below the root node. The string constant for each new node must be followed by the keyword UNDER and by the string constant for some previously declared node. The set of elements of the TREE component, including their UNDER keywords and comma separators, must be enclosed between a pair of parenthesis ((...)) symbols.

The component element specified after the UNDER keyword is called the *parent* of the new element that precedes the same UNDER keyword. The new element is called the *child* of that parent element. The ALTER SECURITY LABEL COMPONENT statement fails with an error, however, if the ADD TREE clause specifies a parent element that is not already defined in the database for this component. The UNDER keyword cannot be followed by an element that is added to the component in the same ADD TREE clause.

The string constant that designates the root node of a TREE component has the highest data sensitivity of all the nodes within the TREE hierarchy. In any subset of successive parent nodes and child nodes in the tree, each non-root element has lower data sensitivity than its parent element or than any ancestor of its parent element, but has higher data sensitivity than any of its child elements or than the descendents of its child elements.

When a user who holds no exemptions attempts to access a data row that is protected by a label that includes a TREE component, a read operation fails if the security label of the user does not include an element that matches one of the TREE elements for the same component of the data row label, or that matches an ancestor of one of those elements. Unless the security policy of the label includes the OVERRIDE clause, a write operation also fails in the same circumstances. If the data row label has multiple TREE components, the user security label must include a matching (or an ancestral) element value for every TREE component of the data row security label.

In the following example, the ALTER SECURITY LABEL COMPONENT statement modifies a tree component called **Oakland** by adding two new nodes to its tree structure that was defined with six nodes by this CREATE SECURITY LABEL COMPONENT statement:

```
CREATE SECURITY LABEL COMPONENT Oakland
TREE ( 'Port' ROOT,
       'Downtown' UNDER 'Port',
       'Airport' UNDER 'Port',
       'Estuary' UNDER 'Airport',
       'Avenues' UNDER 'Downtown',
       'Hills' UNDER 'Avenues');

ALTER SECURITY LABEL COMPONENT Oakland
  ADD TREE ( 'Uptown' UNDER 'Port',
            'Bay' UNDER 'Estuary');
```

Here new **Uptown** node is a child of **Port**, which has the highest data sensitivity because it is the root node. The new **Bay** node is the child of **Estuary**, which is the child of **Airport**, which is the child of **Port**, implying that **Bay** has a lower data sensitivity than these three nodes of the hierarchy. In practice, it is unlikely that any data would be labeled with **Port**, rather than classified at a lower level. The **Port** value might be used for a label granted to a user who is allowed to access all of the data about the **Port**.

If the ALTER SECURITY LABEL COMPONENT statement in this example succeeds, and a subsequently defined data row label specifies **Bay** as its value for

the **Oakland** component, a user with no exemption for the security policy who attempts to read the protected row in a query would need either **Port**, **Airport**, **Estuary**, or **Bay** as a user label value to satisfy this component of the data row label. Values of **Uptown** or **Downtown** for this component in the user label are insufficient, because they do not match **Bay** and are not ancestors of **Bay**. For a query to read the protected row, the security label of the user must also include values that satisfy any other components of the row security label, and the user must also hold Select access privilege on the table and at least Connect access privilege on the database that contains the protected row.

The ADD TREE clause cannot interpose a new node between an existing child node and its parent.

Related Information

Related statements: ALTER TABLE, "CREATE SECURITY LABEL " on page 2-175, "CREATE SECURITY LABEL COMPONENT " on page 2-177, "CREATE SECURITY POLICY " on page 2-181, CREATE TABLE, "DROP SECURITY " on page 2-338, GRANT EXEMPTION, GRANT SECURITY LABEL, "RENAME SECURITY " on page 2-490, REVOKE EXEMPTION, and REVOKE SECURITY LABEL

For a discussion of LBAC security objects, see the Informix Dynamic Server Security Guide.

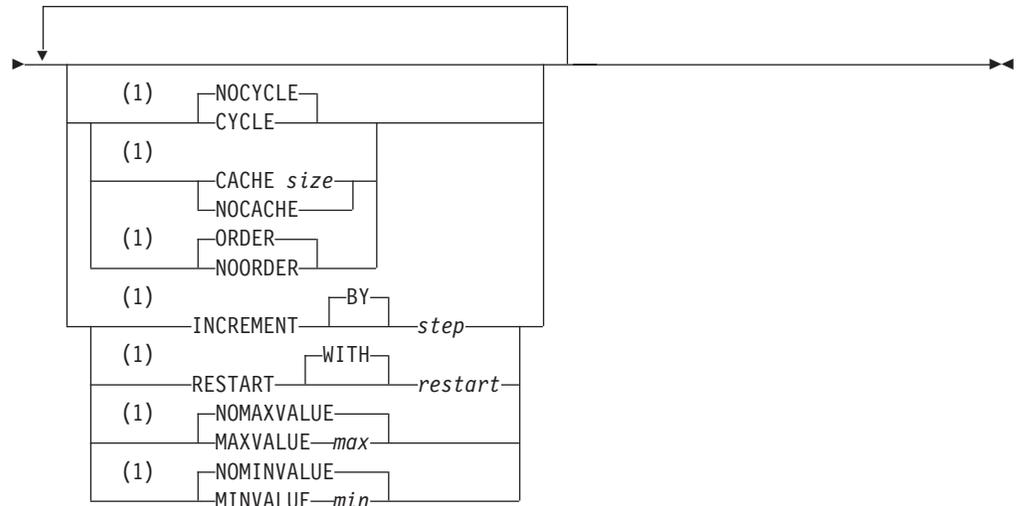
ALTER SEQUENCE

Use the ALTER SEQUENCE statement to modify the definition of a sequence object.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

►► ALTER SEQUENCE owner. *sequence* →



Notes:

- 1 Use path no more than once

Element	Description	Restrictions	Syntax
<i>max</i>	New upper limit on values	Must be integer > CURRVAL and <i>restart</i>	"Literal Number " on page 4-185
<i>min</i>	New lower limit on values	Must be integer < CURRVAL and <i>restart</i>	"Literal Number " on page 4-185
<i>owner</i>	Owner of <i>sequence</i>	Cannot be changed by this statement	" Owner Name " on page 5-45
<i>restart</i>	New first value in sequence	Must be integer in the INT8 range	"Literal Number " on page 4-185
<i>sequence</i>	Name of existing sequence	Must exist. Cannot be a synonym.	"Identifier " on page 5-21
<i>size</i>	New number of values to preallocate in memory	Integer > 2 but < cardinality of values in one cycle (= $\lceil (max - min) / step \rceil$)	"Literal Number " on page 4-185
<i>step</i>	New interval between successive values	Must be a nonzero integer	"Literal Number " on page 4-185

Usage

ALTER SEQUENCE redefines an existing sequence object. It only affects subsequently generated values (and any unused values in the sequence cache). You cannot use the ALTER SEQUENCE statement to rename a sequence nor to change the owner of a sequence.

You must be the owner, or the DBA, or else have the Alter privilege on the sequence to modify its definition. Only elements of the sequence definition that you specify explicitly in the ALTER SEQUENCE statement are modified. An error occurs if you make contradictory changes, such as specifying both MAXVALUE and NOMAXVALUE, or both the CYCLE and NOCYCLE options.

INCREMENT BY Option

Use the INCREMENT BY option to specify a new interval between successive numbers in a sequence. The interval, or *step* value, can be a positive whole number

(for ascending sequences) or a negative whole number (for descending sequences) in the INT8 range. The BY keyword is optional.

RESTART WITH Option

Use the RESTART WITH option to specify a new first number of the sequence. The *restart* value must be an integer within the INT8 range that is greater than or equal to the *min* value (for an ascending sequence) or that is less than or equal to the *max* value (for a descending sequence), if *min* or *max* is specified in the ALTER SEQUENCE statement. The WITH keyword is optional.

When you modify a sequence using the RESTART option, the *restart* value is stored in the **syssequences** system catalog table only until the first use of the sequence object in a NEXTVAL expression. After that, the value is reset in the system catalog. Use of the **dbschema** utility can increment sequence objects in the database, creating gaps in the generated numbers that might not be expected in applications that require serialized integers.

MAXVALUE or NOMAXVALUE Option

Use the MAXVALUE option to specify a new upper limit of values in the sequence. The maximum value, or *max*, must be an integer in the INT8 range that is greater than *sequence.CURRVAL* and *restart* (or greater than the *origin* in the original CREATE SEQUENCE statement, if *restart* is not specified).

Use the NOMAXVALUE option to replace the current limit with a new default maximum of 2e64 for ascending sequences or -1 for descending sequences.

MINVALUE or NOMINVALUE Option

Use the MINVALUE option to specify a new lower limit of values in the sequence. The minimum value, or *min*, must be an integer the INT8 range that is less than *sequence.CURRVAL* and *restart* (or less than the *origin* in the original CREATE SEQUENCE statement, if *restart* is not specified).

Use the NOMINVALUE option to replace the current lower limit with a default of 1 for ascending sequences or -(2e64) for descending sequences.

CYCLE or NOCYCLE Option

Use the CYCLE option to continue generating sequence values after the sequence reaches the maximum (ascending) or minimum (descending) limit, to replace the NOCYCLE attribute. After an ascending sequence reaches *max*, it generates the *min* value for the next value. After a descending sequence reaches *min*, it generates the *max* value for the next sequence value.

Use the NOCYCLE option to prevent the sequence from generating more values after reaching the declared limit. Once the sequence reaches the limit, the next reference to *sequence.NEXTVAL* returns an error message.

CACHE or NOCACHE Option

Use the CACHE option to specify a new number of sequence values that are preallocated in memory for rapid access. The cache size must be a whole number in the INT range that is less than the number of values in a cycle (or less than $(\text{max} - \text{min}) / \text{step}$). The minimum size is 2 preallocated values.

Use NOCACHE to have no values preallocated in memory. (See also the description of SEQ_CACHE_SIZE in "CREATE SEQUENCE" on page 2-184.)

ORDER or NOORDER Option

These keywords have no effect on the behavior of the sequence. The sequence always issues values to users in the order of their requests, as if the ORDER keyword were always specified. The ORDER and NOORDER keywords are accepted by the ALTER SEQUENCE statement, however, for compatibility with implementations of sequence objects in other dialects of SQL.

Examples

The examples below are based on the following sequence object and table:

```
CREATE SEQUENCE seq_2
  INCREMENT BY 1 START WITH 1
  MAXVALUE 30 MINVALUE 0
  NOCYCLE CACHE 10 ORDER;
```

```
CREATE TABLE tab1 (col1 int, col2 int);
INSERT INTO tab1 VALUES (0, 0);
```

```
INSERT INTO tab1 (col1, col2) VALUES (seq_2.NEXTVAL, seq_2.NEXTVAL)
```

```
SELECT * FROM tab1;
```

col1	col2
0	0
1	1

```
ALTER SEQUENCE seq_2
  RESTART WITH 5
  INCREMENT by 2
  MAXVALUE 300;
```

```
INSERT INTO tab1 (col1, col2) VALUES (seq_2.NEXTVAL, seq_2.NEXTVAL)
INSERT INTO tab1 (col1, col2) VALUES (seq_2.NEXTVAL, seq_2.NEXTVAL)
SELECT * FROM tab1;
```

col1	col2
0	0
1	1
5	5
7	7

Related Information

Related statements: "CREATE SEQUENCE" on page 2-184, "DROP SEQUENCE" on page 2-340, "RENAME SEQUENCE" on page 2-491, "CREATE SYNONYM" on page 2-187, "DROP SYNONYM" on page 2-341, "GRANT" on page 2-395, "REVOKE" on page 2-494, "INSERT" on page 2-430, "UPDATE" on page 2-689, and "SELECT" on page 2-528

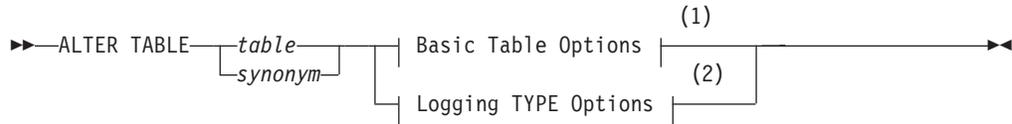
For information about the **syssequences** system catalog table in which sequence objects are registered, see the *IBM Informix Guide to SQL: Reference*.

For information about initializing, generating, or reading values from a sequence, see "NEXTVAL and CURRVAL Operators (IDS)" on page 4-79.

ALTER TABLE

Use the ALTER TABLE statement to modify the definition of an existing table.

Syntax



Notes:

- 1 See "Basic Table Options" on page 2-45
- 2 See "Logging TYPE Options" on page 2-69

Element	Description	Restrictions	Syntax
<i>synonym</i>	Synonym for the table to be altered	Synonym and its table must exist; USETABLENAME must not be set	"Identifier " on page 5-21
<i>table</i>	Name of table to be altered	Must exist in current database	"Identifier " on page 5-21

Usage

In Dynamic Server, the database server performs the actions in the ALTER TABLE statement in the order that you specify. If any action fails, the entire operation is cancelled.

Altering a table on which a view depends might invalidate the view.

Warning: The clauses available with this statement have varying performance implications. Before you undertake alter operations, check corresponding sections in your *IBM Informix Performance Guide* to review effects and strategies.

You cannot alter a temporary table. You also cannot alter a violations or diagnostics table. In addition, you cannot add, drop, or modify a column if the table that contains the column has a violation table or a diagnostics table associated with it. If the **USETABLENAME** environment variable is set, you cannot specify a *synonym* for the table in the ALTER TABLE statement.

In Extended Parallel Server, if a table has range fragmentation, only the Logging TYPE options and LOCK MODE clause are valid. All other ALTER TABLE options return an error.

For a RAW or (in Extended Parallel Server) a STATIC table, the Logging TYPE options are the only part of the ALTER TABLE statement that you can use.

To use ALTER TABLE, you must meet one of the following conditions:

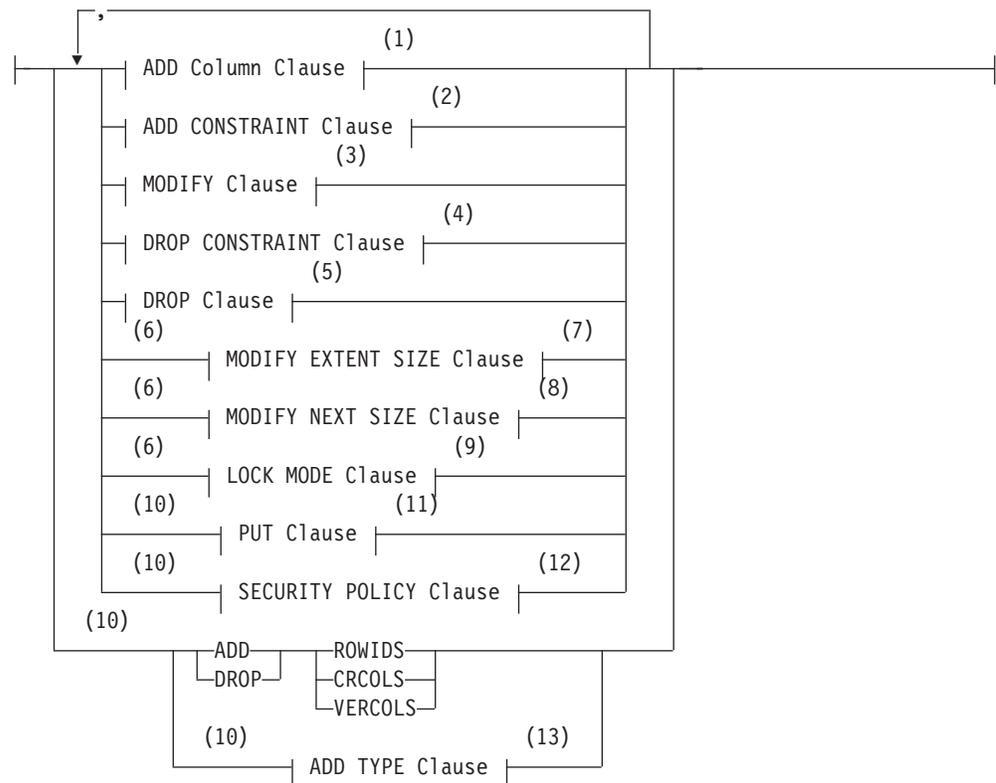
- You must have DBA privilege on the database containing the table.
- You must own the table (and for Extended Parallel Server, you must also have the Resource privilege).
- You must have the Alter privilege on the specified table and the Resource privilege on the database where the table resides.

- To add a referential constraint, you must have the DBA or References privilege on either the referenced columns or the referenced table.
- To drop a constraint, you must have the DBA privilege or be the owner of the constraint. If you are the owner of the constraint but not the owner of the table, you must have Alter privilege on the specified table. You do not need the References privilege to drop a constraint.

Basic Table Options

The Basic Table Options segment of ALTER TABLE has the following syntax.

Basic Table Options:



Notes:

- 1 See "ADD Column Clause" on page 2-47
- 2 See "ADD CONSTRAINT Clause" on page 2-63
- 3 See "MODIFY Clause" on page 2-57
- 4 See "DROP CONSTRAINT Clause" on page 2-66
- 5 See "DROP Clause" on page 2-55
- 6 Use path once
- 7 See "MODIFY EXTENT SIZE" on page 2-66
- 8 See "MODIFY NEXT SIZE Clause" on page 2-67
- 9 See "LOCK MODE Clause" on page 2-68
- 10 Dynamic Server only

- 11 See “PUT Clause (IDS)” on page 2-62
- 12 See “SECURITY POLICY Clause” on page 2-70
- 13 See “ADD TYPE Clause (IDS)” on page 2-72

You can use the Basic Table Options segment to modify the schema of a table by adding, modifying, or dropping columns and constraints, or changing the extent size or locking granularity of a table. The database server performs alterations in the order that you specify. If any of the actions fails, the entire operation is cancelled.

With Dynamic Server, you can also associate a table with a named ROW type or specify a new storage space to store large-object data. You can also add or drop **rowid** columns and shadow columns for Enterprise Replication or to support secondary server update operations of the USELASTCOMMITTED feature. You cannot, however, specify these options in conjunction with any other alterations.

Using the ADD ROWIDS Keywords

Use the ADD ROWIDS keywords to add a new column called **rowid** to a fragmented table. (Fragmented tables do not contain the hidden **rowid** column by default.) When you add a **rowid** column, the database server assigns a unique number to each row that remains stable for the life of the row. The database server creates an index that it uses to find the physical location of the row. After you add the **rowid** column, each row of the table contains an additional 4 bytes to store the **rowid** value.

The following example uses the ADD ROWIDS option to add a new **rowid** column of type INTEGER to a fragmented table called **frag1**:

```
ALTER TABLE frag1 ADD ROWIDS;
```

Tip: Use the ADD ROWIDS clause only on fragmented tables. In nonfragmented tables, the **rowid** column remains unchanged. It is recommended that you use primary keys as an access method rather than exploiting the **rowid** column.

For additional information about the **rowid** column, refer to your *IBM Informix Administrator's Reference*.

Using the DROP ROWIDS Keywords

The DROP ROWIDS keywords can drop a **rowid** column that you added (with either the ALTER TABLE or ALTER FRAGMENT statement) to a fragmented table.

The following example drops the **rowid** column from the **frag1** table:

```
ALTER TABLE frag1 DROP RWIDS;
```

You cannot drop the **rowid** column of a nonfragmented table.

Using the ADD CRCOLS Keywords

The ADD CRCOLS keywords create shadow columns, **cdrserver** and **cdrtime**, that Enterprise Replication uses for conflict resolution. These columns enable the database server to use the timestamp or SPL conflict-resolution rule.

In the following example, the **cdrserver** and **cdrtime**, shadow columns are added to the **customer** table:



Notes:

- 1 See “Data Type ” on page 4-21
- 2 See “DEFAULT Clause” on page 2-49
- 3 See “Single-Column Constraint Format” on page 2-50
- 4 See “Add Column Security” on page 2-55

Element	Description	Restrictions	Syntax
<i>column</i>	Name of column before which <i>new_column</i> is to be placed	Must already exist in the table	“Identifier ” on page 5-21
<i>new_column</i>	Name of column that you are adding	You cannot add a serial column if the table contains data	“Identifier ” on page 5-21

The following restrictions apply to the ADD clause:

- You cannot add a serial column to a table that contains data.
- In Extended Parallel Server, you cannot add a column to a table that has a bit-mapped index.
- You cannot add columns beyond the maximum row size of 32,767 bytes.

The following restrictions affect the use of the ADD Column clause to add a column of the IDSSECURITYLABEL data type to support label-based access control:

- If the table has no security policy, a user who holds the DBSECADM role must also include the ADD SECURITY POLICY statement to specify an existing security policy.
- Only a user who holds the DBSECADM role can add a column of type IDSSECURITYLABEL.
- A table can have at most one column of type IDSSECURITYLABEL.
- The IDSSECURITYLABEL column cannot have column protection.
- The IDSSECURITYLABEL column has an implicit NOT NULL constraint by default. If no *label* name for the default security label is specified in the DEFAULT clause, the default value for this column is the security label for write access that is held by the user.
- The IDSSECURITY LABEL column cannot have any explicit single-column constraints, and it cannot be part of multiple-column referential or check constraints.

Using the BEFORE Option

The BEFORE option specifies the column before which to add the new columns. In the following example, the BEFORE option directs the database server to add the **item_weight** column before the **total_price** column:

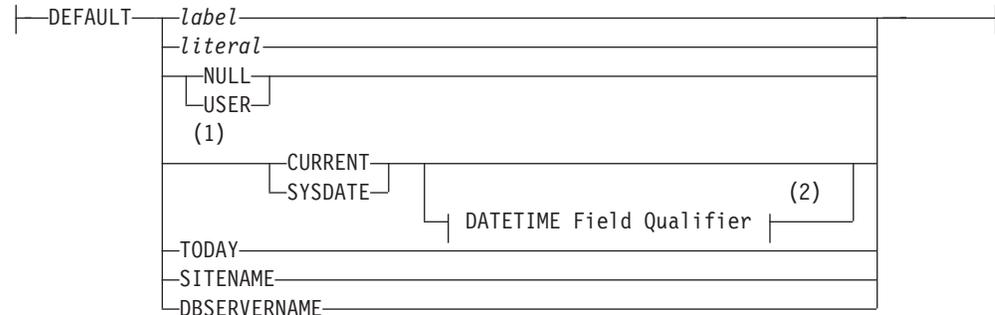
```
ALTER TABLE items
  ADD (item_weight DECIMAL(6,2) NOT NULL BEFORE total_price);
```

If you do not include the BEFORE option, the database server adds the new column or list of columns to the end of the table definition by default.

DEFAULT Clause

Use the DEFAULT clause to specify value that the database server should insert in a column when an explicit value for the column is not specified.

DEFAULT Clause:



Notes:

- 1 Informix extension
- 2 See " DATETIME Field Qualifier" on page 4-38

Element	Description	Restrictions	Syntax
<i>label</i>	Name of a security label	Must exist and must belong to the security policy that protects the table. The column must be of type IDSSECURITYLABEL.	"Identifier " on page 5-21
<i>literal</i>	Literal default value for the column	Must be appropriate for the data type of the column. See "Using a Literal as a Default Value" on page 2-196.	"Identifier " on page 5-21

You cannot specify a default value for serial columns. If the table that you are altering already has rows in it when you add a column that contains a default value, the default values are applied only to rows inserted after the ALTER TABLE MODIFY statement that added the new column. Any rows that existed before the new column was added have a NULL value in the new column, unless you update those rows to insert some non-NULL value. New rows that you insert will have the default value that the MODIFY Clause specified, unless you insert some other value into the new column.

The following example adds a column of data type DECIMAL(6,2) to the **items** table. In **items**, the new column **item_weight** has a literal default value:

```

ALTER TABLE items
  ADD item_weight DECIMAL (6, 2) DEFAULT 2.00
  BEFORE total_price;
  
```

In this example, each existing row in the **items** table has a default value of 2.00 for the **item_weight** column.

For more information about the options of the DEFAULT clause, refer to "DEFAULT Clause" on page 2-195.

DEFAULT Labels

When DBSECADM adds a `IDSSECURITYLABEL` column to a table that is protected by a security policy, the `DEFAULT label` specification is required unless the table is empty. If the table is not empty, the specified `label` is inserted into the existing rows of the table.

An error is issued if the `DEFAULT` clause specifies a security label for a column whose data type is not `IDSSECURITYLABEL`, or if the table has no security policy, or if the security policy of the `label` is not the security policy of the table, .

To define a specific `label` as the default value of an `IDSSECURITYLABEL` column, specify the `label` name without the `policy` qualifier, rather than as `policy.label`, because the security policy of the table is the only valid policy for any security label that protects data in the table.

The statement in the following example adds security policy **MegaCorp** to table **T1** and specifies column-level protection for the table by declaring a new column **D** of type `IDSSECURITYLABEL`, whose default value is a security label called **mylabel**:

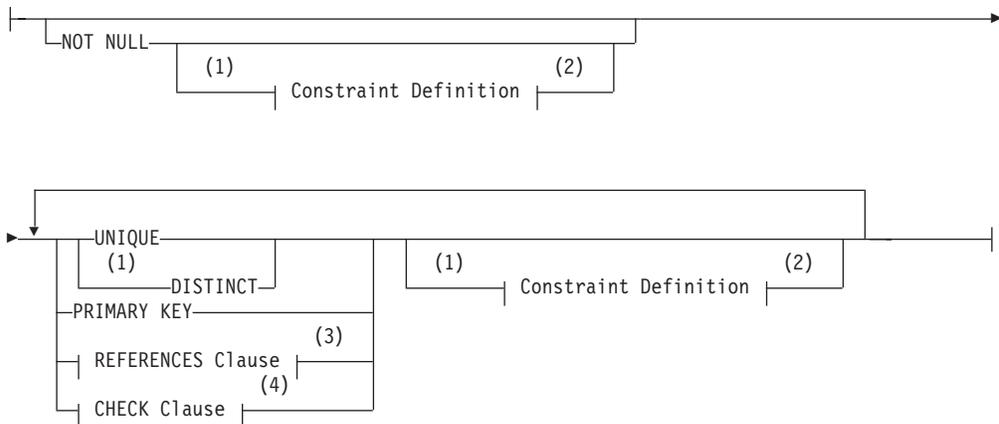
```
ALTER TABLE T1
  ADD (D IDSSECURITYLABEL DEFAULT mylabel1)
  ADD SECURITY POLICY MegaCorp;
```

Because no `BEFORE` clause is included, column **D** is last among the columns in the schema of table **T1**. This statement fails if any of the database objects that it references (except new column **D**) does not already exist in the database.

Single-Column Constraint Format

Use the Single-Column Constraint Format to associate one or more constraints with a single column.

Single-Column Constraint Format:



Notes:

- 1 Informix extension
- 2 See “Constraint Definition” on page 2-51
- 3 See “REFERENCES Clause” on page 2-52
- 4 See “CHECK Clause” on page 2-54

You cannot specify a primary-key or unique constraint on a new column if the table contains data. In the case of a unique constraint, however, the table can contain a *single* row of data. When you want to add a column with a primary-key constraint, the table must be empty when you issue the ALTER TABLE statement.

The following rules apply when you place primary-key or unique constraints on existing columns:

- When you place a primary-key or unique constraint on a column or on a set of columns, the database server creates an internal B-tree index on the constrained column or set of columns, and automatically calculates column statistics, equivalent to distributions created by the UPDATE STATISTICS statement in HIGH mode, unless a user-created index was already defined on the column or set of columns.
- When you place a primary-key or unique constraint on a column or set of columns, and a unique index already exists on that column or set of columns, the constraint shares that index. If the existing index allows duplicates, however, the database server returns an error. You must then drop the existing index before you can add the constraint.
- When you place a primary-key or unique constraint on a column or set of columns, and a referential constraint already exists on that column or set of columns, the duplicate index is upgraded to UNIQUE (if possible), and the index is shared.

You cannot place a unique constraint on a BYTE or TEXT column, nor can you place referential constraints on columns of these data types. A check constraint on a BYTE or TEXT column can check only for IS NULL, IS NOT NULL, or LENGTH.

The IDSECURITYLABEL column has an implicit NOT NULL constraint, but it cannot have explicit single-column constraints nor be part of multiple-column referential constraints or check constraints. If the constraint is on a column that stores encrypted data, Dynamic Server cannot enforce the constraint.

When you place a referential constraint on a column or on a set of columns, and an index already exists on that column or set of columns, the index is upgraded to UNIQUE (if possible) and the index is shared.

Using NOT NULL Constraints with ADD

If a table contains data, when you add a column with a NOT NULL constraint you must also include a DEFAULT clause. If the table is empty, however, you can add a column and apply only the NOT NULL constraint. The following statement is valid whether or not the table contains data:

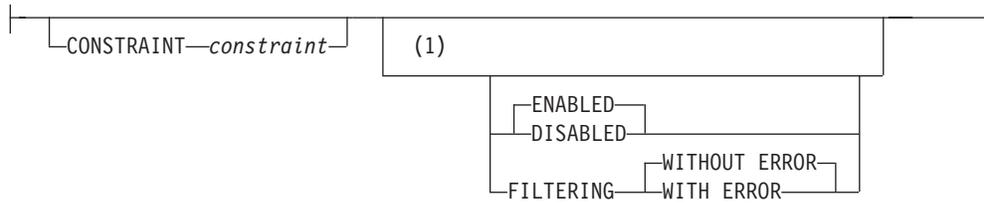
```
ALTER TABLE items
  ADD (item_weight DECIMAL(6,2)
  DEFAULT 2.0 NOT NULL
  BEFORE total_price);
```

Constraint Definition

In Dynamic Server, use the Constraint Definition portion of the ALTER TABLE statement to declare the name of a constraint and to set the mode of the constraint to disabled, enabled, or filtering.

In Extended Parallel Server, use the Constraint Definition portion of the ALTER TABLE statement to declare the name of a constraint.

Constraint Definition:



Notes:

- 1 Dynamic Server only

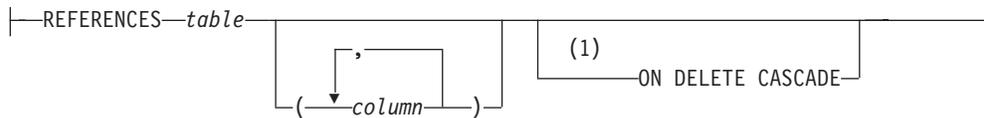
Element	Description	Restrictions	Syntax
<i>constraint</i>	Name declared here for the constraint	Must be unique among the names of indexes and constraints in database	"Identifier " on page 5-21

For more information about constraint-mode options, see "Choosing a Constraint-Mode Option (IDS)" on page 2-204.

REFERENCES Clause

The REFERENCES clause has the following syntax.

REFERENCES Clause:



Notes:

- 1 Informix extension

Element	Description	Restrictions	Syntax
<i>column</i>	Referenced column in the referenced table	See "Restrictions on Referential Constraints."	"Identifier " on page 5-21
<i>table</i>	The referenced table	The referenced and the referencing tables must reside in the same database	"Identifier " on page 5-21

The REFERENCES clause allows you to place a foreign-key constraint on one or more columns. The referenced column can be in the same table as the referencing column, or in a different table in the same database.

If the referenced table is different from the referencing table, the default *column* is the primary-key column. If the referenced table is the same as the referencing table, there is no default.

Restrictions on Referential Constraints

You must have the REFERENCES privilege to create a referential constraint.

The following restrictions apply to the *column* that is specified (the referenced column) in the REFERENCES clause:

- The referenced and referencing tables must be in the same database.
- The referenced column (or set of columns) must have a unique or primary-key constraint.
- The referencing and referenced columns must be the same data type.
(The only exception is that a referencing column must be an integer data type if the referenced column is a serial data type.)
- You cannot place a referential constraint on a BYTE or TEXT column.
- Constraints uses the collation in effect at their time of creation.
- A column-level REFERENCES clause can include only a single column name.
- Maximum number of columns in a table-level REFERENCES clause is 16.
- In Dynamic Server, the total length of the columns in a table-level REFERENCES clause cannot exceed 390 bytes.
- In Extended Parallel Server, the total length of the columns in a table-level REFERENCES clause cannot exceed 255 bytes.

Default Column for the References Clause

If the referenced table is different from the referencing table, you do not need to specify the referenced column; the default column is the primary-key column (or columns) of the referenced table. If the referenced table is the same as the referencing table, you must specify the referenced column.

The following example creates a new column in the **cust_calls** table, **ref_order**. The **ref_order** column is a foreign key that references the **order_num** column in the **orders** table.

```
ALTER TABLE cust_calls
  ADD ref_order INTEGER
  REFERENCES orders (order_num)
  BEFORE user_id;
```

When you place a referential constraint on a column or set of columns, and a duplicate or unique index already exists on that column or set of columns, the index is shared.

Using the ON DELETE CASCADE Option

Use the ON DELETE CASCADE option if you want rows deleted in the child table when corresponding rows are deleted in the parent table. If you do not specify cascading deletes, the default behavior of the database server prevents you from deleting data in a table if other tables reference it.

If you specify this option, when you delete a row in the parent table, the database server also deletes any rows associated with that row (foreign keys) in a child table. The advantage of the ON DELETE CASCADE option is that it allows you to reduce the quantity of SQL statements needed to perform delete actions.

For example, in the **stores_demo** database, the **stock** table contains the **stock_num** column as a primary key. The **catalog** table refers to the **stock_num** column as a foreign key. The following ALTER TABLE statements drop an existing foreign-key constraint (without cascading delete) and add a new constraint that specifies cascading deletes:

```
ALTER TABLE catalog DROP CONSTRAINT aa;

ALTER TABLE catalog ADD CONSTRAINT
  (FOREIGN KEY (stock_num, manu_code) REFERENCES stock
  ON DELETE CASCADE CONSTRAINT ab);
```

With cascading deletes specified on the child table, in addition to deleting a stock item from the **stock** table, the delete cascades to the catalog table that is associated with the **stock_num** foreign key. This cascading delete works only if the **stock_num** that you are deleting was not ordered; otherwise, the constraint from the **items** table would disallow the cascading delete. For more information, see “Restrictions on DELETE When Tables Have Cascading Deletes” on page 2-305.

If a table has a trigger with a DELETE trigger event, you cannot define a cascading-delete referential constraint on that table. You receive an error when you attempt to add a referential constraint that specifies ON DELETE CASCADE to a table that has a delete trigger.

For information about syntax restrictions and locking implications when you delete rows from tables that have cascading deletes, see “Considerations When Tables Have Cascading Deletes” on page 2-305.

Locks Held During Creation of a Referential Constraint

When you create a referential constraint, the database server places an exclusive lock on the referenced table. The lock is released after you finish with the ALTER TABLE statement or at the end of a transaction (if you are altering the table in a database that uses transaction logging).

CHECK Clause

A check constraint designates a condition that must be met *before* data can be inserted into a column.

CHECK Clause:

```
|—CHECK—(—| Condition |———)———|
```

Notes:

- 1 See “Condition ” on page 4-5

During an insert or update, if a row returns *false* for any check constraint defined on a table, the database server returns an error. No error is returned, however, if a row returns NULL for a check constraint. In some cases, you might want to use both a check constraint and a NOT NULL constraint.

Check constraints are defined using *search conditions*. The search condition cannot contain user-defined routines, subqueries, aggregates, host variables, or rowids. In addition, the condition cannot contain the variant built-in functions CURRENT, SYSDATE, USER, SITENAME, DBSERVERNAME, or TODAY.

The check constraint cannot include columns in different tables. When you are using the ADD or MODIFY clause, the check constraint cannot depend upon values in other columns of the same table.

The next example adds a new **unit_price** column to the **items** table and includes a check constraint to ensure that the entered value is greater than 0:

```
ALTER TABLE items
  ADD (unit_price MONEY (6,2) CHECK (unit_price > 0));
```

To create a constraint that checks values in more than one column, use the ADD CONSTRAINT clause. The following example builds a constraint on the column that was added in the previous example. The check constraint now spans two columns in the table.

```
ALTER TABLE items ADD CONSTRAINT CHECK (unit_price < total_price);
```

Add Column Security

The Add Column Security clause associates the new column with a security label.

Add Column Security Clause:



Element	Description	Restrictions	Syntax
<i>label</i>	Name of a security label	Must exist and must belong to the same security policy that protects the table.	"Identifier " on page 5-21

Specify the *label* name without the policy qualifier, rather than as *policy.label*, because the security policy of the table is the only valid policy for any security label that protects data in the table.

The column cannot be of type IDSSECURITYLABEL

When a user who holds appropriate table access privileges attempts to access a value in the protected column, the database server compares this label with the security credentials of the user, and allows or withholds access on the basis of this comparison.

DROP Clause

Use the DROP clause to drop one or more columns from a table.

DROP Clause:



Element	Description	Restrictions	Syntax
<i>column</i>	Name of column to be dropped	Must exist in the table. No fragment expression can reference the column, and it cannot be the last column in the table.	"Identifier " on page 5-21

You cannot issue an ALTER TABLE DROP statement that would drop every column from the table. At least one column must remain in the table.

You cannot drop a column that is part of a fragmentation strategy.

A column that is protected by a security label can be dropped by the ALTER TABLE DROP statement, but the user must be DBSECADM and must also hold the usual CONNECT, RESOURCE, and ALTER access privileges for dropping columns.

In Extended Parallel Server, you cannot use the DROP clause if the table has a dependent GK index.

How Dropping a Column Affects Constraints

When you drop a column, all constraints on that column are also dropped:

- All single-column constraints are dropped.
- All referential constraints that reference the column are dropped.
- All check constraints that reference the column are dropped.
- If the column is part of a multiple-column primary-key or unique constraint, the constraints placed on the multiple columns are also dropped. This action, in turn, triggers the dropping of all referential constraints that reference the multiple columns.

Because any constraints that are associated with a column are dropped when the column is dropped, the structure of other tables might also be altered when you use this clause. For example, if the dropped column is a unique or primary key that is referenced in other tables, those referential constraints also are dropped. Therefore the structure of those other tables is also altered.

How Dropping a Column Affects Triggers

In general, when you drop a column from a table, the triggers based on that table remain unchanged. If the column that you drop appears in the action clause of a trigger, however, dropping the column can invalidate the trigger. The following statements illustrate the possible effects on triggers:

```
CREATE TABLE tab1 (i1 int, i2 int, i3 int);
CREATE TABLE tab2 (i4 int, i5 int);
CREATE TRIGGER col1trig UPDATE OF i2 ON tab1
  BEFORE(INSERT INTO tab2 VALUES(1,1));
ALTER TABLE tab2 DROP i4;
```

After the ALTER TABLE statement, **tab2** has only one column. The **col1trig** trigger is invalidated because the action clause as it is currently defined with values for two columns cannot occur.

If you drop a column that occurs in the triggering column list of an UPDATE trigger, the database server drops the column from the triggering column list. If the column is the only member of the triggering column list, the database server drops the trigger from the table. For more information on triggering columns in an UPDATE trigger, see “CREATE TRIGGER” on page 2-239.

If a trigger is invalidated when you alter the underlying table, drop and then re-create the trigger.

How Dropping a Column Affects Views

When you drop a column from a table, the views based on that table remain unchanged. That is, the database server does not automatically drop the corresponding columns from associated views.

The view is not automatically dropped because ALTER TABLE can change the order of columns in a table by dropping a column and then adding a new column with the same name. In this case, views based on the altered table continue to work, but retain their original sequence of columns.

If a view is invalidated when you alter the underlying table, you must rebuild the view by using the DROP VIEW and CREATE VIEW statements.

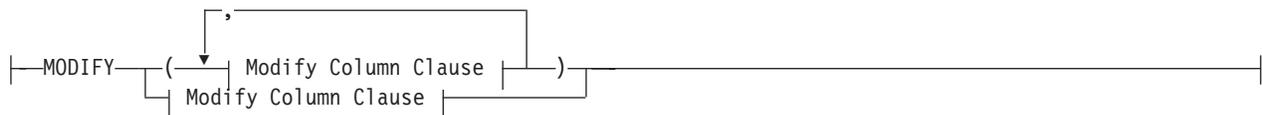
How Dropping a Column Affects a Generalized-Key Index

In Extended Parallel Server, if you drop a column from a table that has a dependent GK index, all GK indexes on the table that refer to the dropped column are dropped. Any GK indexes on other tables that refer to the dropped column are also dropped.

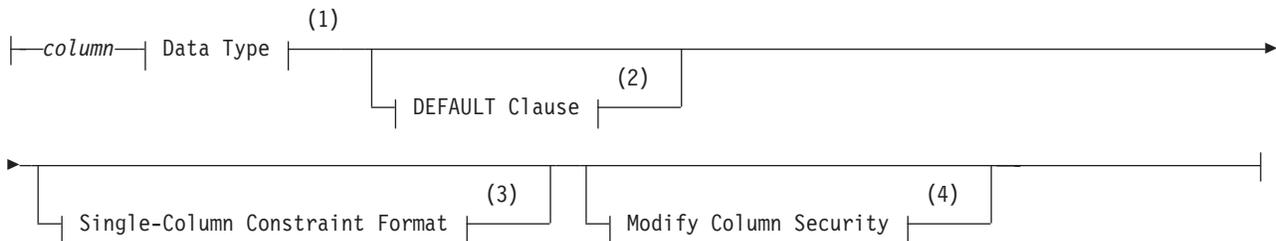
MODIFY Clause

Use the MODIFY clause to change the data type, length, or default value of a column, to add or remove the security label of a column, to allow or disallow NULL values in a column, or to reset the serial counter of a SERIAL or SERIAL8 column.

MODIFY Clause:



Modify Column Clause:



Notes:

- 1 See "Data Type " on page 4-21
- 2 See "DEFAULT Clause" on page 2-49
- 3 See "Single-Column Constraint Format" on page 2-50
- 4 See "Modify Column Security" on page 2-60

Element	Description	Restrictions	Syntax
<i>column</i>	Column to modify	Must exist in table. Cannot be a collection or IDSECURITYLABEL data type.	"Identifier " on page 5-21

In Extended Parallel Server, you cannot use the MODIFY clause if the table has a dependent GK index.

You cannot change the data type of a column to a COLLECTION or a ROW type.

The IDSECURITYLABEL column of a protected table cannot be altered to a different data type, nor can an existing column be altered to be of type IDSECURITYLABEL.

When you modify a column, *all* attributes previously associated with the column (that is, default value, single-column check constraint, or referential constraint) are dropped. When you want certain attributes of the column to remain, such as PRIMARY KEY, you must re-specify those attributes.

For example, if you are changing the data type of an existing column, **quantity**, to SMALLINT, but you want to keep the default value (in this case, 1) and the NOT NULL column attribute, you can issue this statement:

```
ALTER TABLE items MODIFY (quantity SMALLINT DEFAULT 1 NOT NULL);
```

Tip: Both attributes are specified *again* in the modify clause.

When you change the data type of a column, the database server does not perform the modification in place. The next example (for Dynamic Server only) changes a VARCHAR(15) column to an LVARCHAR(3072) column:

```
ALTER TABLE stock MODIFY (description LVARCHAR(3072));
```

When you modify a column that has column constraints associated with it, the following constraints are dropped:

- All single-column constraints are dropped.
- All referential constraints that reference the column are dropped.
- If the modified column is part of a multiple-column primary-key or unique constraint, all referential constraints that reference the multiple columns also are dropped.

For example, if you modify a column that has a unique constraint, the unique constraint is dropped. If this column was referenced by columns in other tables, those referential constraints are also dropped. In addition, if the column is part of a multiple-column primary-key or unique constraint, the multiple-column constraints are not dropped, but any referential constraints placed on the column by other tables are dropped.

For another example, suppose that a column is part of a multiple-column primary-key constraint. This primary key is referenced by foreign keys in two other tables. When this column is modified, the multiple-column primary-key constraint is not dropped, but the referential constraints placed on it by the two other tables are dropped.

Using the MODIFY Clause

The characteristics of the object you are attempting to modify can affect how you handle your modifications.

Altering BYTE and TEXT Columns

You can use the MODIFY clause to change a BYTE column to a TEXT column, and vice versa. In Dynamic Server, you can also use the MODIFY clause to change a BYTE column to a BLOB column and a TEXT column to a CLOB column. Except for these operations, however, you cannot use the MODIFY clause to change a BYTE or TEXT column to any other type of column, nor to change any other type of column to a BYTE or TEXT column.

Altering the Next Serial Value

You can use the MODIFY clause to reset the next value of a SERIAL, BIGSERIAL, or SERIAL8 column. You cannot set the next value below the current maximum value in the column because that action can cause the database server to generate

duplicate numbers. You can set the next value, however, to any value higher than the current maximum, which creates a gap in the series of values.

If the new serial value that you specify is less than the current maximum value in the serial column, the maximum value is not altered. If the maximum value is less than what you specify, the next serial number will be what you specify. The next serial value is not equivalent to one greater than the maximum serial value in the column in two situations:

- There are no rows in the table, and an initial serial value was specified when the table was created (or by a previous ALTER TABLE statement).
- There are rows in the table, but the next serial value was modified by a previous ALTER TABLE statement.

The following example sets the next serial value to 1000:

```
ALTER TABLE my_table MODIFY (serial_num SERIAL (1000));
```

As an alternative, you can use the INSERT statement to create a gap in the series of serial values in the column. For more information, see “Inserting Values into Serial Columns” on page 2-436.

Altering the Next Serial Value in a Typed Table (IDS): You can set the initial serial number or modify the next serial number for a ROW-type field with the MODIFY clause of the ALTER TABLE statement. (You cannot set the initial number for a serial field when you create a ROW data type.)

Suppose you have ROW types **parent**, **child1**, **child2**, and **child3**.

```
CREATE ROW TYPE parent (a int);
CREATE ROW TYPE child1 (s serial) UNDER parent;
CREATE ROW TYPE child2 (b float, s8 serial8) UNDER child1;
CREATE ROW TYPE child3 (d int) UNDER child2;
```

You then create corresponding typed tables:

```
CREATE TABLE OF TYPE parent;
CREATE TABLE OF TYPE child1 UNDER parent;
CREATE TABLE OF TYPE child2 UNDER child1;
CREATE TABLE OF TYPE child3 UNDER child2;
```

To change the next SERIAL and SERIAL8 numbers to 75, you can issue the following statement:

```
ALTER TABLE child3 MODIFY (s serial(75), s8 serial8(75));
```

When the ALTER TABLE statement executes, the database server updates corresponding serial columns in the **child1**, **child2**, and **child3** tables.

Altering the Structure of Tables

When you use the MODIFY clause, you can also alter the structure of other tables. If the modified column is referenced by other tables, those referential constraints are dropped. You must add those constraints to the referencing tables again, using the ALTER TABLE statement.

When you change the data type of an existing column, all data is converted to the new data type, including numbers to characters and characters to numbers (if the characters represent numbers). The following statement changes the data type of the **quantity** column:

```
ALTER TABLE items MODIFY (quantity CHAR(6));
```

When a primary-key or unique constraint exists, however, conversion takes place only if it does not violate the constraint. If a data type conversion would result in duplicate values (by changing FLOAT to SMALLFLOAT, for example, or by truncating CHAR values), the ALTER TABLE statement fails.

Modifying Tables for NULL Values

You can modify an existing column that formerly permitted NULLs to disallow NULLs, provided that the column contains no NULL values. To do this, specify MODIFY with the same column name and data type and the NOT NULL keywords. Those keywords create a NOT NULL constraint on the column.

You can modify an existing column that did not permit NULLs to permit NULLs. To do this, specify MODIFY with the column name and the existing data type, and omit the NOT NULL keywords. The omission of the NOT NULL keywords drops the NOT NULL constraint on the column. If a unique index exists on the column, you can remove it using the DROP INDEX statement.

An alternative method of permitting NULL values in an existing column that did not permit NULL values is to use the DROP CONSTRAINT clause to drop the NOT NULL constraint on the column.

Adding a Constraint on a Non-Opaque Column

ALTER TABLE ... MODIFY operations that use the Single Column Constraint format to implicitly create an index on a non-opaque column also automatically calculate the distribution of the specified column. The distribution statistics are available to the query optimizer when it designs query plans for the table on which the constraint is defined:

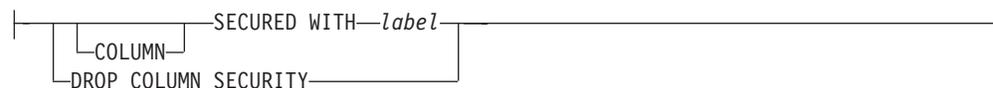
- For columns on which the new constraint is implemented as a B-tree index, the recalculated column distribution statistics are equivalent to distributions created by the UPDATE STATISTICS statement in HIGH mode.
- If the new constraint is not implemented as a B-tree index, the automatically recalculated statistics correspond to distributions created by the UPDATE STATISTICS statement in LOW mode.

See also the section “Automatic Calculation of Distribution Statistics” on page 2-144 in the description of the CREATE INDEX statement for additional information about statistical distributions that are produced automatically when an index or constraint is created on an existing table.

Modify Column Security

The Modify Column Security clause is valid only for tables that are protected by a security policy. Use this clause to add or drop row-level security for the column.

Modify Column Security Clause:



Element	Description	Restrictions	Syntax
<i>label</i>	Name of a security label	Must exist and must belong to the same security policy that protects the table.	“Identifier ” on page 5-21

This clause can add or drop row-level protection.

- To drop the row-level protection of a column, specify the DROP COLUMN SECURITY keywords.
- To provide row-level protection to a column, specify SECURED WITH *label* (or equivalently, COLUMN SECURED WITH *label*).

The security label can be the same label that protects other rows or columns of the table, or it can be a different label of the same security policy. The following restrictions apply to the SECURED WITH label option:

- The column cannot be of type IDSSECURITYLABEL.
- Specify the *label* without the policy qualifier, rather than as *policy.label*.
- The *label* must be a label of the security policy that secures the table.

Adding a Constraint That Existing Rows Violate (IDS)

If you use the MODIFY clause to add a constraint in the enabled mode and receive an error message because existing rows would violate the constraint, take the following steps to add the constraint successfully:

1. Add the constraint in the disabled mode.
Issue the ALTER TABLE statement again, but this time specify the DISABLED keyword in the MODIFY clause.
2. Start a violations and diagnostics table for the target table with the START VIOLATIONS TABLE statement.
3. Issue the SET CONSTRAINTS statement to switch the database object mode of the constraint to the enabled mode.
When you issue this statement, existing rows in the target table that violate the constraint are duplicated in the violations table; however, you receive an integrity-violation error message, and the constraint remains disabled.
4. Issue a SELECT statement on the violations table to retrieve the nonconforming rows that are duplicated from the target table.
You might need to join the violations and diagnostics tables to get all the necessary information.
5. Take corrective action on the rows in the target table that violate the constraint.
6. After you fix all the nonconforming rows in the target table, issue the SET statement again to enable the constraint that was disabled.
Now the constraint is enabled, and no integrity-violation error message is returned because all rows in the target table now satisfy the new constraint.

How Modifying a Column Affects a Generalized-Key Index (XPS)

In Extended Parallel Server, when you modify a column, all GK indexes that reference the column are dropped if the column is used in the GK index in a way that is incompatible with the new data type of the column.

For example, if a numeric column is changed to a character column, any GK indexes involving that column are dropped if they involve arithmetic expressions.

How Modifying a Column Affects Triggers

If you modify a column that appears in the triggering column list of an UPDATE trigger, the trigger is unchanged.

When you modify a column in a table, the triggers based on that table remain unchanged, but the column modification might invalidate the trigger.

When you do not specify a name for a new constraint, the database server provides one. You can find the name of the constraint in the **sysconstraints** system catalog table. For more information about the **sysconstraints** system catalog table, see the *IBM Informix Guide to SQL: Reference*.

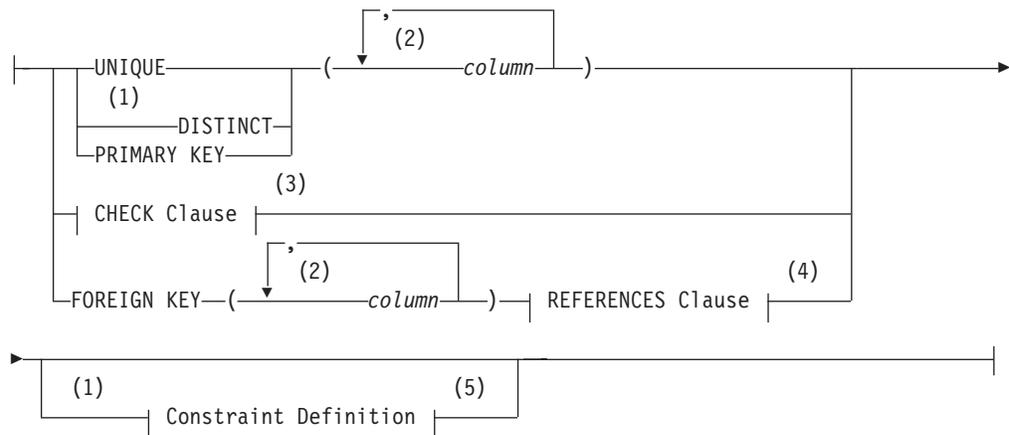
The following restrictions on the ADD CONSTRAINT clause affect columns of certain data types:

- You cannot place a unique constraint nor referential constraints on a BYTE or TEXT column.
- A check constraint on a BYTE or TEXT column can check only for IS NULL, IS NOT NULL, or LENGTH.
- An IDSECURITYLABEL column has an implicit NOT NULL constraint by default, but you cannot use the ADD CONSTRAINT clause to reference an IDSECURITYLABEL column in the definition of a single-column constraints nor as part of a multiple-column referential constraint or check constraint.
- When you add a constraint, the collating order must be the same as when the table was created.

Multiple-Column Constraint Format

Use the Multiple-Column Constraint Format option to assign one or more constraints to one column or to a set of columns.

Multiple-Column Constraint Format:



Notes:

- 1 Informix extension
- 2 Use path no more than 16 times
- 3 See "CHECK Clause" on page 2-54
- 4 See "REFERENCES Clause" on page 2-52
- 5 See "Constraint Definition" on page 2-51

Element	Description	Restrictions	Syntax
<i>column</i>	A column on which the constraint is placed	No more than 16 columns	"Identifier" on page 5-21

As in the CREATE TABLE statement, the Multiple-Column Constraint format for ALTER TABLE differs from the Single-Column Constraint format by requiring the FOREIGN KEY keywords before the REFERENCES clause when you specify a foreign key constraint. In addition, as its name implies, the Multiple-Column format can specify a list of columns as the scope of the new constraint, but this syntax is also valid with a single column.

A multiple-column constraint has these restrictions:

- It can include no more than 16 column names.
- In Dynamic Server, the total length of the list of columns cannot exceed 390 bytes.
- In Extended Parallel Server, the total length of the list of columns cannot exceed 255 bytes.

If the constraint is on a set of columns that includes a column that stores encrypted data, Dynamic Server cannot enforce the constraint. You can declare a name for the constraint and set its mode by means of “Constraint Definition” on page 2-51.

ALTER TABLE ... ADD CONSTRAINT operations that implicitly create an index on a non-opaque column or on a set of columns automatically calculate distributions on the specified column, or on the lead column of a multiple-column constraint.

These distribution statistics are equivalent to distributions created by the UPDATE STATISTICS statement in HIGH mode, and are available to the query optimizer when it designs query plans for the table on which the new constraint was created. See also the section “Automatic Calculation of Distribution Statistics” on page 2-144 in the description of the CREATE INDEX statement for additional information about statistical distributions that are produced when an index or constraint is created on an existing table.

Adding a Primary-Key or Unique Constraint

When you place a primary-key or unique constraint on a column or set of columns, those columns must contain unique values. The database server checks for existing constraints and indexes:

- If a user-created unique index already exists on that column or set of columns, the constraint shares the index.
- If a user-created index that allows duplicates already exists on that column or set of columns, the database server returns an error.

In this case, you must drop the existing index before adding the primary-key or unique constraint.

- If a referential constraint already exists on that column or set of columns, the duplicate index is upgraded to unique (if possible) and the index is shared.
- If no referential constraint or user-created index exists on that column or set of columns, the database server creates an internal B-tree index on the specified columns.

When you place a referential constraint on a column or set of columns, and an index already exists on that column or set of columns, the index is shared.

If you own the table or have the Alter privilege on the table, you can create a check, primary-key, or unique constraint on the table and specify yourself as the owner of the constraint. To add a referential constraint, you must have the

References privilege on either the referenced columns or the referenced table. When you have the DBA privilege, you can create constraints for other users.

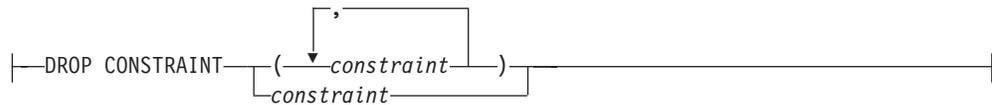
Recovery from Constraint Violations (IDS)

If you use the ADD CONSTRAINT clause to add a constraint in the enabled mode, you receive an error message because existing rows would violate the constraint. For a procedure to add the constraint successfully, see “Adding a Constraint That Existing Rows Violate (IDS)” on page 2-61.

DROP CONSTRAINT Clause

Use the DROP CONSTRAINT clause to drop a named constraint.

DROP CONSTRAINT Clause:



Element	Description	Restrictions	Syntax
<i>constraint</i>	Constraint to be dropped	Must exist.	“Identifier ” on page 5-21

To drop an existing constraint, specify the DROP CONSTRAINT keywords and the name of the constraint. Here is an example of dropping a constraint:

```
ALTER TABLE manufact DROP CONSTRAINT con_name;
```

If no name is specified when the constraint is created, the database server generates the name. You can query the **sysconstraints** system catalog table for the name and owner of a constraint. For example, to find the name of the constraint placed on the **items** table, you can issue the following statement:

```
SELECT constrname FROM sysconstraints
WHERE tabid = (SELECT tabid FROM systables
WHERE tabname = 'items');
```

When you drop a primary-key or unique constraint that has a corresponding foreign key, the referential constraints are dropped. For example, if you drop the primary-key constraint on the **order_num** column in the **orders** table and **order_num** exists in the **items** table as a foreign key, that referential relationship is also dropped.

By default, every IDSSECURITYLABEL column has an implicit NOT NULL constraint, but the DROP CONSTRAINT clause cannot reference an IDSSECURITYLABEL column.

MODIFY EXTENT SIZE

Use the MODIFY EXTENT SIZE clause with the ALTER TABLE statement to change the size of the first extent of a table in a dbspace.

You cannot use the MODIFY EXTENT SIZE clause to change the size of the first extent:

- of a table in a blob space
- of external tables, virtual tables, or system catalog tables

- in the tbspace **tbspace**

MODIFY EXTENT SIZE Clause:

|—MODIFY EXTENT SIZE—*kilobytes*—|

Element	Description	Restrictions	Syntax
<i>kilobytes</i>	Length (in kilobytes) assigned here to the first extent for this table	Specification cannot be a variable, and (4(page size)) ≤ <i>kilobytes</i> ≤ (chunk size)	“Expression” on page 4-40

The minimum extent size is 4 times the disk-page size. For example, on a system with 2-kilobyte pages, the minimum length is 8 kilobytes. The maximum length is the chunk size.

The following example specifies an extent size of 32 kilobytes:

```
ALTER TABLE customer MODIFY EXTENT SIZE 32;
```

When you change the size of the first extent, the database server records the change in the system catalog and on the partition page, but only makes the actual change when the table is rebuilt or a new partition or fragment is created.

For example, if a table has a first extent size of 8 kilobytes and you use the ALTER TABLE statement to change this to 16 kilobytes, the server does not drop the current first extent and recreate it with the new size. Instead, the new first extent size of 16 kilobytes takes effect only when the server rebuilds the table after actions such as creating a cluster index on the table or detaching a fragment from the table.

If a TRUNCATE TABLE statement without the REUSE option is executed before the ALTER TABLE statement with the MODIFY EXTENT SIZE clause, there is no change in the size of the current first extent.

If an existing table in a dbspace has data in it, the first and next extents are already allocated for the table and you will not be able to change the size of the first or next extent. If you want to change the size of existing extents, you must drop the table, recreate it with a storage clause indicating the desired size, and load the data again.

You can change the size of the first and next extent at the same time. The following example specifies changing the size of the first and next extent:

```
ALTER TABLE customer MODIFY EXTENT SIZE 32 NEXT SIZE 32
```

The first and next extent sizes are recorded in the PNSIZES logical log record.

MODIFY NEXT SIZE Clause

Use the MODIFY NEXT SIZE clause to change the size of the next extent.

MODIFY NEXT SIZE Clause:

|—MODIFY NEXT SIZE—*kilobytes*—|

Element	Description	Restrictions	Syntax
<i>kilobytes</i>	Length (in kilobytes) assigned here to the next extent for this table	Specification cannot be a variable, and $(4(\text{page size}) \leq \text{kilobytes} \leq (\text{chunk size}))$	"Expression" on page 4-40

The minimum extent size is 4 times the disk-page size. For example, on a system with 2-kilobyte pages, the minimum length is 8 kilobytes. The maximum length is the chunk size. The following example specifies an extent size of 32 kilobytes:

```
ALTER TABLE customer MODIFY NEXT SIZE 32;
```

This clause cannot change the size of existing extents. You cannot change the size of existing extents without unloading all of the data.

You cannot use the MODIFY NEXT SIZE clause of the ALTER TABLE statement to change the size of the next extent of any system catalog table, even if you are user **informix**.

To change the size of existing extents, you must unload all the data, drop the table, modify the *first-extent* and *next-extent* sizes in the CREATE TABLE definition in the database schema, re-create the table, and reload the data. For information about how to optimize extent sizes, see your *IBM Informix Performance Guide*.

LOCK MODE Clause

Use the LOCK MODE keywords to change the locking granularity of a table.

LOCK MODE Clause:



Notes:

- 1 Extended Parallel Server only

The following table describes the locking-granularity options available.

Granularity	Effect
-------------	--------

PAGE

Obtains and releases one lock on a whole page of rows

This is the default locking granularity. Page-level locking is especially useful when you know that the rows are grouped into pages in the same order that you are using to process all the rows. For example, if you are processing the contents of a table in the same order as its cluster index, page locking is especially appropriate.

ROW

Obtains and releases one lock per row

Row-level locking provides the highest level of concurrency. Only tables with row-level locking support the LAST COMMITTED feature, which can improve performance in the Committed Read and Dirty Read isolation levels when another session holds an exclusive lock on a row that you attempt to read. If you are using many rows at one time, however, the

lock-management overhead of row-level locking can become significant. You can also exceed the maximum number of locks available, depending on the configuration of your database server.

TABLE (XPS only)

Places a lock on the entire table

This type of lock reduces update concurrency in comparison to row and page locks. A table lock reduces the lock-management overhead for a table. Multiple read-only transactions can still access the table.

The following statement changes the lock mode for the **customer** table to page level locking:

```
ALTER TABLE customer LOCK MODE(page);
```

The next example changes the lock mode for the **customer** table to row level locking:

```
ALTER TABLE customer LOCK MODE(row);
```

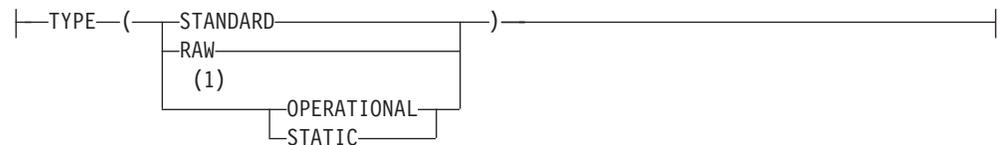
Precedence and Default Behavior

The LOCK MODE setting in an ALTER TABLE statement takes precedence over the settings of the **IFX_DEF_TABLE_LOCKMODE** environment variable and the DEF_TABLE_LOCKMODE configuration parameter. For information about the **IFX_DEF_TABLE_LOCKMODE** environment variable, refer to the *IBM Informix Guide to SQL: Reference*. For information about the DEF_TABLE_LOCKMODE configuration parameter, refer to the *IBM Informix Dynamic Server Administrator's Reference*.

Logging TYPE Options

Use the Logging TYPE options to specify that the table have particular characteristics that can improve various bulk operations on it.

Logging TYPE Options:



Notes:

- 1 Extended Parallel Server only

Here STANDARD, the default option of the CREATE TABLE statement, is used for online transaction processing (OLTP) databases. The OPERATIONAL and STATIC options are used primarily to improve performance in data warehousing databases.

A table can have any of the following logging characteristics.

Option Effect

STANDARD

Logging table that allows rollback, recovery, and restoration from archives. This is the default. Use this type for recovery and constraints functionality on OLTP databases.

RAW Nonlogging table that do not support primary key constraints or unique constraints, but that support referential constraints, and can be indexed and updated. Use this type for quickly loading data. In XPS, raw tables take advantage of light appends and avoid the overhead of logging and checking constraints.

OPERATIONAL (XPS only)

Logging table that uses light appends and cannot be restored from archive. Use this type on tables that are refreshed frequently. Light appends allow the quick addition of many rows.

STATIC (XPS only)

Nonlogging table that can contain index and referential constraints but cannot be updated. Use this type for read-only operations because there is no logging or locking overhead.

Warning: Use raw tables for fast loading of data. It is recommended that you alter the logging type to STANDARD and perform a level-0 backup before you use the table in a transaction or modify the data within the table. If you must use a raw table within a transaction, either set the isolation level to Repeatable Read or lock the table in exclusive mode to prevent concurrency problems.

The Logging TYPE option can convert a non-logging table, such as a RAW table, to a STANDARD table that supports transaction logging. If you use this feature, you should be aware that the database server does not check to see whether a level 0 archive has been performed on the table.

Operations on a RAW table are not logged and are not recoverable, so RAW tables are always at risk. When the database server converts a table that is not logged to a STANDARD table type, it is your responsibility to perform a level-0 backup before using the table in a transaction or modifying data in the table. Failure to do this might lead to recovery problems in the event of a server crash.

For more information on these logging types of tables, refer to your *IBM Informix Administrator's Guide*.

The Logging TYPE options have the following restrictions:

- You must perform a level-0 archive before the logging type of a table can be altered to STANDARD from any other logging type.
- If you have triggers defined on the table, you cannot change the logging type to RAW or STATIC. These nonlogging table types do not support triggers.
- The table cannot be a SCRATCH or TEMP table, and you cannot change any of these types of tables to a SCRATCH or TEMP table.
- The table cannot have a dependent generalized-key (GK) index.

The following example changes a nonlogging table to a table that uses transaction logging:

```
ALTER TABLE tabnolog TYPE (STANDARD);
```

The following example changes a logging table to a nonlogging table.

```
ALTER TABLE tablog TYPE (RAW);
```

SECURITY POLICY Clause

The optional Security Policy clause can use the following syntax to drop the security policy that is currently associated with the table, or to associate a security policy with a table that has none.

SECURITY POLICY Clause:



Element	Description	Restrictions	Syntax
<i>policy</i>	Name of a security policy	Must be the security policy that protects the table	"Identifier " on page 5-21

Only DBSECADM can use this clause to add a security policy to an existing table, or to remove from the table the protection of the security policy that currently protects a table.

The following guidelines apply to tables that can be protected by executing the ADD SECURITY POLICY clause of the ALTER TABLE statement:

- A table is not protected unless it has a security policy associated with it and has either rows secured or at least one column secured. The former indicates that the table is a protected table with row level granularity and the latter indicates that the table is a protected table with column-level granularity.
- Securing rows by using the ALTER TABLE ADD statement to add an IDSSECURITYLABEL column to an existing table fails if the table does not have a security policy associated with it.
- Securing a column with the ALTER TABLE MODIFY ... COLUMN SECURED WITH clause fails if the table does not have a security policy associated with it.
- A table can have at most one security policy. The ADD SECURITY POLICY clause fails if the table already has a security policy.
- A table can have any number of protected columns, and each protected column can have a different security label, or several protected columns can share the same security label.
- A security policy cannot be added to a temporary table, a typed table, a table in a table hierarchy, or an external table outside the current database.
- A table can have at most one column of type IDSSECURITYLABEL.
- The IDSSECURITYLABEL column cannot have column protection.
- The IDSSECURITY LABEL column cannot have single column constraints nor be part of multiple column referential or check constraints.
- The IDSSECURITYLABEL column cannot be encrypted.
- The IDSSECURITYLABEL column has an implicit DEFAULT NOT NULL constraint. The default column value is the value of the security label of the user for write access.
- The IDSSECURITYLABEL column can be dropped only by DBSECADM, who must also hold the usual CONNECT, RESOURCE, and ALTER access privileges for dropping columns.
- The IDSSECURITYLABEL column cannot be modified by the ALTER TABLE statement.
- Attaching a fragment to a protected fragmented table fails if any of these conditions are true:
 - if the source table and the target table are not protected using the same security policy;
 - if the tables do not have the same protection granularity;

- if the tables do not have the same set of protected columns, each protected by the same security label.

For more information on using the ALTER FRAGMENT statement to attach fragments to protected tables, see “Additional Restrictions on the ATTACH Clause in Dynamic Server” on page 2-11.

- Detaching a fragment of a protected table creates a new table that is protected by the same security policy for the same row security label column, and the same set of protected columns.

If the DROP SECURITY POLICY clause executes successfully, it has the following effects:

- A table is not protected unless it has a security policy associated with it and has either rows secured or at least one column secured. The former indicates that the table is a protected table with row level granularity and the latter indicates that the table is a protected table with column level granularity.
- Securing rows with the IDSSECURITYLABEL column clause fails if the table does not have a security policy associated with it.
- Securing a column with the COLUMN SECURED WITH clause of the ALTER TABLE ADD or ALTER TABLE MODIFY statement fails if the table does not have a security policy associated with it.
- When a security policy is dropped from a table by the ALTER TABLE DROP SECURITY POLICY statement, the IDSSECURITYLABEL column is automatically dropped. If the table has one or more protected columns, those columns become unprotected.

Do not confuse this DROP SECURITY POLICY clause of the ALTER table statement with the DROP SECURITY POLICY statement.

- When the DROP SECURITY POLICY clause of the ALTER TABLE statement executes successfully, it terminates the association of the table with the security policy, drops the IDSSECURITYLABEL column, and removes LBAC protection from data that had been protected in that table. It has no effect, however, on the security policy, nor on other tables protected by the policy.
- When the DROP SECURITY POLICY statement executes successfully, the effects depend on whether the policy is dropped in RESTRICT or CASCADE mode, but in either mode, it destroys the specified policy. See the description of “DROP SECURITY ” on page 2-338 for more information on the DROP SECURITY POLICY statement of SQL, and about restrictions on that statement.

ADD TYPE Clause (IDS)

Use the ADD TYPE clause to convert a table that is not based on a named ROW data type into a typed table. This clause is an extension to the ANSI/ISO standard for SQL.

ADD TYPE Clause:

|—ADD TYPE—*row_type*—|

Element	Description	Restrictions	Syntax
<i>row_type</i>	Identifier of an existing named ROW data type for the table	The <i>row_type</i> fields must match the column data type in their order and number	“Identifier ” on page 5-21

When you use the ADD TYPE clause, you assign the specified named ROW data type to a table whose columns match the fields of that data type.

In addition to the requirements common to all ALTER TABLE operations (namely DBA privilege on the database, Alter privilege on the table, and ownership of the table), all of the following must be also true when you use the ADD TYPE clause to convert an untyped table to the specified named ROW data type:

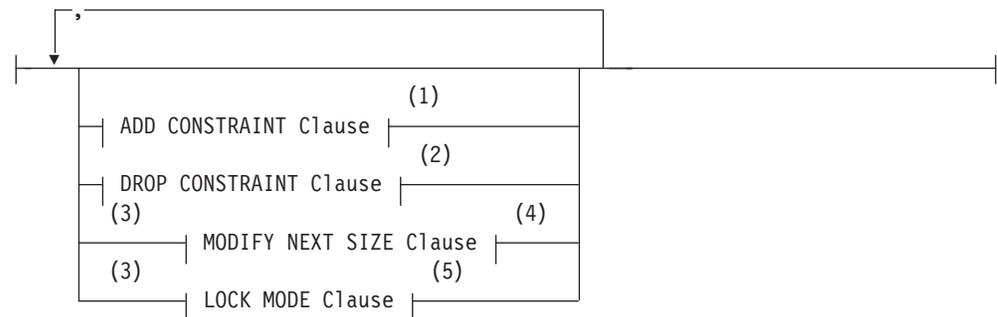
- The named ROW data type is already registered in the database.
- You hold the Usage privilege on the named ROW data type.
- There must be a 1-to-1 correspondence between the ordered set of column data types of the untyped table and the ordered set of field data types of the named ROW data type.
- The table cannot be a fragmented table that has **rowid** values.

You cannot combine the ADD TYPE clause with any clause that changes the schema of the table. No other ADD, DROP, or MODIFY clause is valid in the same ALTER TABLE statement that has the ADD TYPE clause. The ADD TYPE clause does not allow you to change column data types. (To change the data type of a column, use the MODIFY clause.)

Options Valid on Typed Tables

ALTER TABLE supports only the following options for tables of ROW data types.

Typed-Table Options:



Notes:

- 1 See "ADD CONSTRAINT Clause" on page 2-63
- 2 See "DROP CONSTRAINT Clause" on page 2-66
- 3 Use path no more than once
- 4 See "MODIFY NEXT SIZE Clause" on page 2-67
- 5 See "LOCK MODE Clause" on page 2-68

Altering Subtables and Supertables

Two considerations apply to typed tables that are part of inheritance hierarchies:

- For subtables, ADD CONSTRAINT and DROP CONSTRAINT are not valid on inherited constraints.
- For supertables, ADD CONSTRAINT and DROP CONSTRAINT propagate to all subtables.

Related Information

Related statements: "CREATE TABLE " on page 2-191, "DROP TABLE " on page 2-342, "LOCK TABLE " on page 2-449, and "SET Database Object Mode" on page 2-591.

For discussions of data-integrity constraints and the ON DELETE CASCADE option, see the *IBM Informix Guide to SQL: Tutorial*.

For a discussion of database and table creation, see the *IBM Informix Database Design and Implementation Guide*.

For information on how to maximize performance when you make table modifications, see your *IBM Informix Performance Guide*.

BEGIN WORK

Use the BEGIN WORK statement to start a *transaction* (a series of database operations that the COMMIT WORK or ROLLBACK WORK statement terminates).

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 Dynamic Server and ESQL/C only

Usage

The BEGIN WORK statement is valid only in a database that supports transaction logging.

Each row that an UPDATE, DELETE, or INSERT statement affects during a transaction is locked and remains locked throughout the transaction. A transaction that contains many such statements or that contains statements that affect many rows can exceed the limits that your operating system or the database server configuration imposes on the number of simultaneous locks.

If no other user is accessing the table, you can avoid locking limits and reduce locking overhead by locking the table with the LOCK TABLE statement after you begin the transaction. Like other locks, this table lock is released when the transaction terminates. The example of a transaction on "Example of BEGIN WORK" on page 2-75 includes a LOCK TABLE statement.

Important: Issue the BEGIN WORK statement only if a transaction is not in progress. If you issue a BEGIN WORK statement while you are in a transaction, the database server returns an error.

The WORK keyword is optional. The following two statements are equivalent:

```
BEGIN;  
BEGIN WORK;
```

In Informix ESQL/C, if you use the BEGIN WORK statement within a UDR called by a WHENEVER statement, specify WHENEVER SQLERROR CONTINUE and WHENEVER SQLWARNING CONTINUE before the ROLLBACK WORK statement. These statements prevent the program from looping endlessly if the ROLLBACK WORK statement encounters an error or a warning.

BEGIN WORK and ANSI-Compliant Databases

In an ANSI-compliant database, you do not need the BEGIN WORK statement because transactions are implicit; every SQL statement occurs within a transaction. The database server generates a warning when you use a BEGIN WORK statement immediately after any of the following statements:

- DATABASE
- COMMIT WORK
- CREATE DATABASE
- ROLLBACK WORK

The database server returns an error when you use a BEGIN WORK statement after any other statement in an ANSI-compliant database.

BEGIN WORK WITHOUT REPLICATION (IDS, ESQL/C)

When you use Enterprise Replication for data replication, you can use the BEGIN WORK WITHOUT REPLICATION statement to start a transaction that does not replicate to other database servers.

You cannot execute BEGIN WORK WITHOUT REPLICATION as a stand-alone embedded statement in an Informix ESQL/C application. Instead you must execute this statement indirectly. You can use either of the following methods:

- You can use a combination of the PREPARE and EXECUTE statements to prepare and execute the BEGIN WORK WITHOUT REPLICATION statement.
- You can use the EXECUTE IMMEDIATE statement to prepare and execute BEGIN WORK WITHOUT REPLICATION in a single step.

You cannot use the DECLARE *cursor* CURSOR WITH HOLD statement with the BEGIN WORK WITHOUT REPLICATION statement.

For more information about data replication, see the *IBM Informix Dynamic Server Enterprise Replication Guide*.

Example of BEGIN WORK

When consecutive SQL statements perform what is logically a single unit of work, you can define a transaction by grouping them between the BEGIN WORK and COMMIT WORK statements. If the business requirements dictate that either all of the statements be performed successfully, or else that none of them be performed, you can enclose the statements of the transaction between BEGIN WORK to start a transaction and COMMIT WORK to complete the transaction successfully (or ROLLBACK WORK, to cancel the transaction, if the program detects an error).

In the following program fragment, the transaction locks the **stock** table (LOCK TABLE), updates rows in the **stock** table (UPDATE), deletes rows from the **stock**

table (DELETE), and inserts a row into the **manufact** table (INSERT). In this example (with no error handling), the database server executes each of these SQL statements in sequence:

```
BEGIN WORK;
  LOCK TABLE stock;
  UPDATE stock SET unit_price = unit_price * 1.10
    WHERE manu_code = 'KAR';
  DELETE FROM stock WHERE description = 'baseball bat';
  INSERT INTO manufact (manu_code, manu_name, lead_time)
    VALUES ('LYM', 'LYMAN', 14);
COMMIT WORK;
```

Each statement itself is atomic; it either completes successfully or else the database is unchanged afterwards. If any of these statements fail, the other statements will still be executed and the net result is as if the failed statement was never attempted. When the COMMIT WORK statement is executed, the successful changes are made permanent.

Typically, however, transactions are defined with error handling, so that the database server must perform a sequence of operations either completely or not at all. In this case, when you include all of the operations within a single transaction, the database server guarantees that all the statements are completely and perfectly committed to disk, or else the database can be restored to the same state that it was in before the transaction began.

By adding appropriate error handling (for example, by setting the **DBACCNOIGN** environment variable in DB-Acess, or by adding EXEC SQL WHENEVER ERROR STOP in ESQL/C), the transaction can be implicitly rolled back because the program stops on an error without executing COMMIT WORK. More careful conditional coding in a programming language such as ESQL/C allows the programmer to explicitly roll back the transaction while continuing the larger program.

| Error-handling and business logic in applications and UDRs can also delimit one
| or more portions of a transaction by including SAVEPOINT and ROLLBACK TO
| SAVEPOINT statements. If the ROLLBACK TO SAVEPOINT statement is issued
| after an error is encountered, or after the results of part of the transaction indicate
| a conflict with a business rule or with some other criterion, only the changes that
| were made to the database between the ROLLBACK statement and its specified or
| default savepoint are cancelled, rather than the entire transaction. The current
| transaction continues at the statement that follows the ROLLBACK statement, with
| any uncommitted changes to the data or to the schema of the database from
| operations that preceded the savepoint remain pending, until the entire transaction
| is either committed or rolled back. Any locks held by statements that were rolled
| back are retained until the complete transaction ends.

Related Information

Related statements: “COMMIT WORK ” on page 2-81, “ROLLBACK WORK” on page 2-520, “SAVE EXTERNAL DIRECTIVES ” on page 2-522, “SAVEPOINT ” on page 2-526, and “WHENEVER ” on page 2-718.

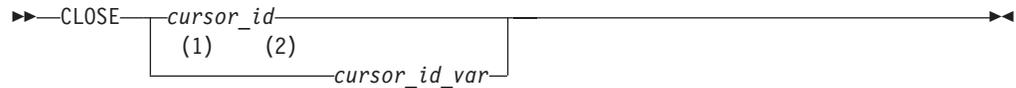
For discussions of transactions and locking, see the *IBM Informix Guide to SQL: Tutorial*.

CLOSE

Use the CLOSE statement when you no longer need to refer to the set of rows associated with a Select cursor or with a Function cursor. With ESQL/C, this statement can also flush and close an Insert cursor.

Use this statement with Informix ESQL/C or SPL.

Syntax



Notes:

- 1 Informix extension
- 2 ESQL/C only

Element	Description	Restrictions	Syntax
<i>cursor_id</i>	Name of cursor to be closed	Must have been declared	"Identifier " on page 5-21
<i>cursor_id_var</i>	Host variable that contains the value of <i>cursor_id</i>	Must be of a character data type	Must conform to language-specific rules for names.

Usage

Closing a cursor makes the cursor unusable in any statements except OPEN or FREE and releases resources that the database server had allocated to the cursor.

In a database that is not ANSI-compliant, you can close a cursor that has not been opened or that has already been closed. No action is taken in these cases.

In an ANSI-compliant database, the database server returns an error if you close a cursor that was not open.

Closing a Select or Function Cursor

When a cursor is associated with a SELECT, EXECUTE FUNCTION, or EXECUTE PROCEDURE statement of SQL, closing the cursor terminates the associated SQL statement.

The database server releases all resources that it might have allocated to the active set of rows, for example, a temporary table that the cursor used to hold an ordered set. The database server also releases any locks that it might have held on rows that were selected through the cursor. If a transaction contains the CLOSE statement, however, the database server does not release the locks until you issue the COMMIT WORK or ROLLBACK WORK statement.

After you close a Select cursor or a Function cursor, the FETCH statement cannot reference that cursor until you reopen it.

In an SPL routine, the built-in SQLCODE function can indicate the result of the CLOSE statement for a Select cursor or a Function cursor. This function returns a

value equivalent to the **SQLCODE** field of the **sqlca** structure. Dynamic Server issues an error, however, if you invoke the built-in **SQLCODE** function outside the calling context of an SPL routine.

Closing an Insert Cursor

Because Dynamic Server does not support Insert cursors in SPL routines, the discussion of Insert cursors in this section applies only to Informix ESQL/C. In SPL routines, the **CLOSE** statement can reference only a Select cursor or a Function cursor that the **DECLARE** statement defined. (A **FOREACH** statement of SPL that has an **INSERT** statement in its statement block can declare a *direct cursor* that functionally resembles an Insert cursor, but the **CLOSE** statement cannot reference a direct cursor that **FOREACH** declared. Dynamic Server closes the direct cursor automatically at runtime when program control exits from the **FOREACH** loop where the direct cursor was defined.)

In Informix ESQL/C, the **CLOSE** statement treats a cursor that is associated with an **INSERT** statement differently from one that is associated with a **SELECT**, **EXECUTE FUNCTION**, or **EXECUTE PROCEDURE** statement. When a cursor identifier is associated with an **INSERT** statement, the **CLOSE** statement writes any remaining buffered rows into the database. The number of rows that were successfully inserted into the database is returned in the third element of the **sqlerrd** array, **sqlca.sqlerrd[2]**, in the **sqlca** structure. For information on how to use **SQLERRD** to count the total number of rows that were inserted, see “Error Checking” on page 2-484.

The **SQLCODE** field of the **sqlca** structure indicates the result of the **CLOSE** statement for an Insert cursor. If all buffered rows are successfully inserted, **SQLCODE** is set to zero. If an error is encountered, the **SQLCODE** field is set to a negative error message number.

When **SQLCODE** is zero, the row buffer space is released, and the cursor is closed; that is, you cannot execute a **PUT** or **FLUSH** statement that names the cursor until you reopen it.

Tip: When you encounter an **sqlca.SQLCODE** error, a corresponding **SQLSTATE** error value also exists. For information about how to get the message text, check the **GET DIAGNOSTICS** statement.

If the insert is not successful, the number of successfully inserted rows is stored in **sqlerrd**. Any buffered rows that follow the last successfully inserted row are discarded. Because the insert fails, the **CLOSE** statement fails also, and the cursor is not closed. For example, a **CLOSE** statement can fail if insufficient disk space prevents some of the rows from being inserted. In this case, a second **CLOSE** statement can be successful because no buffered rows exist. An **OPEN** statement can also be successful because the **OPEN** statement performs an implicit close.

Closing a Collection Cursor

You can declare both Select and Insert cursors on collection variables. Such cursors are called Collection cursors. Use the **CLOSE** statement to deallocate resources that have been allocated for the Collection cursor. Only ESQL/C routines can use **CLOSE** to reference Insert cursors on collection variables. The **CLOSE** statement in SPL routines cannot reference direct Collection cursors that the **FOREACH** statement of SPL can declare.

For more information on how to use a Collection cursor, see “ Fetching from a Collection Cursor (IDS)” on page 2-374 and “Inserting into a Collection Cursor (IDS)” on page 2-483.

Using End of Transaction to Close a Cursor

The COMMIT WORK and ROLLBACK WORK statements close all cursors except those that are declared with a hold. It is better to close all cursors explicitly, however. For Select or Function cursors, this action simply makes the intent of the program clear. It also helps to avoid a logic error if the WITH HOLD clause is later added to the declaration of a cursor.

For an Insert cursor in ESQL/C routines, it is important to use the CLOSE statement explicitly so that you can test the error code. Following the COMMIT WORK statement, **SQLCODE** reflects the result of the COMMIT statement, not the result of closing cursors. If you use a COMMIT WORK statement without first using a CLOSE statement, and if an error occurs while the last buffered rows are being written to the database, the transaction is still committed.

For how to use Insert cursors and the WITH HOLD clause, see “DECLARE ” on page 2-286.

In an ANSI-compliant database, a cursor cannot be closed implicitly. You must issue the CLOSE statement.

Examples

```
EXEC SQL close democursor;
```

The following is ESQL/C Source code example from demo1.ec:

```
#include <stdio.h>

EXEC SQL define FNAME_LEN      15;
EXEC SQL define LNAME_LEN     15;

main()
{
EXEC SQL BEGIN DECLARE SECTION;
    char fname[ FNAME_LEN + 1 ];
    char lname[ LNAME_LEN + 1 ];
EXEC SQL END DECLARE SECTION;

    printf( "DEM01 Sample ESQL Program running.\n\n");

EXEC SQL WHENEVER ERROR STOP;

EXEC SQL connect to 'stores7';

EXEC SQL declare democursor cursor for
    select fname, lname
        into :fname, :lname
        from customer
        where lname < "C";

EXEC SQL open democursor;
for (;;)
{
EXEC SQL fetch democursor;
if (strncmp(SQLSTATE, "00", 2) != 0)
    break;
printf("%s %s\n",fname, lname);
}
```

```

    }
    if (strncmp(SQLSTATE, "02", 2) != 0)
        printf("SQLSTATE after fetch is %s\n", SQLSTATE);

    EXEC SQL close democursor;
    EXEC SQL free democursor;
    EXEC SQL create routine from 'del_ord.sql';
    EXEC SQL disconnect current;
    printf("\nDEMO1 Sample Program over.\n\n");
    exit(0);
}

```

Related Information

Related statements: “DECLARE ” on page 2-286, “FETCH ” on page 2-367, “FLUSH ” on page 2-377, “FREE ” on page 2-379, “OPEN ” on page 2-459, “PUT ” on page 2-479, and “ SET AUTOFREE ” on page 2-582

For an introductory discussion of cursors, see the *IBM Informix Guide to SQL: Tutorial*.

For a more advanced discussion of cursors, see the *IBM Informix ESQL/C Programmer’s Manual*.

CLOSE DATABASE

Use the CLOSE DATABASE statement to close the implicit connection to the current database.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax

►►—CLOSE DATABASE—◄◄

Usage

When you issue a CLOSE DATABASE statement, you can issue only the following SQL statements immediately after it:

- CONNECT
- CREATE DATABASE
- DATABASE
- DROP DATABASE
- DISCONNECT

(The DISCONNECT statement is valid here only if an explicit connection existed before CLOSE DATABASE was executed.)

Issue the CLOSE DATABASE statement before you drop the current database.

If your current database supports transaction logging, and if you have started a transaction, you must issue a COMMIT WORK or ROLLBACK WORK statement before you can use the CLOSE DATABASE statement.

The following example shows how to use the CLOSE DATABASE statement to drop the current database:

```
DATABASE stores_demo
. . .
CLOSE DATABASE
DROP DATABASE stores_demo
```

In Informix ESQL/C, the CLOSE DATABASE statement cannot appear in a multistatement PREPARE operation.

If a previous CONNECT statement has established an explicit connection to a database, and that connection is still your current connection, you cannot use the CLOSE DATABASE statement to close that explicit connection. (You can use the DISCONNECT statement to close the explicit connection.)

If you use the CLOSE DATABASE statement within a UDR called by a WHENEVER statement, specify WHENEVER SQLERROR CONTINUE and WHENEVER SQLWARNING CONTINUE before the ROLLBACK WORK statement. This action prevents the program from looping if the ROLLBACK WORK statement encounters an error or a warning.

When you issue the CLOSE DATABASE statement, any declared cursors are no longer valid. You must re-declare any cursors that you want to use.

In an ANSI-compliant database, if no error is encountered while you exit from DB–Access in non-interactive mode without issuing the CLOSE DATABASE, COMMIT WORK, or DISCONNECT statement, the database server automatically commits any open transaction.

Related Information

Related statements: “CONNECT ” on page 2-83, “CREATE DATABASE ” on page 2-97, “DATABASE ” on page 2-281, “DISCONNECT ” on page 2-320, and “DROP DATABASE ” on page 2-325

COMMIT WORK

Use the COMMIT WORK statement to commit all modifications made to the database from the beginning of a transaction.

Syntax

```
▶▶—COMMIT—┐—————▶▶
             └─WORK─┘
```

Usage

The COMMIT WORK statement informs the database server that you reached the end of a series of statements that must succeed as a single unit. The database server takes the required steps to make sure that all modifications that the transaction makes are completed correctly and saved to disk.

Use COMMIT WORK only at the end of a multistatement operation in a database with transaction logging, when you are sure that you want to keep all changes made to the database from the beginning of a transaction.

The COMMIT WORK statement releases all row and table locks.

The WORK keyword is optional in a COMMIT WORK statement. The following two statements are equivalent:

```
COMMIT;  
COMMIT WORK;
```

The following example shows a transaction bounded by BEGIN WORK and COMMIT WORK statements.

```
BEGIN WORK;  
  DELETE FROM call_type WHERE call_code = '0';  
  INSERT INTO call_type VALUES ('S', 'order status');  
COMMIT WORK;
```

In this example, the user first deletes the row from the **call_type** table where the value of the **call_code** column is 0. The user then inserts a new row in the **call_type** table where the value of the **call_code** column is S. The database server guarantees that both operations succeed or else neither succeeds.

In Informix ESQL/C, the COMMIT WORK statement closes all open cursors except those that were declared using the WITH HOLD option.

Issuing COMMIT WORK in a Database That Is Not ANSI Compliant

In a database that is not ANSI compliant, but that supports transaction logging, if you initiate a transaction with a BEGIN WORK statement, you must issue a COMMIT WORK statement at the end of the transaction. If you fail to issue a COMMIT WORK statement in this case, the database server rolls back any modifications that the transaction made to the database.

If you do not issue a BEGIN WORK statement, however, each statement executes within its own transaction. These single-statement transactions do not require either a BEGIN WORK statement or a COMMIT WORK statement.

Explicit DB-Access Transactions

When you use DB-Access in interactive mode with a database that is not ANSI-compliant but that supports transaction logging, if you select the **Commit** menu but do not issue the COMMIT WORK statement after a transaction has been started by the BEGIN WORK statement, DB-Access automatically commits the data, but issues the following warning:

```
Warning: Data commit is a result of unhandled exception in TXN PROC/FUNC
```

The purpose of this warning is to remind you to issue COMMIT WORK explicitly to end a transaction that BEGIN WORK initiated.

In non-interactive mode, however, DB-Access rolls back the current transaction if you end a session without issuing the COMMIT WORK statement.

Issuing COMMIT WORK in an ANSI-Compliant Database

In an ANSI-compliant database, you do not need BEGIN WORK to mark the beginning of a transaction. You only need to mark the end of each transaction, because a transaction is always in effect. A new transaction starts automatically after each COMMIT WORK or ROLLBACK WORK statement.

You must, however, issue an explicit COMMIT WORK statement to mark the end of each transaction. If you fail to do so, the database server rolls back any modifications that the transaction made to the database.

In an ANSI-compliant database, however, if no error is encountered while you exit from DB–Access in non-interactive mode without issuing the CLOSE DATABASE, COMMIT WORK, or DISCONNECT statement, the database server automatically commits any open transaction.

Related Information

Related statements: “BEGIN WORK ” on page 2-74, “SAVE EXTERNAL DIRECTIVES ” on page 2-522, and “DECLARE ” on page 2-286

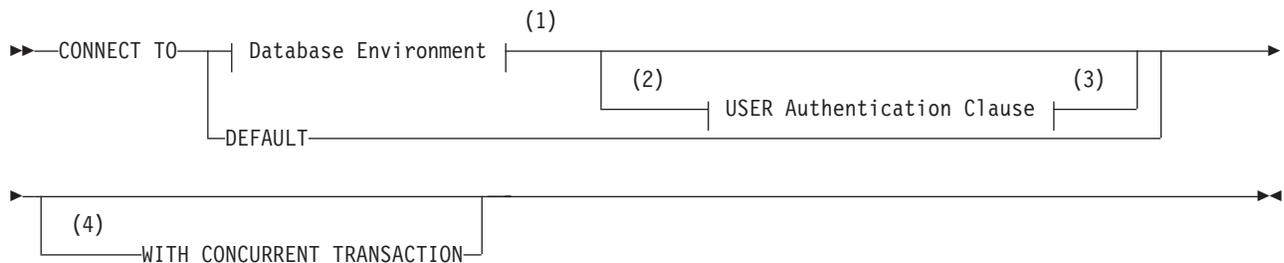
For a discussion of concepts related to transactions, see the *IBM Informix Guide to SQL: Tutorial*.

CONNECT

Use the CONNECT statement to connect to a database environment.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 See “Database Environment” on page 2-84
- 2 ESQL/C and DB-Access only
- 3 See “USER Validation Clause” on page 2-87
- 4 ESQL/C only

Usage

The CONNECT statement connects an application to a *database environment*, which can be a database, a database server, or a database and a database server. If the application successfully connects to the specified database environment, the connection becomes the *current connection* for the application. SQL statements fail if the application has no current connection to a database server. If you specify a database name, the database server opens that database. You cannot include CONNECT within a PREPARE statement.

An application can connect to several database environments at the same time, and it can establish multiple connections to the same database environment, provided each connection has a unique connection name.

On UNIX, the only restriction on establishing multiple connections to the same database environment is that an application can establish only one connection to each local server that uses the shared-memory connection mechanism. To find out whether a local server uses the shared-memory connection mechanism or the local-loopback connection mechanism, examine the `$INFORMIXDIR/etc/sqlhosts` file. For more information on the `sqlhosts` file, refer to your *IBM Informix Administrator's Guide*.

On Windows, the local connection mechanism is named pipes. Multiple connections to the local server from one client can exist.

Only one connection is current at any time; other connections are dormant. The application cannot interact with a database through a dormant connection. When an application establishes a new connection, that connection becomes current, and the previous current connection becomes dormant. You can make a dormant connection current with the `SET CONNECTION` statement. See also “`SET CONNECTION`” on page 2-587.

Privileges for Executing the `CONNECT` Statement

The current user, or `PUBLIC`, must hold the `Connect` privilege on the database that the `CONNECT` statement specifies. The user who executes the `CONNECT` statement cannot have the same authorization identifier as an existing role in that database.

For information on how to use the `USER Authentication` clause to specify an alternate user name when the `CONNECT` statement connects to a database server on a remote host, see “`USER Validation Clause`” on page 2-87.

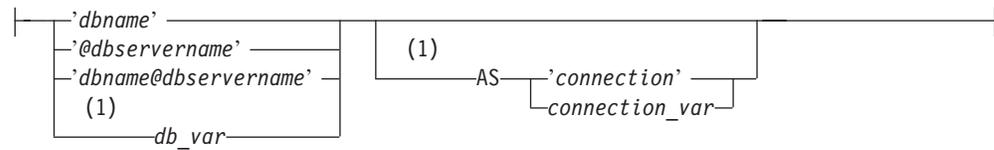
Connection Context

Each connection encompasses a set of information that is called the *connection context*. The connection context includes the name of the current user, the information that the database environment associates with this name, and information on the state of the connection (such as whether an active transaction is associated with the connection). The connection context is saved when an application becomes dormant, and this context is restored when the application becomes current again. (For more information, see “`Making a Dormant Connection the Current Connection`” on page 2-587.)

Database Environment

The `CONNECT` statement, like the `SET CONNECTION` statement, can use the Database Environment syntax segment to specify the database or the database server to which the application is attempting to establish a connection. Unlike the `SET CONNECTION` statement, the `CONNECT` statement can also declare a name for this connection to the specified database environment.

Database Environment:



Notes:

- 1 ESQ/C only

Element	Description	Restrictions	Syntax
<i>connection</i>	Optional case-sensitive name that you declare here for a connection	Must be unique among connection names	"Identifier " on page 5-21
<i>connection_var</i>	Host variable that stores the name of <i>connection</i>	Must be a fixed-length character data type	Language specific
<i>db_var</i>	Host variable that contains a valid database environment (in one of the formats in the syntax diagram)	Must be a fixed-length character data type, whose contents are in a format from the syntax diagram	Language specific
<i>dbname</i>	Database to which to connect	Must already exist	"Identifier " on page 5-21
<i>dbservername</i>	Name of the database server to which a connection is made	Must already exist; blank space is not valid between @ symbol and <i>dbservername</i> . See also "Restrictions on dbservername."	"Identifier " on page 5-21

If the **DELIMITED** environment variable is set, any quotation (') marks in the database environment must be single. If **DELIMITED** is not set, then either single (') or double (") quotation marks are valid here.

Restrictions on dbservername

If you specify *dbservername*, it must satisfy the following restrictions.

- If the database server that you specify is not online, you receive an error.
- On UNIX, the database server that you specify in *dbservername* must match the name of a database server in the **sqlhosts** file.
- On Windows, *dbservername* must match the name of a database server in the **sqlhosts** subkey in the registry. It is recommended that you use the **setnet32** utility to update the registry.

If the name of a database server is a delimited identifier, or if it includes uppercase letters, that database server cannot participate in cross-server distributed DML operations. To avoid this restriction, use only undelimited names that include no uppercase letters when you declare the name or the alias of a database server that will participate in cross-server queries.

Specifying the Database Environment

You can specify a database server and a database, or a database server only, or a database only. How a database is located and opened depends on whether you specify a database server name in the database environment expression.

Only Database Server Specified: The *@dbservername* option establishes a connection to the database server only; it does not open a database. When you use

this option, you must subsequently use the `DATABASE` or `CREATE DATABASE` statement (or a `PREPARE` statement for one of these statements and an `EXECUTE` statement) to open a database.

Database Server and Database Specified: If you specify both a database server and a database, your application connects to the database server, which locates and opens the database.

Only Database Specified: The `dbname` option establishes a connection to the default database server or to another database server in the `DBPATH` environment variable. It also locates and opens the specified database. (The same is true of the `db_var` option if this specifies only a database name.)

If you specify only `dbname`, its database server is read from the `DBPATH` environment variable. The database server in the `INFORMIXSERVER` environment variable is always added before the `DBPATH` value.

On UNIX, set the `INFORMIXSERVER` and `DBPATH` environment variables as the following example (for the C shell) shows:

```
setenv INFORMIXSERVER srvA
setenv DBPATH //srvB://srvC
```

On Windows, choose **Start > Programs > Informix > setnet32** from the Task Bar and set the `INFORMIXSERVER` and `DBPATH` environment variables:

```
set INFORMIXSERVER = srvA
set DBPATH = //srvA://srvB://srvC
```

The next example shows the resulting `DBPATH` that your application uses:
`//srvA://srvB://srvC`

The application first establishes a connection to the database server that `INFORMIXSERVER` specifies. The database server uses parameters in the configuration file to locate the database. If the database does not reside on the default database server, or if the default database server is not online, the application connects to the next database server in `DBPATH`. In the previous example, that database server would be `srvB`.

Declaring a Connection Name

In ESQL/C applications, you can declare an identifier for the connection to the database environment by including the `AS` keyword, followed by a quoted string or by a host variable that stores the identifier. The host variable must be a fixed-length character data type.

Connection Identifiers

The optional *connection* name is a unique identifier that an ESQL/C application can use to refer to a connection in subsequent `SET CONNECTION` and `DISCONNECT` statements. If the application does not provide a connection name (or a connection host variable), it can refer to the connection using the database environment. If the application makes more than one connection to the same database environment, however, each connection must have a unique name.

Only the `CONNECT` statement can use the `AS` keyword to declare a connection name. The `CONNECT` statement cannot, however, reference a previously declared connection name to specify a connection to a database environment.

USER Validation Clause

The USER Authentication clause specifies information that is used to determine whether the application can access the target computer on a remote host.

USER Authentication Clause:

```
|—USER—'user_id'——|—USING—validation_var——|
      |user_id_var—|
```

Element	Description	Restrictions	Syntax
<i>user_id</i>	Valid login name	See “Restrictions on the User Identifier Parameter.”	“Quoted String” on page 4-189
<i>user_id_var</i>	Host variable that contains <i>user_id</i>	Must be a fixed-length character data type; same restrictions as <i>user_id</i>	Language specific
<i>validation_var</i>	Host variable that contains a valid password for login name in <i>user_id</i> or <i>user_id_var</i>	Must be a fixed-length character type. See “Restrictions on the Validation Variable Parameter.”	Language specific

The USER Authentication clause is required when the CONNECT statement connects to the database server on a remote host. Subsequent to the CONNECT statement, all database operations on the remote host use the specified user name.

In DB-Access, the USING clause is valid within files executed from DB-Access. In interactive mode, DB-Access prompts you for a password, so the USING keyword and *validation_var* are not used.

Restrictions on the Validation Variable Parameter

On UNIX, the password stored in *validation_var* must be a valid password and must exist in the `/etc/passwd` file. If the application connects to a remote database server, the password must exist in this file on both the local and remote database servers.

On Windows, the password stored in *validation_var* must be valid and must be the password entered in **User Manager**. If the application connects to a remote database server, the password must exist in the domain of both the client and the server.

Restrictions on the User Identifier Parameter

The connection is rejected if any of the following conditions occur:

- The specified user lacks the privileges to access the database specified in the database environment.
- The specified user lacks the permissions to connect to the remote host.
- You supply a USER Authentication clause but omit the USING *validation_var* specification.

In compliance with the X/Open standard for the CONNECT statement, the Informix ESQL/C preprocessor supports a CONNECT statement that has a USER Authentication clause without the USING *validation_var* specification. If the *validation_var* is not present, however, the database server rejects the connection at runtime.

On UNIX, the *user_id* that you specify must be a valid login name and must exist in the `/etc/passwd` file. If the application connects to a remote server, the login name must exist in this file on both the local and remote database servers.

On Windows, the *user_id* that you specify must be a valid login name and must exist in **User Manager**. If the application connects to a remote server, the login name must exist in the domain of both the client and the server.

Use of the Default User ID

If you do not supply the USER Authentication clause, the default user ID is used to attempt the connection. The default user ID is the login name of the user running the application. In this case, you obtain network permissions with the standard authorization procedures. For example, on UNIX, the default user ID must match a user ID in the `/etc/hosts.equiv` file. On Windows, you must be a member of the domain, or if the database server is installed locally, you must be a valid user on the computer where it is installed.

The DEFAULT Connection Specification

Instead of specifying an explicit database environment, you can use the DEFAULT keyword to request a *default connection* to a default database server. The default database server can be local or remote. To designate the default database server, set its name in the **INFORMIXSERVER** environment variable. This option of CONNECT does not open a database.

If the CONNECT TO DEFAULT statement succeeds, you must use the DATABASE statement or the CREATE DATABASE statement to open or create a database in the default database environment.

The Implicit Connection with DATABASE Statements

If you do not execute a CONNECT statement in your application, the first SQL statement must be one of the following *database statements* (or a single statement PREPARE for one of the following statements):

- DATABASE
- CREATE DATABASE
- DROP DATABASE

If one of these database statements is the first SQL statement in an application, the statement establishes a connection to a database server, which is known as an *implicit* connection. If the database statement specifies only a database name, the database server name is obtained from the **DBPATH** environment variable. This situation is described in “Specifying the Database Environment” on page 2-85.

An application that makes an implicit connection can establish other connections explicitly (using the CONNECT statement) but cannot establish another implicit connection unless the original implicit connection is closed. An application can terminate an implicit connection using the DISCONNECT statement. After you create an explicit connection, you cannot use any database statement to create implicit connections until after you close the explicit connection.

After *any* implicit connection is made, that connection is considered to be the default connection, regardless of whether the database server is the default that the **INFORMIXSERVER** environment variable specifies. This feature allows the application to refer to the implicit connection if additional explicit connections are made, because the implicit connection has no identifier.

For example, if you establish an implicit connection followed by an explicit connection, you can make the implicit connection current by issuing the SET CONNECTION DEFAULT statement. This means, however, that once you establish an implicit connection, you cannot use the CONNECT DEFAULT statement, because the implicit connection is now the default connection.

The database statements can always be used to open a database or create a new database on the current database server.

WITH CONCURRENT TRANSACTION Option

The WITH CONCURRENT TRANSACTION clause enables you to switch to a different connection while a transaction is active in the current connection. If the current connection was *not* established using the WITH CONCURRENT TRANSACTION clause, you cannot switch to a different connection if a transaction is active; the CONNECT or SET CONNECTION statement fails, returning an error, and the transaction in the current connection continues to be active.

In this case, the application must commit or roll back the active transaction in the current connection before it switches to a different connection.

The WITH CONCURRENT TRANSACTION clause supports the concept of multiple concurrent transactions, where each connection can have its own transaction and the COMMIT WORK and ROLLBACK WORK statements affect only the current connection. The WITH CONCURRENT TRANSACTION clause does not support global transactions in which a single transaction spans databases over multiple connections. The COMMIT WORK and ROLLBACK WORK statements do not act on databases across multiple connections.

The following example illustrates how to use the WITH CONCURRENT TRANSACTION clause:

```
main()
{
EXEC SQL connect to 'a@srv1' as 'A';
EXEC SQL connect to 'b@srv2' as 'B' with concurrent transaction;
EXEC SQL connect to 'c@srv3' as 'C' with concurrent transaction;

/*
   Execute SQL statements in connection 'C' , starting a transaction
*/
EXEC SQL set connection 'B'; -- switch to connection 'B'

/*
   Execute SQL statements starting a transaction in 'B'.
   Now there are two active transactions, one each in 'B' and 'C'.
*/

EXEC SQL set connection 'A'; -- switch to connection 'A'

/*
   Execute SQL statements starting a transaction in 'A'.
   Now there are three active transactions, one each in 'A', 'B' and 'C'.
*/

EXEC SQL set connection 'C'; -- ERROR, transaction active in 'A'

/*
   SET CONNECTION 'C' fails (current connection is still 'A')
```

```

    The transaction in 'A' must be committed or rolled back because
    connection 'A' was started without the CONCURRENT TRANSACTION
    clause.
*/
EXEC SQL commit work;  -- commit tx in current connection ('A')

/*
    Now, there are two active transactions, in 'B' and in 'C',
    which must be committed or rolled back separately
*/

EXEC SQL set connection 'B'; -- switch to connection 'B'
EXEC SQL commit work;        -- commit tx in current connection ('B')

EXEC SQL set connection 'C'; -- go back to connection 'C'
EXEC SQL commit work;        -- commit tx in current connection ('C')

EXEC SQL disconnect all;
}

```

Warning: When an application uses the WITH CONCURRENT TRANSACTION clause to establish multiple connections to the same database environment, a deadlock condition can occur.

Related Information

Related Statements: “DISCONNECT ” on page 2-320, “SET CONNECTION” on page 2-587, “DATABASE ” on page 2-281, and “CREATE DATABASE ” on page 2-97

For more information about `sqlhosts`, refer to your *IBM Informix Administrator's Guide*.

CREATE ACCESS_METHOD

Use the CREATE ACCESS_METHOD statement to register a new primary or secondary access method in the `sysams` system catalog table.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```

▶▶ CREATE SECONDARY  
PRIMARY ACCESS_METHOD access_method

```



```

▶▶ ( (1)  
Purpose Options )

```

Notes:

- 1 See “Purpose Options” on page 5-50

Element	Description	Restrictions	Syntax
<i>access method</i>	Name declared here for the new access method	Must be unique among access-method names in the sysams system catalog table	"Identifier " on page 5-21

Usage

The CREATE ACCESS_METHOD statement adds a user-defined access method to a database. To create an access method, you specify *purpose functions* (or *purpose methods*), *purpose flags*, or *purpose values* as attributes of the access method, and you associate keywords (based on column names in the **sysams** system catalog table) with UDRs. You must have the DBA or Resource privilege to create an access method.

For information on setting purpose options, including a list of all the purpose function keywords, refer to "Purpose Options" on page 5-50.

The PRIMARY keyword specifies a user-defined primary-access method for a virtual table. The SECONDARY keyword specifies creating a user-defined secondary-access method for a virtual index. The SECONDARY keyword (and creating virtual indexes) is not supported in the Java Virtual-Table Interface.

The following statement creates a secondary-access method named **T_tree**:

```
CREATE SECONDARY ACCESS_METHOD T_tree
(
  am_getnext = ttree_getnext,

  . . .
  am_unique,
  am_cluster,
  am_sptype = 'S'
);
```

In the preceding example, the **am_getnext** keyword in the Purpose Options list is associated with the **ttree_getnext()** UDR as the name of a method to scan for the next item that satisfies a query. This example indicates that the **T_tree** secondary access method supports unique keys and clustering, and resides in an sbspace.

Any UDR that the CREATE ACCESS_METHOD statement associates with the keyword for a purpose function task, such as the association of **ttree_getnext()** with **am_getnext** in the preceding example, must already have been registered in the database by the CREATE FUNCTION statement (or by a functionally equivalent statement, such as CREATE PROCEDURE FROM).

The following statement creates a primary-access method named **am_tabprops** that resides in an extspace.

```
CREATE PRIMARY ACCESS_METHOD am_tabprops
(
  am_open = FS_open,
  am_close = FS_close,
  am_beginscan = FS_beginScan,
  am_create = FS_create,
  am_scancost = FS_scanCost,
  am_endscan = FS_endScan,
  am_getnext = FS_getNext,
  am_getbyid = FS_getById,
  am_drop = FS_drop,
```

```

am_truncate = FS_truncate,
am_rowids,
am_sptype = 'x'
);

```

Related Information

Related statements: “ALTER ACCESS_METHOD ” on page 2-5 and “DROP ACCESS_METHOD” on page 2-322

For the schema of the **sysams** table, see the *IBM Informix Guide to SQL: Reference*.

For information about how to set purpose-option specifications, see “Purpose Options” on page 5-50.

For more information on primary-access methods, see the *IBM Informix Virtual-Table Interface Programmer’s Guide*.

For more information on secondary-access methods, see the *IBM Informix Virtual-Index Interface Programmer’s Guide* and the *IBM Informix User-Defined Routines and Data Types Developer’s Guide*.

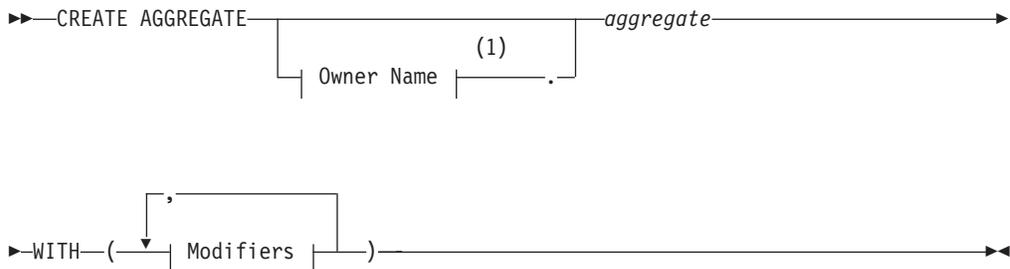
For a discussion of privileges, see the “GRANT ” on page 2-395 or “REVOKE” on page 2-494 statements or the *IBM Informix Database Design and Implementation Guide*.

CREATE AGGREGATE

Use the CREATE AGGREGATE statement to create a new aggregate function and register it in the **sysaggregates** system catalog table. User-defined aggregates extend the functionality of the database server by performing aggregate computations that the user implements.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Modifiers:



Notes:

- 1 See “ Owner Name ” on page 5-45

Element	Description	Restrictions	Syntax
<i>aggregate</i>	Name of the new aggregate	Must be unique among names of built-in aggregates and UDRs	“Identifier ” on page 5-21
<i>comb_func</i>	Function that merges one partial result into the other and returns the updated partial result	Must specify the combined function both for parallel queries and for sequential queries	“Identifier ” on page 5-21
<i>final_func</i>	Function that converts a partial result into the result type	If this is omitted, then the returned value is the final result of <i>iter_func</i>	“Identifier ” on page 5-21
<i>init_func</i>	Function that initializes the data structures required for the aggregate computation	Must be able to handle NULL arguments	“Identifier ” on page 5-21
<i>iter_func</i>	Function that merges a single value with a partial result and returns updated partial result	Must specify an iterator function. If <i>init_func</i> is omitted, <i>iter_func</i> must be able to handle NULL arguments	“Identifier ” on page 5-21

Usage

You can specify the INIT, ITER, COMBINE, FINAL, and HANDLESNULLS modifiers in any order.

Important: You must specify the ITER and COMBINE modifiers in a CREATE AGGREGATE statement. You do not need to specify the INIT, FINAL, and HANDLESNULLS modifiers in a CREATE AGGREGATE statement.

The ITER, COMBINE, FINAL, and INIT modifiers specify the support functions for a user-defined aggregate. These support functions do not need to exist at the time when you create the user-defined aggregate.

If you omit the HANDLESNULLS modifier, rows with NULL aggregate argument values do not contribute to the aggregate computation. If you include the HANDLESNULLS modifier, you must define all the support functions to handle NULL values as well.

Extending the Functionality of Aggregates

Dynamic Server provides two ways to extend the functionality of aggregates. Use the CREATE AGGREGATE statement only for the second of the two cases.

- Extensions of built-in aggregates

A built-in aggregate is an aggregate that the database server provides, such as COUNT, SUM, or AVG. These only support built-in data types. To extend a built-in aggregate so that it supports a user-defined data type (UDT), you must create user-defined routines that overload the binary operators for that aggregate. For further information on extending built-in aggregates, see the *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

- Creation of user-defined aggregates

A user-defined aggregate is an aggregate that you define to perform an aggregate computation that the database server does not provide. You can use user-defined aggregates with built-in data types, extended data types, or both. To create a user-defined aggregate, use the CREATE AGGREGATE statement. In this statement, you name the new aggregate and specify the support functions

that compute the aggregate result. These support functions perform initialization, sequential aggregation, combination of results, and type conversion.

Example of Creating a User-Defined Aggregate

The following example defines a user-defined aggregate named **average**:

```
CREATE AGGREGATE average
  WITH (
    INIT = average_init,
    ITER = average_iter,
    COMBINE = average_combine,
    FINAL = average_final
  );
```

Before you use the **average** aggregate in a query, you must also use CREATE FUNCTION statements to create the support functions specified in the CREATE AGGREGATE statement.

The following table gives an example of the task that each support function might perform for **average**.

Keyword	Support Function	Effect
INIT	average_init	Allocates and initializes an extended data type storing the current sum and the current row count
ITER	average_iter	For each row, adds the value of the expression to the current sum and increments the current row count by one
COMBINE	average_combine	Adds the current sum and the current row count of one partial result to the other and returns the updated result
FINAL	average_final	Returns the ratio of the current sum to the current row count and converts this ratio to the result type

Parallel Execution

The database server can break up an aggregate computation into several pieces and compute them in parallel. The database server uses the INIT and ITER support functions to compute each piece sequentially. Then the database server uses the COMBINE function to combine the partial results from all the pieces into a single result value. Whether an aggregate is parallel is an optimization decision that is transparent to the user.

Related Information

Related statements: “CREATE FUNCTION ” on page 2-113 and “ DROP AGGREGATE ” on page 2-323

For information about how to invoke a user-defined aggregate, see “User-Defined Aggregates” on page 4-167 in the Expression segment.

For a description of the **sysaggregates** system catalog table that stores data about user-defined aggregates, see the *IBM Informix Guide to SQL: Reference*.

For a discussion of user-defined aggregates, see *IBM Informix User-Defined Routines and Data Types Developer’s Guide*.

CREATE CAST

Use the CREATE CAST statement to register a cast that converts data from one data type to another.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```
► CREATE EXPLICIT  
IMPLICIT CAST WITH function ( source_type AS target_type ) ►
```

Element	Description	Restrictions	Syntax
<i>function</i>	UDR that you register to implement the cast	See “WITH Clause” on page 2-97.	“Identifier ” on page 5-21
<i>source_type</i>	Data type to be converted	Must exist in the database at the time the cast is registered. See also “Source and Target Data Types.”	“Data Type ” on page 4-21
<i>target_type</i>	Data type that results from the conversion	The same restrictions that apply for the <i>source_type</i> (as listed above) also apply for the <i>target_type</i>	“Data Type ” on page 4-21

Usage

A cast is a mechanism that the database server uses to convert one data type to another. The database server uses casts to perform the following tasks:

- To compare two values in the WHERE clause of a SELECT, UPDATE, or DELETE statement
- To pass values as arguments to user-defined routines
- To return values from user-defined routines

To create a cast, you must have the necessary privileges on both the *source* data type and the *target* data type. All users have access privileges to use the built-in data types. To create a cast to or from an OPAQUE, DISTINCT, or named ROW data type, however, requires the Usage privilege on that data type.

The CREATE CAST statement registers a cast in the **syscasts** system catalog table. For more information on **syscasts**, see the chapter on system catalog tables in the *IBM Informix Guide to SQL: Reference*.

Source and Target Data Types

The CREATE CAST statement defines a cast that converts a *source* type to a *target* type. Both the *source* and *target* data types must exist in the database when you execute the CREATE CAST statement to register the cast. The *source* and the *target* data types have the following restrictions:

- Either the *source* or the *target* type, but not both, can be a built-in data type.
- Neither the *source* nor the *target* type can be a DISTINCT type of the other.
- Neither the *source* nor the *target* types can be a COLLECTION data type.

Explicit and Implicit Casts

To process queries with multiple data types often requires casts that convert data from one data type to another. You can use the CREATE CAST statement to create the following kinds of casts:

- Use the CREATE EXPLICIT CAST statement to define an *explicit* cast.
- Use the CREATE IMPLICIT CAST statement to define an *implicit* cast.

Explicit Casts

An explicit cast is a cast that you must specifically invoke, with either the CAST AS keywords or with the cast operator (::). The database server does *not* automatically invoke an explicit cast to resolve data type conversions. The EXPLICIT keyword is optional; by default, the CREATE CAST statement creates an explicit cast.

The following CREATE CAST statement defines an explicit cast from the **rate_of_return** opaque data type to the **percent** distinct data type:

```
CREATE EXPLICIT CAST (rate_of_return AS percent
    WITH rate_to_prct);
```

The following SELECT statement explicitly invokes this explicit cast in its WHERE clause to compare the **bond_rate** column (of type **rate_of_return**) to the **initial_APR** column (of type **percent**):

```
SELECT bond_rate FROM bond
    WHERE bond_rate::percent > initial_APR;
```

Implicit Casts

The database server invokes built-in casts to convert from one built-in data type to another built-in type that is not directly substitutable. For example, the database server performs conversion of a character type such as CHAR to a numeric type such as INTEGER through a built-in cast.

An implicit cast is a cast that the database server can invoke automatically when it encounters data types that cannot be compared with built-in casts. This type of cast enables the database server to automatically handle conversions between other data types.

To define an implicit cast, specify the IMPLICIT keyword in the CREATE CAST statement. For example, the following CREATE CAST statement specifies that the database server should automatically use the **prcnt_to_char()** function to convert from the CHAR data type to a distinct data type, **percent**:

```
CREATE IMPLICIT CAST (CHAR AS percent WITH char_to_prct);
```

This cast only supports automatic conversion *from* the CHAR data type *to* **percent**. For the database server to convert *from* **percent** *to* CHAR, you also need to define another implicit cast, as follows:

```
CREATE IMPLICIT CAST (percent AS CHAR WITH prcnt_to_char);
```

The database server automatically invokes the **char_to_prct()** function to evaluate the WHERE clause of the following SELECT statement:

```
SELECT commission FROM sales_rep WHERE commission > "25"
```

Users can also invoke implicit casts explicitly. For more information on how to explicitly invoke a cast function, see “Explicit Casts.”

When a built-in cast does not exist for conversion between data types, you can create user-defined casts to make the necessary conversion.

WITH Clause

The WITH clause of the CREATE CAST statement specifies the name of the user-defined function to invoke to perform the cast. This function is called the cast function. You must specify a *function name* unless the *source data type* and the *target data type* have identical representations. Two data types have identical representations when the following conditions are met:

- Both data types have the same length and alignment.
- Both data types are passed by reference or both are passed by value.

The cast function must be registered in the same database as the cast at the time the cast is invoked, but need not exist when the cast is created. The CREATE CAST statement does not check privileges on the specified *function name*, or even verify that the cast function exists. Each time a user invokes the cast explicitly or implicitly, the database server verifies that the user has the Execute privilege on the cast function.

Related Information

Related statements: "CREATE FUNCTION " on page 2-113, "CREATE DISTINCT TYPE " on page 2-100, "CREATE OPAQUE TYPE " on page 2-145, "CREATE ROW TYPE " on page 2-168, and "DROP CAST " on page 2-324

For more information about data types, casting, and conversion, see the Data Types segment in this document and the *IBM Informix Guide to SQL: Reference*.

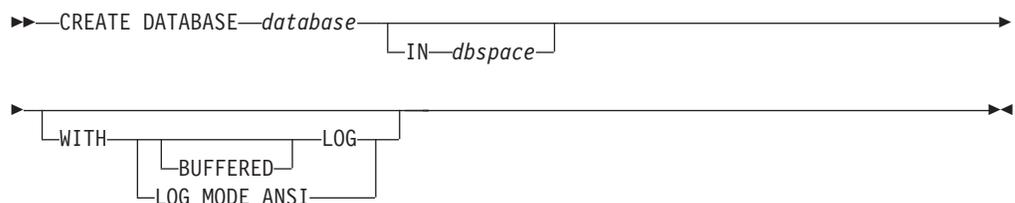
For examples that show how to create and use casts, see the *IBM Informix Database Design and Implementation Guide*.

CREATE DATABASE

Use the CREATE DATABASE statement to create a new database.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax



Element	Description	Restrictions	Syntax
<i>database</i>	Name that you declare here for the new database that you are creating	Must be unique among names of databases of the database server	"Database Name" on page 5-15
<i>dbspace</i>	The dbspace to store the data for this database; default is the root dbspace	Must already exist	"Identifier " on page 5-21

Usage

This statement is an extension to ANSI-standard syntax. (The ANSI/ISO standard for the SQL language does not specify any syntax for construction of a database, the process by which a database comes into existence and has its name declared.)

The *database* that CREATE DATABASE specifies becomes the current database.

If the DBCREATE_PERMISSION configuration parameter is not set, any user can create a database. If the configuration file includes one or more DBCREATE_PERMISSION specifications, however, only the specified users can create databases. Whether or not DBCREATE_PERMISSION is set, user **informix** can use the CREATE DATABASE statement. For additional information about how to set the DBCREATE_PERMISSION parameter to control which users can create new databases, see the *IBM Informix Administrator's Reference*.

The *database* name that you declare must be unique within the database server environment in which you are working. The database server creates the system catalog tables that describe the structure of the new database.

When you create a database, you alone can access it. It remains inaccessible to other users until you, as DBA, grant database privileges. For information on how to grant database privileges, see "GRANT " on page 2-395.

If a previous CONNECT statement has established an explicit connection to a database, and that connection is still your current connection, you cannot use the CREATE DATABASE statement (nor any SQL statement that creates an implicit connection) until after you use DISCONNECT to close the explicit connection.

In Informix ESQL/C, the CREATE DATABASE statement cannot appear in a multistatement PREPARE operation.

If you do not specify a *dbspace*, the database server creates the system catalog tables in the **root** dbspace by default. The following statement creates the **vehicles** database in the **root** dbspace:

```
CREATE DATABASE vehicles;
```

The following statement creates the **vehicles** database in the **research** dbspace:

```
CREATE DATABASE vehicles IN research;
```

In Extended Parallel Server, you can create a database in the dbspace of the primary coserver (coserver 1) only.

Logging Options

The logging options of the CREATE DATABASE statement determine the type of logging that is done for the database. In the event of a failure, the database server uses the log to re-create all committed transactions in your database. If you do not specify the WITH LOG option, the unbuffered log type is used.

The following statement uses the WITH LOG option to create a database with unbuffered logging:

```
CREATE DATABASE unbufDatabase WITH LOG;
```

If you do not specify the WITH LOG keywords, you cannot use transactions nor the statements that support transaction logging (BEGIN WORK, COMMIT WORK, ROLLBACK WORK, SET IMPLICIT TRANSACTION, SET LOG, and SET ISOLATION).

If you are using Extended Parallel Server, the CREATE DATABASE statement always creates a database with logging.

Specifying Buffered Logging

The following example creates a database that uses a buffered log:

```
CREATE DATABASE vehicles WITH BUFFERED LOG;
```

If you use a buffered log, you marginally enhance the performance of logging at the risk of not being able to re-create the last few transactions after a failure. (See the discussion of buffered logging in the *IBM Informix Database Design and Implementation Guide*.)

ANSI-Compliant Databases

When you use the LOG MODE ANSI option in the CREATE DATABASE statement, the database that you create is an ANSI-compliant database that conforms to the ANSI/ISO standard for the SQL language. The following example creates an ANSI-compliant database:

```
CREATE DATABASE employees WITH LOG MODE ANSI;
```

ANSI-compliant databases are different from databases that are not ANSI compliant in several ways, including the following features:

- All SQL statements are automatically contained in transactions.
- All databases use unbuffered logging.
- Owner naming is enforced.

You must qualify with the owner name any table, view, synonym, index, or constraint that you do not own. Unless you enclose the owner name between quotation marks, alphabetic characters in owner names default to uppercase. (To prevent this upshifting of lowercase letters in undelimited owner names, you can set the ANSIOWNER environment variable to 1.)

In addition, the routine signature of a UDR includes the name of the owner; in databases that are not ANSI compliant, this is true only for the **sysdbopen()** and **sysdbclose()** procedures.

- For sessions, the default isolation level is REPEATABLE READ.
- Default privileges on objects differ from those in databases that are not ANSI compliant. When you create a table or a synonym, other users do not receive access privileges (as members of the PUBLIC group) on the object by default.
- All DECIMAL data types are fixed-point values. If you declare a column as DECIMAL(*p*), the default scale is zero, meaning that only integer values can be stored. (In a database that is not ANSI compliant, DECIMAL(*p*) is a floating-point data type of a scale large enough to store the exponential notation for a value.)

Other slight differences exist between databases that are ANSI compliant and those that are not. These differences are noted with the related SQL statement in this document. For a detailed discussion of the differences between ANSI compliant databases and databases that are not ANSI-compliant, see the *IBM Informix Database Design and Implementation Guide*.

Creating an ANSI-compliant database does not mean that you automatically receive warnings for Informix extensions to the ANSI/ISO standard for SQL syntax when you run the database. You must also use the **-ansi** flag or the **DBANSIWARN** environment variable to receive such warnings.

For additional information about **-ansi** and **DBANSIWARN**, see the *IBM Informix Guide to SQL: Reference*.

Related Information

Related statements: “CLOSE DATABASE ” on page 2-80, “CONNECT ” on page 2-83, “DATABASE ” on page 2-281, and “DROP DATABASE ” on page 2-325

CREATE DISTINCT TYPE

Use the CREATE DISTINCT TYPE statement to create a new distinct data type.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

►►—CREATE DISTINCT TYPE—*distinct_type*—AS—*source_type*—◀◀

Element	Description	Restrictions	Syntax
<i>distinct_type</i>	Name that you declare here for the new distinct data type	In an ANSI-compliant database, the combination of the owner and data type must be unique within the database. In a database that is not ANSI compliant, the name must be unique among names of data types in the database.	“Data Type ” on page 4-21
<i>source_type</i>	Name of existing type on which the new type is based	Must be either a built-in data type or one created with the CREATE DISTINCT TYPE, CREATE OPAQUE TYPE, or CREATE ROW TYPE statement	“Data Type ” on page 4-21

Usage

A *distinct type* is a data type based on a built-in data type or on an existing opaque data type, a named ROW data type, or another distinct data type. Distinct data types are strongly typed. Although the distinct type has the same physical representation of data as its source type, values of the two types cannot be compared without an explicit cast from one type to the other.

To create a distinct data type, you must have the Resource privilege on the database. Any user with the Resource privilege can create a distinct type from one of the built-in data types, which user **informix** owns.

Important: You cannot create a distinct type on the SERIAL, BIGSERIAL, or SERIAL8 data types.

To create a distinct type from an opaque type, from a named ROW type, or from another distinct type, you must be the owner of the data type or have the Usage privilege on the data type.

By default, after a distinct type is defined, only the owner of the distinct type and the DBA can use it. The owner of the distinct type, however, can grant to other users the Usage privilege on the distinct type.

A distinct type has the same storage structure as its source type. The following statement creates the distinct type **birthday**, based on the built-in DATE data type:

```
CREATE DISTINCT TYPE birthday AS DATE;
```

Although Dynamic Server uses the same storage format for the distinct type as it does for its source type, a distinct type and its source type cannot be compared in an operation unless one type is cast explicitly to the other type.

Privileges on Distinct Types

To create a distinct type, you must have the Resource privilege on the database. When you create the distinct type, only you, the owner, have Usage privilege on this type. Use the GRANT or REVOKE statements to grant or revoke Usage privilege to other database users.

To find out what privileges exist on a particular type, check the **sysxdtypes** system catalog table for the owner name and the **sysxdtypeauth** system catalog table for additional data type privileges that might have been granted. For more information on system catalog tables, see the *IBM Informix Guide to SQL: Reference*.

The DB–Access utility can also display privileges on distinct types.

Support Functions and Casts

When you create a distinct type, Dynamic Server automatically defines two explicit casts:

- A cast from the distinct type to its source type
- A cast from the source type to the distinct type

Because the two data types have the same representation (the same length and alignment), no support functions are required to implement the casts.

You can create an implicit cast between a distinct type and its source type. To create an implicit cast, use the Table Options clause to specify the format of the external data. You must first drop the default explicit cast, however, between the distinct type and its source type.

All support functions and casts that are defined on the source type can be used on the distinct type. Casts and support functions that are defined on the distinct type, however, are not available to the source type. Use the Table Options clause to specify the format of the external data.

Manipulating Distinct Types

When you compare or manipulate data of a distinct type and its source type, you must explicitly cast one type to the other in the following situations:

- To insert or update a column of one type with values of the other type
- To use a relational operator to add, subtract, multiply, divide, compare, or otherwise manipulate two values, one of the source type and one of the distinct type

Usage

If the original *table* resides entirely on a single coserver, you can create duplicate copies of small tables across coservers for read-only use. For each attached index of the original table, a similarly defined index is built on each table duplicate, using the same dbspaces as the table.

Because query operators read the local copy of the table, duplicating small tables across coservers might improve the performance of some queries.

If a local copy of a duplicated table exists but is not available because the dbspace that stores it is offline (or for some similar reason), a query that requires access to the table fails. The database server does not attempt to access the original table.

The location of a duplicated table can be either a dbslice or a comma-separated list of dbspaces. You can combine dbslices and lists of dbspaces in a single CREATE DUPLICATE statement.

- If the original table is not fragmented, the dbspace list need provide only a single dbspace on each coserver.

For example, if the table **tab1** is not fragmented, enter the following statement to create a duplicate on the remaining three of the four coservers. If the original table is stored in the dbspace **db1** on coserver 1 and **db2** is on coserver 2, **db3** is on coserver 3, and **db4** is on coserver 4.

```
CREATE DUPLICATE OF TABLE tab1 IN (db2, db3, db4);
```

- If the original table is fragmented with one fragment in the first dbspace of several dbslices that contain dbspaces on all coservers, you can create duplicate copies of the table in the remaining dbspaces of the dbslice.

For example, you might create the **tab3** table in the first dbspace of three dbslices, each of which contains a dbspace on each coserver, as follows:

```
CREATE TABLE tab3 (...)  
  FRAGMENT BY HASH (...) IN dbs11.1, dbs12.1, dbs13.1;
```

To duplicate the **tab3** table across the remaining coservers, use the following statement:

```
CREATE DUPLICATE OF TABLE tab3 IN dbs11, dbs12, dbs13;
```

- You can mix dbslice names and dbspace lists in the same CREATE DUPLICATE statement. For example, instead of using dbspaces in a dbslice, for the previous example, you might enter the following statement in which **dbsp2a** is on coserver 2, **dbsp3a** is on coserver 3, and **dbsp4a** is on coserver 4:

```
CREATE DUPLICATE OF TABLE tab3 IN  
  dbs11, dbs12, (dbsp2a, dbsp3a, dbsp4a);
```

The first fragment of the original table is duplicated into **dbs11**, which contains a dbspace on each coserver, the second fragment into **dbs12**, which also contains a dbspace on each coserver, and the third fragment into the list of dbspaces.

Only one fragment of a duplicated table can reside in any single dbspace. You cannot list an existing dbspace of the duplicated table in the list of dbspaces into which it is duplicated, but it is not an error for an existing dbspace to be a member of a dbslice that specifies duplication dbspaces. Matching dbspaces in the dbslice are ignored.

The CREATE DUPLICATE statement requires the ALTER privilege.

Supported Operations

The following operations are permitted on duplicated tables:

- SELECT
- UPDATE STATISTICS
- LOCK and UNLOCK
- SET Residency
- DROP TABLE

You cannot duplicate a table in certain circumstances. The table must not:

- Have GK or detached indexes
- Use range fragmentation
- Be a temporary table
- Be a violations or diagnostic table
- Contain BYTE, TEXT, SERIAL, or SERIAL8 columns
- Have referential constraints

The CREATE DUPLICATE statement does not support incremental duplication. It also does not support multiple duplicates of the same table on a single coserver, nor duplication of tables that are fragmented across coservers.

If you need to take a dbspace offline and it contains a copy of a duplicated table, or if you need to update data in a duplicated table, first drop all duplicates of the table, as described in "DROP DUPLICATE " on page 2-326.

Related Information

"DROP DUPLICATE " on page 2-326

CREATE EXTERNAL TABLE (XPS)

Use the CREATE EXTERNAL TABLE statement to define an external source that is not part of your database to load and unload data for your database.

Only Extended Parallel Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```
►► CREATE EXTERNAL TABLE table | Column Definition | (1)
►► USING ( | Table Options | (2) | DATAFILES Clause | (3) | Table Options | (2) | )
```

Notes:

- 1 See "Column Definition" on page 2-105
- 2 See "Table Options" on page 2-109
- 3 See "DATAFILES Clause" on page 2-108

Element	Description	Restrictions	Syntax
<i>table</i>	Name declared here for a table to store external data	Must be unique among names of tables, views, and synonyms in the current database	"Identifier " on page 5-21

Usage

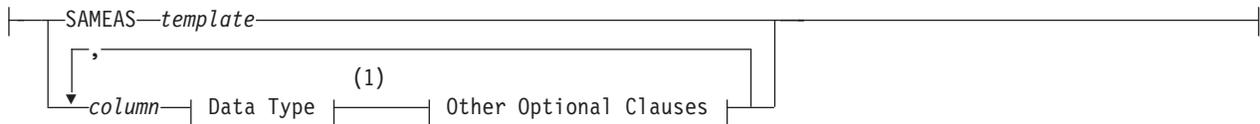
The first portion of the syntax diagram declares the name of the table and defines its columns and any column-level constraints.

The portion that follows the USING keyword identifies external files that the database server opens when you use the external table, and specifies additional options for characteristics of the external table.

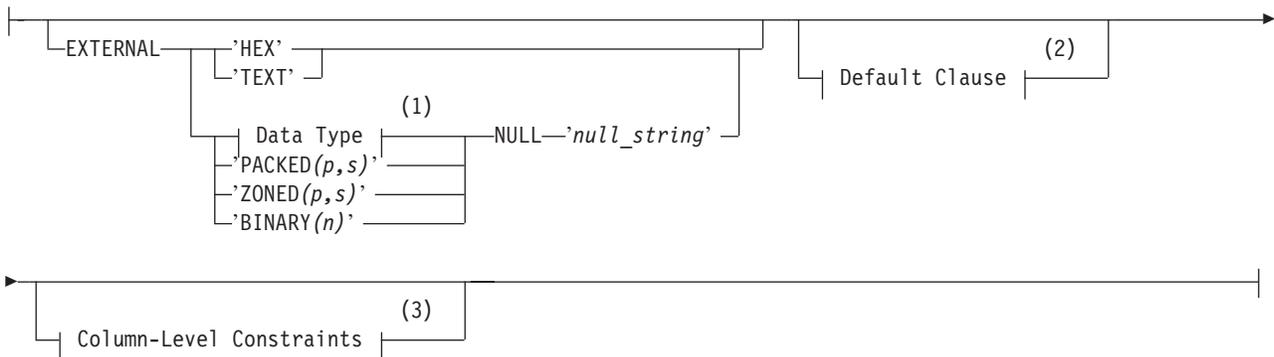
After executing the CREATE EXTERNAL TABLE statement, you can move data to and from the external source with an INSERT INTO ... SELECT statement. See the section " INTO EXTERNAL Clause (XPS)" on page 2-577 for more information about loading the results of a query into an external table.

Column Definition

Column Definition:



Other Optional Clauses:



Notes:

- 1 See "Data Type " on page 4-21
- 2 See "DEFAULT Clause" on page 2-195
- 3 See "Column-Level Constraints" on page 2-107

Element	Description	Restrictions	Syntax
<i>column</i>	One column name for each column of the external table	For each column, you must specify a built-in data type	"Identifier " on page 5-21
<i>n</i>	Number of 8-bit bytes to represent the integer	For FIXED format binary integers; <i>big-endian</i> byte order	<i>n</i> =2 for 16-bit integers <i>n</i> =4 for 32-bit integers

Element	Description	Restrictions	Syntax
<i>p</i>	Precision (number of digits)	For FIXED-format files only	"Literal Number " on page 4-185
<i>s</i>	Scale (digits in fractional part)	For FIXED-format files only	"Literal Number " on page 4-185
<i>null_string</i>	Value to represent NULL	See "Defining NULL Values."	"Quoted String" on page 4-189
<i>template</i>	Existing table with the same schema as the external table	Cannot be subset of columns nor differ in any column data type	"Database Object Name " on page 5-16

Using the SAMEAS Clause

The SAMEAS *template* clause uses all the column names and data types from the *template* table in the definition of new table. You cannot, however, use indexes in the external table definition, and you cannot use the SAMEAS clause for FIXED-format files.

Using the EXTERNAL Keyword

Use the EXTERNAL keyword to specify a data type for each column of your external table that has a data type different from the internal table. For example, you might have a VARCHAR column in the internal table that you want to map to a CHAR column in the external table.

You must specify an external type for every column that is in fixed format. You cannot specify an external type for delimited format columns except for BYTE and TEXT columns where your specification is optional. For more information, see "TEXT and HEX External Types" on page 2-107.

Integer Data Types: Besides valid Informix integer data types, you can specify packed decimal, zoned decimal, and IBM-format binary representation of integers. For packed or zoned decimal, specify *precision* (total number of digits in the number) and *scale* (number of digits that are to the right of the decimal point). Packed decimal representation can store two digits, or a digit and a sign, in each byte. Zoned decimal requires (*p* + 1) bytes to store *p* digits and the sign.

Big-Endian Format: The database server also supports two IBM-format binary representations of integers: BINARY(2) for 16-bit integer storage and BINARY(4) for 32-bit integer storage. The most significant byte of each number has the lowest address; that is, binary-format integers are stored big-end first (big-endian format) in the manner of IBM and Motorola processors. Intel® processors and some others store binary-format integers little-end first, a storage method that the database server does not support for external data.

Defining NULL Values: The packed decimal, zoned decimal, and binary data types do not have a natural NULL value, so you must define a value to be interpreted as a NULL when loading or unloading data from an external source. You can define the *null_string* as a number outside the set of numbers stored in the data file (for example, -9999.99). You can also define a bit pattern in the field as a hexadecimal pattern, such as 0xffff, that is to be interpreted as a NULL.

The database server uses the NULL representation for a FIXED-format external table to both interpret values as the data is loaded into the database and to format NULL values into the appropriate data type when data is unloaded to an external table.

The following examples are of column definitions with NULL values for a FIXED-format external table:

```
i smallint external "binary (2)" null "-32767"  
li integer external "binary (4)" null "-99999"  
d decimal (5,2) external "packed (5,2)" null "0xffffffff"  
z decimal (4,2) external "zoned (4,2)" null "0x0f0f0f0f"  
zl decimal (3,2) external "zoned (3,2)" null "-1.00"
```

If the packed decimal or zoned decimal is stored with all bits cleared to represent a NULL value, the *null_string* can be defined as 0x0. The following guidelines apply to the value assigned to a *null_string*:

- The NULL representation must fit into the length of the external field.
- If a bit pattern is defined, the *null_string* is not case sensitive.
- If a bit pattern is defined, the *null_string* must begin with 0x.
- For numeric fields, the left-most fields are assigned zeros by the database server if the bit pattern does not fill the entire field.
- If the NULL representation is not a bit pattern, the NULL value must be a valid number for that field.

Warning: If a row that contains a NULL value is unloaded into an external table and the column that receives the NULL value has no NULL value defined, the database server inserts a zero into the column.

TEXT and HEX External Types: An Informix BYTE or TEXT column can be encoded in either the TEXT or HEX external type. You can use only delimited BYTE and TEXT formats with these external types. Fixed formats are not allowed. In addition, you cannot use these external types with any other type of delimited-format columns (such as character columns).

You do not need to specify these external types. If you do not define an external column specifically, Informix TEXT columns default to TEXT and Informix BYTE columns default to HEX.

The database server interprets two adjacent field delimiters as a NULL value.

User-defined delimiters are limited to one byte each. During unloading, delimiters and backslash (\) symbols are escaped. During loading, any character that follows a backslash is interpreted as a literal. In TEXT format, nonprintable characters are directly embedded in the data file. For delimiter rules in a multibyte locale, see the *IBM Informix GLS User's Guide*.

For more information on BYTE and TEXT data, see your *IBM Informix Administrator's Guide*.

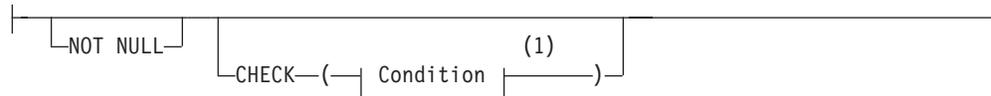
Manipulating Data in Fixed Format Files

For files in FIXED format, you must declare the column name and the EXTERNAL item for each column to set the name and number of characters. For FIXED-format files, the only data type allowed is CHAR. You can use the keyword NULL to specify what string to interpret as a NULL value.

Column-Level Constraints

Use column-level constraints to limit the type of data that is allowed in a column. Constraints at the column level are limited to a single column.

Column-Level Constraints:



Notes:

1 See "Condition " on page 4-5

Using the NOT NULL Constraint:

If you do not indicate a default value for a column, the default is NULL *unless* you place a not-NULL constraint on the column. In that case, no default value exists for the column. If you place a NOT NULL constraint on a column (and no default value is specified), the data in the external table must have a value set for the column when loading through the external table.

When no reject file exists and no value is encountered, the database server returns an error and the loading stops. When a reject file exists and no value is encountered, the error is reported in the reject file and the load continues.

Using the CHECK Constraint:

Check constraints allow you to designate conditions that must be met *before* data can be assigned to a column during an INSERT or UPDATE statement. When a reject file does not exist and a row evaluates to *false* for any check constraint defined on a table during an insert or update, the database server returns an error. When there is a reject file and a row evaluates to *false* for a check constraint defined on the table, the error is reported in the reject file and the statement continues to execute.

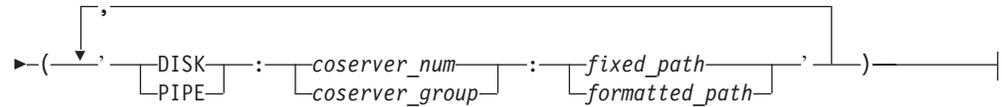
Check constraints are defined with *search conditions*. The search condition cannot contain subqueries, aggregates, host variables, or SPL routines. In addition, it cannot include the built-in functions CURRENT, USER, SITENAME, DBSERVERNAME, or TODAY. When you define a check constraint at the column level, the only column that the check constraint can check against is the column itself. In other words, the check constraint cannot depend upon values in other columns of the table.

DATAFILES Clause

The DATAFILES clause specifies external files that are opened when you use external tables.

DATAFILES Clause:





Element	Description	Restrictions	Syntax
<i>coserver_group</i>	Coserver group that contains the external data	Must exist	“Identifier ” on page 5-21
<i>coserver_num</i>	Numeric ID of coserver containing the external data	Must exist	“Literal Number ” on page 4-185
<i>fixed_path</i>	Pathname for input or output files in the definition of the external table	Must exist	Must conform to operating-system rules.
<i>formatted_path</i>	Formatted pathname that uses pattern-matching characters	Must exist	Must conform to operating-system rules.

You can use cogroup names and coserver numbers when you describe the input or output files for the external table definition. You can identify the DATAFILES either by coserver number or by cogroup name. A coserver number contains only digits. A cogroup name is a valid identifier that begins with a letter but otherwise contains any combination of letters, digits, and underscore symbols.

If you use only some of the available coservers for reading or writing files, you can designate these coservers as a cogroup using **onutil** and then use the cogroup name, rather than explicitly specifying each coserver and file separately. Whenever you use all coservers to manage external files, you can use the predefined *coserver_group*.

For examples of the DATAFILES clause, see “Examples” on page 2-112.

Using Formatting Characters

You can use a formatted pathname to designate a filename. If you use a formatted pathname, you can take advantage of the substitution characters %c, %n, and %r (*first ... last*).

Formatting String Effect

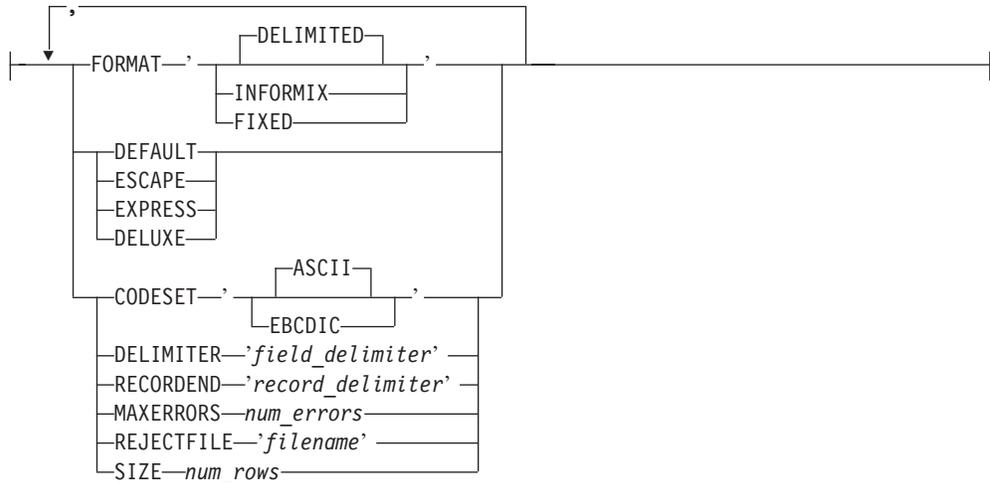
- %c Replaced with the number of the coserver that manages the file
- %n Replaced with the name of the node on which the coserver that manages the file resides
- %r(*first ... last*) Specifies multiple files on a single coserver

Important: The formatted pathname option does not support the %o formatting string.

Table Options

These options specify additional characteristics that define the table.

Table Options:



Element	Description	Restrictions	Syntax
<i>field_delimiter</i>	Character to separate fields. Default is pipe () character	For nonprinting characters, use octal	“Quoted String” on page 4-189
<i>filename</i>	Full pathname for conversion error messages from coservers	See “Reject Files” on page 2-111.	Must conform to operating-system rules.
<i>num_errors</i>	Errors per coserver before load operations are terminated	Value is ignored unless MAXERRORS is set	“Literal Number ” on page 4-185
<i>num_rows</i>	Approximate number of rows contained in the external table	Cannot be a negative number	“Literal Number ” on page 4-185
<i>quoted_string</i>	ASCII character specified here as the escape character	Only a single character is valid	“Quoted String” on page 4-189
<i>record_delimiter</i>	Character to separate records Default is Newline (\n)	For nonprinting characters, use octal	“Quoted String” on page 4-189

The *num_errors* specification is ignored during unload tasks. If the **MAXERRORS** environment variable is not set, the database server processes all data in load operations, regardless of the number of errors or *num_errors* value.

If the **RECORDEND** environment variable is not set, *record_delimiter* defaults to the Newline character (\n). To specify a nonprinting character as the record delimiter or field delimiter, you must encode it as the octal representation of the ASCII character. For example, \006 can represent CTRL-F.

Use the table options keywords as the following table describes. You can use each keyword whenever you plan to load or unload data unless only one of the two modes is specified.

Keyword
Description

CODESET

Specifies the type of code set of the data

DEFAULT (load only)

Specifies replacing missing values in delimited input files with column defaults (if they are defined) instead of NULLs, so input files can be sparsely populated. Files do not need an entry for every column in the file where a default is the value to be loaded.

DELIMITER

Specifies the character that separates fields in a delimited text file

DELUXE (load only)

Sets a flag causing the database server to load data in deluxe mode

Deluxe mode is required for loading into STANDARD tables.

ESCAPE

Defines a character to mark ASCII special characters in ASCII-text-based data files

EXPRESS

Sets a flag that causes the database server to attempt to load data in express mode. If you request express mode but indexes or unique constraints exist on the table or the table contains BYTE or TEXT data, or the target table is not RAW or OPERATIONAL, the load stops with an error message that reports the problem.

FORMAT

Specifies the format of the data in the data files

MAXERRORS

Sets the number of errors that are allowed per coserver before the database server stops the load

RECORDEND

Specifies the character that separates records in a delimited text file

REJECTFILE

Sets the full pathname where all coservers write data-conversion errors. If not specified or if files cannot be opened, any error ends the load job abnormally. See also "Reject Files."

SIZE The approximate number of rows in the external table. This can improve performance when external table is used in a join query.

Important: Check constraints on external tables are designed to be evaluated only when loading data. The database server cannot enforce check constraints on external tables because the data can be freely altered outside the control of the database server. If you want to restrict rows that are written to an external table during unload, use a WHERE clause to filter the rows.

Reject Files

Rows that have conversion errors during a load or rows that violate check constraints on the external table are written to a reject file on the coserver that performs the conversion. Each coserver manages its own reject file. The REJECTFILE clause declares the name of the reject file on each coserver.

You can use the formatting characters %c and %n (but not %r) in the filename format. Use the %c formatting characters to make the filenames unique. For more information on how to format characters, see the section "Using Formatting Characters" on page 2-109.

If you perform another load to the same table during the same session, any earlier reject file of the same name is overwritten.

Reject file entries have the following format:

*coserver-number, filename, record, reason-code,
field-name: bad-line*

The following table describes these elements of the reject file:

Element	Description
<i>coserver-number</i>	Number of the coserver from which the file is read
<i>filename</i>	Name of the input file
<i>record</i>	Record number in the input file where the error was detected
<i>reason-code</i>	Description of the error
<i>field-name</i>	External field name where the first error in the line occurred, or <none> if the rejection is not specific to a particular column
<i>bad-line</i>	Line that caused the error (delimited or fixed-position character files only), up to 80 characters

The reject file writes the *coserver-number*, *filename*, *record*, *field-name*, and *reason-code* in ASCII. The *bad-line* information varies with the type of input file.

- For delimited files or fixed-position character files, up to 80 characters of the bad line are copied directly into the reject file.
- For Informix internal data files, the bad line information is not placed in the reject file because you cannot edit the binary representation in a file; but the *coserver-number*, *filename*, *record*, *reason-code*, and *field-name* are still reported in the reject file so you can isolate the problem. Use the Table Options clause to specify the format of the external data.

Errors that can cause a row to be rejected include the following.

Error Text	Explanation
CONSTRAINT <i>constraint name</i>	This constraint was violated.
CONVERT_ERR	Any field encounters a conversion error.
MISSING_DELIMITER	No delimiter was found.
MISSING_RECORDEND	No record end was found.
NOT NULL	A NULL was found in <i>field-name</i> .
ROW_TOO_LONG	The input record is longer than 32 kilobytes.

Examples

The examples in this section show how to specify the DATAFILES field.

Assume that the database server is running on four nodes, and one file is to be read from each node. All files have the same name. The DATAFILES specification can then be as follows:

```
DATAFILES ("DISK:cogroup_all:/work2/unload.dir/mytbl")
```

Now, consider a system with 16 coservers where only three coservers have tape drives attached (for example, coservers 2, 5, and 9). If you define a cogroup for these coservers before you run LOAD and UNLOAD commands, you can use the cogroup name rather than a list of individual coservers when you execute the commands. To set up the cogroup, run **onutil**.

```
% onutil
1> create cogroup tape_group
2> from coserver.2, coserver.5, coserver.9;
Cogroup successfully created.
```

Then define the file locations for named pipes:

```
DATAFILES ("PIPE:tape_group:/usr/local/TAPE.%c")
```

The filenames expand as follows:

```
DATAFILES ("pipe:2:/usr/local/TAPE.2",
           "pipe:5:/usr/local/TAPE.5",
           "pipe:9:/usr/local/TAPE.9")
```

If, instead, you want to process three files on each of two coservers, define the files as follows:

```
DATAFILES ("DISK:1:/work2/extern.dir/mytbl.%r(1..3)",
           "DISK:2:/work2/extern.dir/mytbl.%r(4..6)")
```

The expanded list follows:

```
DATAFILES ("disk:1:/work2/extern.dir/mytbl.1",
           "disk:1:/work2/extern.dir/mytbl.2",
           "disk:1:/work2/extern.dir/mytbl.3",
           "disk:2:/work2/extern.dir/mytbl.4",
           "disk:2:/work2/extern.dir/mytbl.5",
           "disk:2:/work2/extern.dir/mytbl.6")
```

Related Information

Related statements: "INSERT " on page 2-430 and " SET PLOAD FILE " on page 2-645

See also the "INTO Table Clauses" on page 2-574 of "SELECT" on page 2-528.

For more information on external tables, refer to your *IBM Informix Administrator's Reference*.

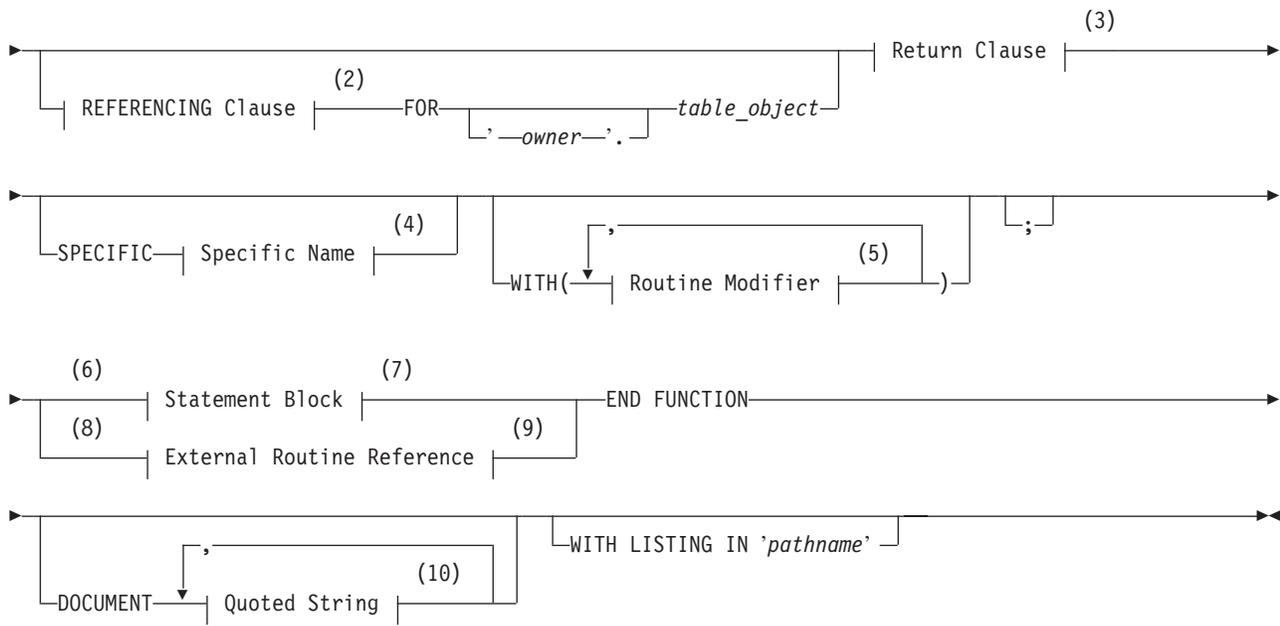
CREATE FUNCTION

Use the CREATE FUNCTION statement to create a user-defined function, to register an external function, or to write and register an SPL function.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```
►► CREATE DBA FUNCTION function ( Routine Parameter List (1) )
```



Notes:

- 1 See "Routine Parameter List" on page 5-67
- 2 See "The REFERENCING and FOR Clauses" on page 2-116
- 3 See "Return Clause " on page 5-55
- 4 See "Specific Name " on page 5-73
- 5 See "Routine Modifier" on page 5-59
- 6 Stored Procedure Language only
- 7 See "Statement Block " on page 5-74
- 8 External routines only
- 9 See " External Routine Reference" on page 5-19
- 10 See "Quoted String" on page 4-189

Element	Description	Restrictions	Syntax
<i>function</i>	Name of new function that is defined here	You must have the appropriate language privileges. See "GRANT " on page 2-395 and "Overloading the Name of a Function" on page 2-118.	"Identifier " on page 5-21
<i>owner</i>	Owner of <i>table_object</i>	Must own <i>table_object</i>	" Owner Name " on page 5-45
<i>pathname</i>	Pathname to a file in which compile-time warnings are stored	The specified pathname must exist on the computer where the database resides	The path and filename must conform to your operating-system rules.
<i>table_object</i>	Name or synonym of the table or view whose triggers can call <i>function</i>	Must exist in the local database	"Identifier " on page 5-21

Tip: If you are trying to create a function from text of source code that is in a separate file, use the CREATE FUNCTION FROM statement.

Usage

Dynamic Server supports user-defined functions written in these languages:

- Stored Procedure Language (SPL)
- One of the external languages (C or Java™) that Dynamic Server supports (*external functions*)

When the IFX_EXTEND_ROLE configuration parameter of is set to ON, only users to whom the DBSA grants the built-in EXTEND role can create external functions.

How many values a function can return is language-dependent. Functions written in SPL can return one or more values. External functions written in the C or Java languages must return exactly one value. But a C function can return a collection type, and external functions in queries can return additional values indirectly from OUT parameters (and for the Java language, from INOUT parameters) that Dynamic Server can process as statement-local variables (SLVs).

For information on how this document uses the terms UDR, function, and procedure as well as recommended usage, see “Relationship Between Routines, Functions, and Procedures” on page 2-155 and “Using CREATE PROCEDURE Versus CREATE FUNCTION” on page 2-155, respectively.

The entire length of a CREATE FUNCTION statement must be less than 64 kilobytes. This length is the literal length of the statement, including white space characters such as blank spaces and tabs.

In ESQ/C, you can use a CREATE FUNCTION statement only within a PREPARE statement. If you want to create a user-defined function for which the text is known at compile time, you must put the text in a file and specify this file with the CREATE FUNCTION FROM statement.

Functions use the collating order that was in effect when they were created. See “SET COLLATION” on page 2-584 for information about using non-default collation.

Privileges Necessary for Using CREATE FUNCTION

You must have the Resource privilege on a database to create a function within that database.

Before you can create an SPL function, you must also have the Usage privilege on the SPL language . For more information, see “Usage Privilege in Stored Procedure Language” on page 2-407.

By default, Usage privilege on SPL is granted to PUBLIC. You must also have at least Resource privilege on a database to create an SPL function in that database.

DBA Keyword and Privileges on the Created Function

If you create a UDR with the DBA keyword, it is known as a *DBA-privileged UDR*. You need the DBA privilege to create a DBA-privileged UDR.

Among users who do not hold the DBA privilege, only those to whom the DBA grants the Execute privilege can invoke the DBA-privileged UDR. If the DBA grants the Execute privilege to PUBLIC, however, then all users can use the

DBA-privileged UDR. For additional information about DBA-privileged UDRs, see “Ownership of Created Database Objects” on page 2-162.

If you omit the DBA keyword, the UDR is known as an *owner-privileged* UDR.

When you create an owner-privileged UDR in an ANSI-compliant database, only you can execute the UDR. Before other users can execute an owner-privileged UDR, its owner must grant the Execute privilege, either to individual users, or to roles, or to PUBLIC.

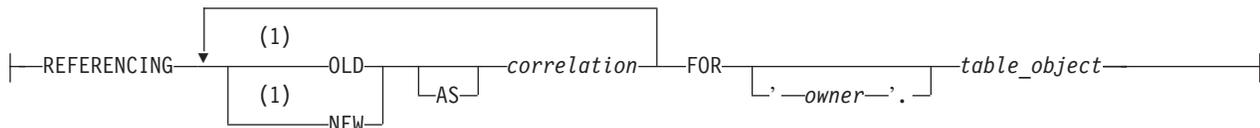
If you create an owner-privileged UDR in a database that is not ANSI compliant, anyone can execute the UDR because PUBLIC is granted the Execute privilege by default. To restrict access to an owner-privileged UDR to specific users, the owner must revoke the Execute privilege on the UDR from PUBLIC, and then grant it to specified users or roles. Setting the **NODEFDAC** environment variable to **yes** prevents privileges on any UDR from being granted to PUBLIC. If this environment variable is set to **yes**, no one besides the owner of the UDR can invoke it unless the owner grants the Execute privilege for that UDR to other users.

If an external C or Java language function has a negator function, you must grant the Execute privilege on both the external function and on its negator function before users can execute the external function.

The REFERENCING and FOR Clauses

The REFERENCING clause can declare correlation names for the original value and for the updated value in columns of the *table_object* that the FOR clause specifies.

REFERENCING and FOR Clauses:



Notes:

- 1 Use path no more than once

Element	Description	Restrictions	Syntax
<i>correlation</i>	Name that you declare here to qualify an old or new column value (as <i>correlation.column</i>) in a trigger routine	Must not be <i>table_object</i>	“Identifier ” on page 5-21
<i>owner</i>	Owner of <i>table_object</i>	Must own <i>table_object</i>	“ Owner Name ” on page 5-45
<i>table_object</i>	Name or synonym of a table or view whose triggers can call <i>function</i>	Must exist in the local database	“Identifier ” on page 5-21

If you include the REFERENCING and FOR *table_object* clauses immediately after the parameter list of the CREATE FUNCTION statement, the function that you

create is known as a *trigger UDR* or *trigger routine*. The FOR clause specifies the table or view whose triggers can invoke the function from the FOR EACH ROW section of their Triggered Action list.

In the REFERENCING clause, the OLD *correlation* specifies a prefix by which the trigger routine can reference the value that a column of *table_object* had before the trigger routine modifies that column value. The NEW *correlation* specifies a prefix for referencing the new value that the trigger routine assigns to the column. Whether the trigger routine can use correlation names to reference the OLD column value, the NEW column value, or both values depends on the type of triggering event:

- A trigger routine invoked by an Insert trigger can reference only the NEW correlation name.
- A trigger routine invoked by a Delete trigger or by a Select trigger can reference only the OLD correlation name.
- A trigger routine invoked by an Update trigger can reference both the OLD and the NEW correlation names.

For information about how to use the *correlation.column* notation in triggered actions, see “REFERENCING Clauses” on page 2-254.

Besides the general requirements for any Dynamic Server UDR that is written in the SPL language, trigger routines can support certain additional syntax features, and are subject to certain restrictions, that are not features (or that are not restrictions) for ordinary UDRs that are not trigger routines:

- A trigger routine must include the FOR *table_object* clause that specifies the name of the table or view in the local database whose triggers can invoke this routine.
- A trigger routine can also include the REFERENCING clause to declare correlation names for OLD and NEW values that SPL statements in the UDR can reference.
- Trigger routines can be invoked only in the FOR EACH ROW section of the Triggered Action list in the trigger definition.
- Correlated variables for OLD or NEW values can appear in the IF statement of SPL and in CASE expressions.
- Correlated variables for OLD values cannot be on the left-hand side of a LET expression.
- Correlated variables for NEW values cannot be on the left-hand side of a LET expression if the FOR clause specifies a view whose INSTEAD OF trigger action list invokes the trigger routine.
- Only correlated variables for NEW values can be on the left-hand side of a LET expression that references correlated variables. In this case, however, the FOR clause must specify a table, rather than a view.
- Both OLD and NEW values can be on the right-hand side of a LET expression.
- The Boolean operators SELECTING, INSERTING, DELETING, and UPDATING are valid in trigger routines (and only in trigger routines and in other UDRs that are invoked in triggered action statements) in contexts where Boolean expressions are valid. These operators return TRUE ('t') if the triggering event matches the DML operation referenced by the name of the operator, and they return FALSE ('f') otherwise.
- If a single triggering event activates multiple triggers on the same table or view, then all of the BEFORE actions take place before any of the FOR EACH ROW actions, and all of the AFTER actions follow the FOR EACH ROW actions. The order of execution of different triggers on the same event is not guaranteed.

- Trigger routines must be written in the SPL language. They cannot be written in an external language, like the C or Java language, but they can include calls to external language routines, such as the **mi_trigger** application programming interface for trigger introspection.
- Trigger functions cannot reference savepoint objects. Any changes to the data values or to the schema of the database by a triggered action must be committed or rolled back in their entirety. Dynamic Server does not support the partial rollback of triggered actions.

For more information about the **mi_trigger** API, refer to the *IBM Informix Datablade API Programmer's Guide* and to the *IBM Informix Datablade API Function Reference*.

If you include the REFERENCING clause but omit the FOR clause, or if you include the FOR clause but omit the REFERENCING clause, the CREATE FUNCTION statement fails with an error.

If you omit the REFERENCING and FOR clauses, the UDR cannot use the SELECTING, INSERTING, DELETING, and UPDATING operators, and cannot declare variables that can represent and manipulate column values in triggered actions on the table or view that the trigger definition specifies.

See the “REFERENCING Clauses” on page 2-254 section in the CREATE TRIGGER statement description for the syntax of the REFERENCING clause for Delete, Insert, Select, and Update triggers.

Overloading the Name of a Function

Because Dynamic Server supports *routine overloading*, you can define more than one function with the same name, but different parameter lists. You might want to overload functions in the following situations:

- You create a user-defined function with the same name as a built-in function (such as **equal**()) to process a new user-defined data type.
- You create *type hierarchies*, in which subtypes inherit data representation and functions from supertypes.
- You create *distinct types*, which are data types that have the same internal storage representation as an existing data type, but have different names and cannot be compared to the source type without casting. Distinct types inherit support functions from their source types.

For a brief description of the routine signature that uniquely identifies each user-defined function, see “Routine Overloading and Routine Signatures (IDS)” on page 5-19.

Using the SPECIFIC Clause to Specify a Specific Name

You can declare a specific name, unique to the database, for a user-defined function. A specific name is useful when you are overloading a function.

DOCUMENT Clause

The quoted string in the DOCUMENT clause provides a synopsis and description of the UDR. The string is stored in the **sysprocbody** system catalog table and is intended for the user of the UDR. Anyone with access to the database can query the **sysprocbody** system catalog table to obtain a description of one or all of the UDRs stored in the database.

For example, the following query obtains a description of the SPL function **update_by_pct**, that “SPL Functions” shows:

```
SELECT data FROM sysprocbody b, sysprocedures p
WHERE b.procid = p.procid
      --join between the two catalog tables
      AND p.procname = 'update_by_pct'
      -- look for procedure named update_by_pct
      AND b.datakey = 'D'-- want user document;
```

The preceding query returns the following text:

```
USAGE: Update a price by a percentage
Enter an integer percentage from 1 - 100
and a part id number
```

A UDR or application program can query the system catalog tables to fetch the DOCUMENT clause and display it for a user.

For C and Java language functions, you can include a DOCUMENT clause at the end of the CREATE FUNCTION statement, whether or not you use the END FUNCTION keywords.

WITH LISTING IN Clause

The WITH LISTING IN clause specifies a filename where compile time warnings are sent. After you compile a UDR, this file holds one or more warning messages.

If you do not use the WITH LISTING IN clause, the compiler does not generate a list of warnings.

On UNIX platforms, if you specify a filename but not a directory, this listing file is created in your home directory on the computer where the database resides. If you do not have a home directory on this computer, the file is created in the root directory (the directory named “/”).

On Windows systems, if you specify a filename but not a directory, this listing file is created in your current working directory if the database is on the local computer. Otherwise, the default directory is %INFORMIXDIR%\bin.

SPL Functions

SPL functions are UDRs written in SPL that return one or more values. To write and register an SPL function, use a CREATE FUNCTION statement. Embed appropriate SQL and SPL statements between the CREATE FUNCTION and END FUNCTION keywords. You can also follow the function with the DOCUMENT and WITH FILE IN options.

SPL functions are parsed, optimized (as far as possible), and stored in the system catalog tables in executable format. The body of an SPL function is stored in the **sysprocbody** system catalog table. Other information about the function is stored in other system catalog tables, including **sysprocedures**, **sysproplan**, and **sysprocauth**. For more information about these system catalog tables, see the *IBM Informix Guide to SQL: Reference*.

The END FUNCTION keywords are required in every SPL function, and a semicolon (;) must follow the clause that immediately precedes the statement block. The following code example creates an SPL function:

```

CREATE FUNCTION update_by_pct ( pct INT, pid CHAR(10))
  RETURNING INT;
  UPDATE inventory SET price = price + price * (pct/100)
    WHERE part_id = pid;
  return (select price from inventory where part_id = pid);
END FUNCTION
DOCUMENT "USAGE: Update a price by a percentage",
  "Enter an integer percentage from 1 - 100",
  "and a part id number"
WITH LISTING IN '/tmp/warn_file';

```

For more information on how to write SPL functions, see the chapter about SPL routines in *IBM Informix Guide to SQL: Tutorial*.

See also the section “Transactions in SPL Routines” on page 5-79.

You can include valid SQL or SPL language statements in SPL functions. See, however, the following sections in Chapter 5, “Other Syntax Segments,” on page 5-1 that describe restrictions on SQL and SPL statements within SPL routines: “Subset of SPL Statements Valid in the Statement Block” on page 5-75; “SQL Statements Valid in SPL Statement Blocks” on page 5-76; and “Restrictions on SPL Routines in Data-Manipulation Statements” on page 5-78.

External Procedures (IDS)

External functions are functions you write in an external language (that is, a programming language other than SPL) that Dynamic Server supports.

To create a C user-defined function

1. Write the C function.
2. Compile the function and store the compiled code in a shared library (the shared-object file for C).
3. Register the function in the database server with the CREATE FUNCTION statement.

To create a user-defined function written in the Java language

1. Write a Java static method, which can use the JDBC functions to interact with the database server.
2. Compile the Java source file and create a **.jar** file (the shared-object file for Java).
3. Execute the **install_jar()** procedure with the EXECUTE PROCEDURE statement to install the JAR file in the current database.
4. If the UDR uses user-defined types, create a map between SQL data types and Java classes. Use the **setUDTextName()** procedure that is explained in “EXECUTE PROCEDURE ” on page 2-364.
5. Register the UDR with the CREATE FUNCTION statement.

Rather than storing the body of an external routine directly in the database, the database server stores only the pathname of the shared-object file that contains the compiled version of the routine. When it executes the external routine, the database server invokes the external object code.

The database server stores information about an external function in system catalog tables, including **sysprocbody** and **sysprocauth**. For more information on the system catalog, see the *IBM Informix Guide to SQL: Reference*.

Example of Registering a C User-Defined Function

The following example registers an external C user-defined function named `equal()` in the database. This function takes two arguments of the type `basetype1` and returns a single Boolean value. The external routine reference name specifies the path to the C shared library where the function object code is actually stored. This library contains a C function `basetype1_equal()`, which is invoked during execution of the `equal()` function.

```
CREATE FUNCTION equal ( arg1 basetype1, arg2 basetype1)
  RETURNING BOOLEAN;
  EXTERNAL NAME
    "/usr/lib/basetype1/lib/libbtype1.so(basetype1_equal)"
  LANGUAGE C
END FUNCTION;
```

Example of Registering a UDR Written in the Java Language

The following `CREATE FUNCTION` statement registers the user-defined function, `sql_explosive_reaction()`. This function is discussed in “`sqlj.install_jar`” on page 6-58.

```
CREATE FUNCTION sql_explosive_reaction(INT) RETURNS INT WITH (class="jvp")
  EXTERNAL NAME "course_jar:Chemistry.explosiveReaction" LANGUAGE JAVA;
```

This function returns a single `INTEGER` value. The `EXTERNAL NAME` clause specifies that the Java implementation of the `sql_explosive_reaction()` function is a method called `explosiveReaction()`, which resides in the `Chemistry` Java class that resides in the `course_jar` JAR file.

Ownership of Created Database Objects

The user who creates an owner-privileged UDR, rather than the user who executes the UDR, owns any database objects that are created by the UDR when the UDR is executed, unless another owner is specified for the created object.

For example, assume that user `mike` creates this user-defined function:

```
CREATE FUNCTION func1 () RETURNING INT;
  CREATE TABLE tab1 (colx INT);
  RETURN 1;
END FUNCTION;
```

If user `joan` now executes function `func1`, user `mike`, not user `joan`, is the owner of the newly created table `tab1`.

In the case of a DBA-privileged UDR, however, the user who executes a UDR (rather than the UDR owner) owns any database objects created by the UDR, unless another owner is specified for the database object within the UDR.

For example, assume that user `mike` creates this user-defined function:

```
CREATE DBA FUNCTION func2 () RETURNING INT;
  CREATE TABLE tab2 (coly INT);
  RETURN 1;
END FUNCTION;
```

If user `joan` now executes function `func2`, user `joan`, not user `mike`, is the owner of the newly created table `tab2`.

See also the section “Support for Roles and User Identity” on page 5-79.

Examples

Overloaded functions are uniquely identified by the name and the input parameter list. Instead of providing a long unique identifier, it is possible to provide specific name and use it later. The following example illustrates an overloaded function, whose identifier is **getArea**, that has the specific names **getSquareArea** and **getRectangleArea**:

```
CREATE FUNCTION getArea
  (i INT DEFAULT 0)
RETURNING INT SPECIFIC getSquareArea;
DEFINE j INT;
LET j = i * i;
RETURN j;
END FUNCTION;

CREATE FUNCTION getArea
  (i INT DEFAULT 0, j INT DEFAULT 0)
RETURNING INT SPECIFIC getRectangleArea;
DEFINE k INT;
LET k = i * j;
RETURN k;
END FUNCTION;
```

Now you can use the specific name, as in the following example:

```
GRANT EXECUTE ON SPECIFIC FUNCTION getSquareArea TO informix;
GRANT EXECUTE ON SPECIFIC FUNCTION getRectangleArea TO informix;
```

Without the specific name, you would need to issue the following:

```
GRANT EXECUTE ON FUNCTION getArea (INTEGER) TO informix;
GRANT EXECUTE ON FUNCTION getArea (INTEGER,INTEGER) TO informix;
```

Related Information

Related statements: “ALTER FUNCTION ” on page 2-28, “ALTER ROUTINE” on page 2-33, “CREATE PROCEDURE ” on page 2-153, “CREATE FUNCTION FROM ,” “DROP FUNCTION ” on page 2-327, “DROP ROUTINE ” on page 2-335, “GRANT ” on page 2-395, “EXECUTE FUNCTION ” on page 2-356, “PREPARE ” on page 2-468, “REVOKE” on page 2-494, and “UPDATE STATISTICS ” on page 2-705

Chapter 3, “SPL Statements,” on page 3-1 of this document describes the SPL language. For a discussion of how to create and use SPL routines, see the *IBM Informix Guide to SQL: Tutorial*.

For a discussion of how to create and use external routines, see *IBM Informix User-Defined Routines and Data Types Developer’s Guide*. For information about how to create C UDRs, see the *IBM Informix Datablade API Programmer’s Guide*.

CREATE FUNCTION FROM

Use the CREATE FUNCTION FROM statement to access a user-defined function whose CREATE FUNCTION statement resides in a separate file.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL. Use this statement with ESQL/C.

Syntax

► CREATE FUNCTION FROM 'file' file_var ►

Element	Description	Restrictions	Syntax
<i>file</i>	Path and filename of a file that contains the full CREATE FUNCTION statement text. Default pathname is current directory.	Must exist and contain exactly one CREATE FUNCTION statement	Must conform to operating-system rules.
<i>file_var</i>	Variable storing value of <i>file</i>	Same as for <i>file</i>	Language specific

Usage

Functions written in the C or Java language are called *external* functions. When the IFX_EXTEND_ROLE configuration parameter is set to ON, only users who have been granted the built-in EXTEND role can create external functions.

An Informix ESQL/C program cannot directly create a user-defined function. That is, it cannot contain the CREATE FUNCTION statement.

To create these functions within an Informix ESQL/C program

1. Create a source file with the CREATE FUNCTION statement.
2. Use the CREATE FUNCTION FROM statement to send the contents of this source file to the database server for execution.

The file that you specify in the *file* parameter can contain only one CREATE FUNCTION statement.

For example, suppose that the following CREATE FUNCTION statement is in a separate file, called **del_ord.sql**:

```
CREATE FUNCTION delete_order( p_order_num INT) RETURNING INT, INT;
  DEFINE item_count INT;
  SELECT count(*) INTO item_count FROM items
    WHERE order_num = p_order_num;
  DELETE FROM orders WHERE order_num = p_order_num;
  RETURN p_order_num, item_count;
END FUNCTION;
```

In the Informix ESQL/C program, you can access the **delete_order()** SPL function with the following CREATE FUNCTION FROM statement:

```
EXEC SQL create function from 'del_ord.sql';
```

If you are not sure whether the UDR in the file is a user-defined function or a user-defined procedure, use the CREATE ROUTINE FROM statement.

The filename that you provide is relative. If you provide a simple filename with no pathname (as in the preceding example), the client application looks for the file in the current directory.

Important: The Informix ESQL/C preprocessor does not process the contents of the file that you specify. It only sends the contents to the database server for execution. Therefore, there is no syntactic check that the file that you specify in CREATE FUNCTION FROM actually contains a CREATE FUNCTION statement. To improve readability of the code, however, it is recommended that you match these two statements.

Related Information

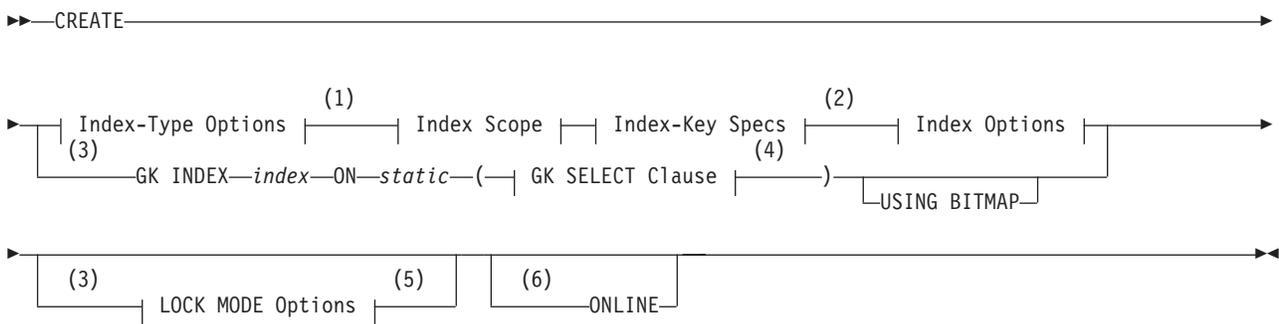
Related statements: “CREATE FUNCTION ” on page 2-113, “CREATE PROCEDURE ” on page 2-153, “CREATE PROCEDURE FROM ” on page 2-163, and “CREATE ROUTINE FROM ” on page 2-166

CREATE INDEX

Use the CREATE INDEX statement to create an index for one or more columns in a table, or on values returned by a UDR that uses column values as arguments.

This statement is an extension to the ANSI/ISO standard for SQL.

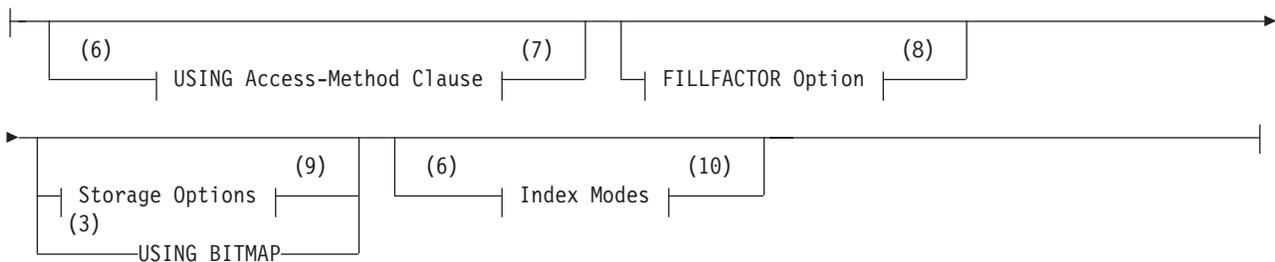
Syntax



Index Scope:



Index Options:



Notes:

- 1 See “ Index-Type Options” on page 2-125
- 2 See “ Index-Key Specification” on page 2-127
- 3 Extended Parallel Server only
- 4 See “ SELECT Clause for Generalized-Key Index” on page 2-141
- 5 See “ LOCK MODE Options (XPS)” on page 2-141
- 6 Dynamic Server only

- 7 See “USING Access-Method Clause (IDS)” on page 2-132
- 8 See “FILLFACTOR Option” on page 2-133
- 9 See “Storage Options” on page 2-134
- 10 See “ Index Modes (IDS)” on page 2-138

Element	Description	Restrictions	Syntax
<i>index</i>	Name declared here for a new index	Must be unique among names of indexes in the database	“Identifier ” on page 5-21
<i>static</i>	Table on which a Generalized Key index is created (XPS)	Table must exist and be static; it cannot be a virtual table	“Identifier ” on page 5-21
<i>synonym, table</i>	Name or synonym of a standard or temporary <i>table</i> to be indexed	Synonym and its table must exist in the current database	“Identifier ” on page 5-21

Usage

When you issue the CREATE INDEX statement, the table is locked in exclusive mode. If another process is using the table, CREATE INDEX returns an error. (For an exception, however, see “The ONLINE Keyword (IDS)” on page 2-143.)

If the index is on a column that stores encrypted data, Dynamic Server cannot use the index.

Indexes use the collation that was in effect when CREATE INDEX executed.

A *secondary-access method* (sometimes referred to as an *index-access method*) is a set of database server functions that build, access, and manipulate an index structure such as a B-tree, R-tree, or an index structure that a DataBlade® module provides, in order to speed up the retrieval of data.

In Dynamic Server, neither *synonym* nor *table* can refer to a virtual table.

You cannot directly base a functional index on a built-in function, but you can create an SPL wrapper that calls and returns a value from a built-in function. The arguments to a user-defined function that defines a functional index cannot be the values from a column of a collection data type.

If you are using Extended Parallel Server, use the USING BITMAP keywords to store the list of records in each key of the index as a compressed bitmap. The storage option is not compatible with a bitmap index because bitmap indexes must be fragmented in the same way as the table.

Index-Type Options

The index-type options let you specify attributes of the index.

Index-Type Options:



UNIQUE or DISTINCT Option

Use the **UNIQUE** or **DISTINCT** keyword to require that the columns on which the index is based accept only unique data. If you do not specify the **UNIQUE** or **DISTINCT** keyword, the index allows duplicate values in the indexed column (or in the set of indexed columns). The following example creates a unique index:

```
CREATE UNIQUE INDEX c_num_ix ON customer (customer_num);
```

A unique index prevents duplicate values in the **customer_num** column. A column with a unique index can have, at most, one **NULL** value.

The **DISTINCT** and **UNIQUE** keywords are synonyms in this context, so the following statement has the same effect as the previous example:

```
CREATE DISTINCT INDEX c_num_ix ON customer (customer_num);
```

The index in both examples is maintained in ascending order, which is the default order.

You can also prevent duplicates in a column or set of columns by creating a unique constraint with the **CREATE TABLE** or **ALTER TABLE** statement. You cannot specify an R-tree secondary-access method for a **UNIQUE** index key. For more information on how to create unique constraints, see the **CREATE TABLE** or **ALTER TABLE** statements.

See also the section “Differences Between a Unique Constraint and a Unique Index” on page 2-199.

How Indexes Affect Primary-Key, Unique, and Referential Constraints

The database server creates internal B-tree indexes for primary-key, unique, and referential constraints. If a primary-key, unique, or referential constraint is added after the table is created, any user-created indexes on the constrained columns are used, if appropriate. An appropriate index is one that indexes the same columns that are used in the primary-key, referential, or unique constraint. If an appropriate user-created index is not available, the database server creates a nonfragmented internal index on the constrained column or columns.

CLUSTER Option

Use the **CLUSTER** keyword to reorder the rows of the table in the order that the index designates. The **CREATE CLUSTER INDEX** statement fails if a **CLUSTER** index already exists on the same table.

```
CREATE CLUSTER INDEX c_clust_ix ON customer (zipcode);
```

This statement creates an index on the **customer** table and physically orders the rows according to their postal code values, in (by default) ascending order.

If the **CLUSTER** option is specified in addition to **FRAGMENTS ON**, the data values are clustered only within each fragment, and not globally across the entire table.

In Dynamic Server, you cannot specify the **CLUSTER** option and the **ONLINE** keyword in the same statement. In addition, some secondary-access methods (such as R-tree) do not support clustering. Before you specify **CLUSTER** for your index, be sure that the index uses an access method that supports clustering.

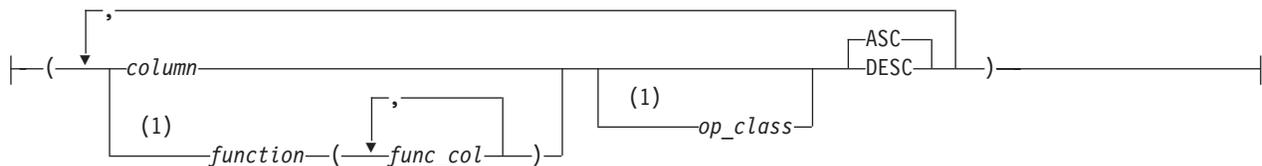
If you are using Extended Parallel Server, you cannot use the **CLUSTER** option on **STANDARD** tables. In addition, you cannot specify the **CLUSTER** option and

storage options in the same CREATE INDEX statement (see “Storage Options” on page 2-134). When you create a clustered index, the **constrid** of any unique or referential constraints on the associated table changes. The **constrid** is stored in the **sysconstraints** system catalog table.

Index-Key Specification

Use the Index-Key Specification of the CREATE INDEX statement to define the key value for the index, to specify whether to sort the index in ascending or descending order, and to identify a default operator class if the secondary access method in the USING clause has no default operator class, or to override its default operator class.

Index-Key Specification:



Notes:

1 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>column</i>	Column or columns used as a key to this index	See “Using a Column or Column List as the Index Key” on page 2-128.	“Identifier ” on page 5-21
<i>function</i>	User-defined function used as a key to this index	Must be a nonvariant function that does not return a large object data type. Cannot be a built-in algebraic, exponential, log, or hex function.	“Identifier ” on page 5-21
<i>func_col</i>	Columns whose values are arguments to <i>function</i>	Cannot be of a collection data type. See “Using a Function as an Index Key (IDS)” on page 2-128.	“Identifier ” on page 5-21
<i>op_class</i>	Operator class associated with <i>column</i> or <i>function</i> for this index	If the secondary-access method in the USING clause has no default operator class, you must specify one here. (See “Using an Operator Class (IDS)” on page 2-132.)	“Identifier ” on page 5-21

The index-key value can be one or more columns that contain built-in data types. If you specify multiple columns, the concatenation of values from the set of columns is treated as a single composite column for indexing.

In Dynamic Server, the index-key value also can be one of the following:

- A column of type `LVARCHAR(size)`, if *size* is smaller than 387 bytes
- One or more columns that contain user-defined data types
- One or more values that a user-defined function returns (referred to as a *functional index*)
- A combination of columns and functions

The 387-byte `LVARCHAR` size limit is for dbspaces of the default (2 kilobyte) page size, but dbspaces of larger page sizes can support larger index key sizes, as listed in the following table.

Table 2-1. Maximum Index Key Size for Selected Page Sizes

Page Size	Maximum Index Key Size
2 kilobytes	387 bytes
4 kilobytes	796 bytes
8 kilobytes	1,615 bytes
12 kilobytes	2,435 bytes
16 kilobytes	3,245 bytes

Using a Column or Column List as the Index Key

These restrictions apply to a column or column list specified as the index key:

- All the columns must exist and must be in the table being indexed.
- The maximum number of columns and total width of all columns depends on the database server. See “Creating Composite Indexes” on page 2-129.
- You cannot add an ascending index to a column or column list that already has a unique constraint on it. See “Using the ASC and DESC Sort-Order Options” on page 2-129.
- You cannot add a unique index to a column or to a column list that has a primary-key constraint. The reason is that defining the column or column list as the primary key causes the database server to create a unique internal index on the column or column list; you cannot create another unique index on this column or column list with the CREATE INDEX statement.
- The number of indexes that you can create on the same column or the same set of columns is restricted. See “Restrictions on the Number of Indexes on a Set of Columns” on page 2-131.

In Dynamic Server, these additional restrictions apply to indexes:

- You cannot create an index on a column of an external table.
- The column cannot be of a collection data type.

Using a Function as an Index Key (IDS)

A *functional index* is indexed on the value that a specified function returns from a column-value argument, rather than on the value of a column. For example, the following statement creates a functional index on table **zones** using the value that the user-defined function **Area()** returns as the key:

```
CREATE INDEX zone_func_ind ON zones (Area(length,width));
```

You can create functional indexes within an SPL routine. You can also create an index on a nonvariant user-defined function that does not return a large object. The arguments to the UDF cannot be from a column of a collection data type.

The function must be a user-defined function. You cannot create a functional index on any built-in function of SQL. You can, however, create a functional index on a user-defined function that calls a built in function and uses the value returned by the built-in function as the index key of a functional index.

The functional index can be a B-tree index, an R-tree index, or a user-defined secondary-access method. The ONLINE keyword, however, is not valid when you create a functional index; see “The ONLINE Keyword (IDS)” on page 2-143.

Creating Composite Indexes

A *simple* index lists only one *column* (or for IDS, only one *column* or *function*) in its Index Key Specification. Any other index is a *composite* index. You should list the columns in a composite index in the order from most frequently used to least frequently used.

In Dynamic Server, if you use SET COLLATION to specify the collating order of a non-default locale, you can create multiple indexes on the same set of columns, using different collations. (Such indexes are useful only on NCHAR or NVARCHAR columns.)

The following example creates a composite index using the **stock_num** and **manu_code** columns of the **stock** table:

```
CREATE UNIQUE INDEX st_man_ix ON stock (stock_num, manu_code);
```

The UNIQUE keyword prevents any duplicates of a given combination of **stock_num** and **manu_code**. The index is in ascending order by default.

You can include up to 16 columns in a composite index. The total width of all indexed columns in a single composite index cannot exceed 380 bytes.

In Dynamic Server, an *index key part* is either a column in a table, or the result of a user-defined function on one or more columns. A composite index can have up to 16 key parts that are columns, or up to 341 key parts that are values returned by a UDR. This limit is language-dependent and applies to UDRs written in SPL or Java; functional indexes based on C language UDRs can have up to 102 key parts. A composite index can have any of the following items as an index key:

- One or more columns
- One or more values that a user-defined function returns (referred to as a functional index)
- A combination of columns and user-defined functions

For dbspaces of the default page size of 2 kilobytes, the total width of all indexed columns in a single CREATE INDEX statement cannot exceed 387 bytes, except for functional indexes of Dynamic Server, whose language-dependent limits are described earlier in this section. For the maximum sizes in dbspaces larger than 2 kilobytes, see “Index-Key Specification” on page 2-127.

Whether the index is based directly on column values in the table, or on functions that take column values as arguments, the maximum size of the index key depends only on page size. The maximum index key size for functional indexes in dbspaces larger than 2 kilobytes are the same as for non-functional indexes. The only difference between limits on column indexes and functional indexes is the number of key parts. A column based index can have no more than 16 key parts and a functional index has different language-dependent limits on key parts. For a given page size, the maximum index key size is the same for both column-based and functional indexes.

Using the ASC and DESC Sort-Order Options

The ASC option specifies an index maintained in ascending order; this is the default order. The DESC option can specify an index that is maintained in descending order. These ASC and DESC options are valid with B-trees only.

Effects of Unique Constraints on Sort Order Options

When a column or list of columns is defined as unique in a CREATE TABLE or ALTER TABLE statement, the database server implements that UNIQUE CONSTRAINT by creating a unique ascending index. Thus, you cannot use the CREATE INDEX statement to add an ascending index to a column or column list that is already defined as unique.

However, you can create a descending index on such columns, and you can include such columns in composite ascending indexes in different combinations. For example, the following sequence of statements is valid:

```
CREATE TABLE customer (  
  customer_num SERIAL(101) UNIQUE,  
  fname          CHAR(15),  
  lname          CHAR(15),  
  company        CHAR(20),  
  address1       CHAR(20),  
  address2       CHAR(20),  
  city           CHAR(15),  
  state          CHAR(2),  
  zipcode        CHAR(5),  
  phone          CHAR(18)  
);  
  
CREATE INDEX c_temp1 ON customer (customer_num DESC);  
CREATE INDEX c_temp2 ON customer (customer_num, zipcode);
```

In this example, the **customer_num** column has a unique constraint placed on it. The first CREATE INDEX statement places an index sorted in descending order on the **customer_num** column. The second CREATE INDEX includes the **customer_num** column as part of a composite index. For more information on composite indexes, see “Creating Composite Indexes” on page 2-129.

Bidirectional Traversal of Indexes

If you do not specify the ASC or DESC keywords when you create an index on a single column, key values are stored in ascending order by default; but the bidirectional-traversal capability of the database server lets you create just one index on a column and use that index for queries that specify sorting of results in either ascending or descending order of the sort column.

Because of this capability, it does not matter whether you create a single-column index as an ascending or descending index. Whichever storage order you choose for an index, the database server can traverse that index in ascending or descending order when it processes queries.

If you create a composite index on a table, however, the ASC and DESC keywords might be required. For example, if you want to enter a SELECT statement whose ORDER BY clause sorts on multiple columns and sorts each column in a different order, and you want to use an index for this query, you need to create a composite index that corresponds to the ORDER BY columns. For example, suppose that you want to enter the following query:

```
SELECT stock_num, manu_code, description, unit_price  
FROM stock ORDER BY manu_code ASC, unit_price DESC;
```

This query sorts first in ascending order by the value of the **manu_code** column and then in descending order by the value of the **unit_price** column. To use an index for this query, you need to issue a CREATE INDEX statement that corresponds to the requirements of the ORDER BY clause. For example, you can enter either of the following statements to create the index:

```
CREATE INDEX stock_idx1 ON stock
  (manu_code ASC, unit_price DESC);
CREATE INDEX stock_idx2 ON stock
  (manu_code DESC, unit_price ASC);
```

The composite index that was used for this query (**stock_idx1** or **stock_idx2**) cannot be used for queries in which you specify the same sort direction for the two columns in the ORDER BY clause. For example, suppose that you want to enter the following queries:

```
SELECT stock_num, manu_code, description, unit_price
  FROM stock ORDER BY manu_code ASC, unit_price ASC;
SELECT stock_num, manu_code, description, unit_price
  FROM stock ORDER BY manu_code DESC, unit_price DESC;
```

If you want to use a composite index to improve the performance of these queries, you need to enter one of the following CREATE INDEX statements. You can use either one of the created indexes (**stock_idx3** or **stock_idx4**) to improve the performance of the preceding queries.

```
CREATE INDEX stock_idx3 ON stock
  (manu_code ASC, unit_price ASC);
CREATE INDEX stock_idx4 ON stock
  (manu_code DESC, unit_price DESC);
```

You can create no more than one ascending index and one descending index on a column. Because of the bidirectional-traversal capability of the database server, you only need to create one of the indexes. Creating both would achieve exactly the same results for an ascending or descending sort on the **stock_num** column.

After INSERT or DELETE operations are performed on an indexed table, the number of index entries can vary within a page, and the number of index pages that a table requires can depend on whether the index specifies ascending or descending order. For some load and DML operations, a descending single-column or multi-column index might cause the database server to allocate more index pages than an ascending index requires.

Restrictions on the Number of Indexes on a Set of Columns

You can create multiple indexes on a set of columns, provided that each index has a unique combination of ascending and descending columns. For example, to create all possible indexes on the **stock_num** and **manu_code** columns of the **stock** table, you could create four indexes:

- The **ix1** index on both columns in ascending order
- The **ix2** index on both columns in descending order
- The **ix3** index on **stock_num** in ascending order and on **manu_code** in descending order
- The **ix4** index on **stock_num** in descending order and on **manu_code** in ascending order

Because of the bidirectional-traversal capability of the database server, you do not need to create these four indexes. You only need to create two indexes:

- The **ix1** and **ix2** indexes achieve the same results for sorts in which the user specifies the same sort direction (ascending or descending) for both columns, so you only need one index of this pair.
- The **ix3** and **ix4** indexes achieve the same results for sorts in which the user specifies different sort directions for the two columns (ascending on the first

column and descending on the second column or vice versa). Thus, you only need to create one index of this pair. (See also "Bidirectional Traversal of Indexes" on page 2-130.)

Dynamic Server can also support multiple indexes on the same combination of ascending and descending columns, if each index has a different collating order; see "SET COLLATION" on page 2-584.

Using an Operator Class (IDS)

An *operator class* is the set of operators associated with a secondary-access method for query optimization and building the index. You must specify an operator class when you create an index if either one of the following is true:

- No default operator class for the secondary-access method exists. (A user-defined access method can provide no default operator class.)
- You want to use an operator class that is different from the default operator class that the secondary-access method provides.

If you use an alternative access method, and if the access method has a default operator class, you can omit the operator class here; but if you do not specify an operator class and the secondary-access method does not have a default operator class, the database server returns an error. For more information, see "Default Operator Classes" on page 2-153. The following CREATE INDEX statement creates a B-tree index on the **cust_tab** table that uses the **abs_btree_ops** operator class for the **cust_num** key:

```
CREATE INDEX c_num1_ix ON cust_tab (cust_num abs_btree_ops);
```

USING Access-Method Clause (IDS)

The USING clause specifies the secondary-access method for the new index.

USING Access-Method Clause:

```
|—USING—sec_acc_method—(—parameter = value—)
```

Element	Description	Restrictions	Syntax
<i>parameter</i>	Secondary-access-method parameter for this index	See the user documentation for your user-defined access method	"Quoted String" on page 4-189
<i>sec_acc_method</i>	Secondary-access method for this index	Method can be a B-tree, R-tree, or user-defined access method, such as one that a DataBlade module defines	"Identifier " on page 5-21
<i>value</i>	Value of the specified <i>parameter</i>	Must be a valid literal value for <i>parameter</i> in this secondary-access method	"Quoted String" on page 4-189 or "Literal Number " on page 4-185

A *secondary-access method* is a set of routines that perform all of the operations needed for an index, such as create, drop, insert, delete, update, and scan.

The database server provides the following secondary-access methods:

- The generic B-tree index is the built-in secondary-access method.

A B-tree index is good for a query that retrieves a range of data values. The database server implements this secondary-access method and registers it as **btree** in the system catalog tables.

- The R-tree method is a registered secondary-access method.

An R-tree index is good for searches on multidimensional data. The database server registers this secondary-access method as **rtree** in the system catalog tables of a database. An R-tree secondary-access method is not valid for a UNIQUE index key. An R-tree index cannot be clustered, and cannot be stored in a dbspace that has a non-default page size. For more information on R-tree indexes, see the *IBM Informix R-Tree Index User's Guide*.

The access method that you specify must be registered in the **sysams** system catalog table. The default secondary-access method is B-tree.

If the access method is B-tree, you can create only one index for each unique combination of ascending and descending columnar or functional keys with operator classes. (This restriction does not apply to other secondary-access methods.) By default, CREATE INDEX creates a generic B-tree index. If you want to create an index with a secondary-access method other than B-tree, you must specify the name of the secondary-access method in the USING clause.

Some user-defined access methods are packaged as DataBlade modules. Some DataBlade modules provide indexes that require specific parameters when you create them. For more information about user-defined access methods, refer to the documentation of your secondary access-method or DataBlade module.

The following example (for a database that implements R-tree indexes) creates an R-tree index on the **location** column that contains an opaque data type, **point**, and performs a query with a filter on the **location** column.

```
CREATE INDEX loc_ix ON TABLE emp (location) USING rtree;
SELECT name FROM emp WHERE location N_equator_equals point('500, 0');
```

The following CREATE INDEX statement creates an index that uses the **fulltext** secondary-access method, which takes two parameters: WORD_SUPPORT and PHRASE_SUPPORT. It indexes a table **t**, which has two columns: **i**, an integer column, and **data**, a TEXT column.

```
CREATE INDEX tx ON t(data)
  USING fulltext (WORD_SUPPORT='PATTERN',
  PHRASE_SUPPORT='MAXIMUM');
```

FILLFACTOR Option

The FILLFACTOR option takes effect only in the following cases:

- when you build an index on a table that contains more than 5,000 rows and that uses more than 100 table pages
- when you create an index on a fragmented table
- when you create a fragmented index on a nonfragmented table.

Use the FILLFACTOR option to provide for expansion of an index at a later date or to create compacted indexes.

FILLFACTOR Option:

|—FILLFACTOR—*percent*—|

Element	Description	Restrictions	Syntax
<i>percent</i>	Percentage of each index page that is filled by index data when the index is created. The default is 90.	$1 \leq \textit{percent} \leq 100$	"Literal Number " on page 4-185

When the index is created, the database server initially fills only that percentage of the nodes specified with the FILLFACTOR value.

The FILLFACTOR can also be set as a parameter in the ONCONFIG file. The FILLFACTOR clause on the CREATE INDEX statement overrides the setting in the ONCONFIG file. For more information about the ONCONFIG file and the parameters you can use, see your *IBM Informix Administrator's Guide*.

Providing a Low Percentage Value

If you provide a low percentage value, such as 50, you allow room for growth in your index. The nodes of the index initially fill to a certain percentage and contain space for inserts. The amount of available space depends on the number of keys in each page as well as the percentage value.

For example, with a 50-percent FILLFACTOR value, the page would be half full and could accommodate doubling in size. A low percentage value can result in faster inserts and can be used for indexes that you expect to grow.

Providing a High Percentage Value

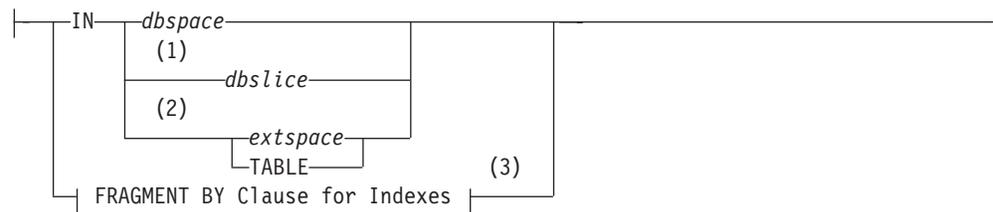
If you provide a high percentage value, such as 99, indexes are compacted, and any new index inserts result in splitting nodes. The maximum density is 100 percent. With a 100-percent FILLFACTOR value, the index has no room available for growth; any addition to the index results in splitting the nodes.

A 99-percent FILLFACTOR value allows room for at least one insertion per node. A high percentage value can result in faster queries and is appropriate for indexes that you do not expect to grow, or for mostly read-only indexes.

Storage Options

The storage options specify the distribution scheme of an index. You can use the IN clause to specify a storage space for the entire index, or you can use the FRAGMENT BY clause to fragment the index across multiple storage spaces.

Storage Options:



Notes:

- 1 Extended Parallel Server only
- 2 Dynamic Server only
- 3 See "FRAGMENT BY Clause for Indexes" on page 2-136

Element	Description	Restrictions	Syntax
<i>dbslice</i>	The dbslice that contains all of the index fragments	Must exist	“Identifier ” on page 5-21
<i>dbspace</i>	The dbspace in which to store the index	Must exist	“Identifier ” on page 5-21
<i>extspace</i>	Name assigned by the onspaces command to a storage area outside the database server	Must exist	See the documentation for your access method.

If you specify any storage option (except IN TABLE), you create a *detached index*. Detached indexes are indexes that are created with a specified distribution scheme. Even if the distribution scheme specified for the index is identical to that specified for the table, the index is still considered to be detached. If the distribution scheme of a table changes, all detached indexes continue to use the distribution scheme that the Storage Option clause specified.

For information on locally detached and globally detached indexes, see “FRAGMENT BY Clause for Indexes” on page 2-136. If you are using Extended Parallel Server, you cannot use the CLUSTER option and storage options in the same CREATE INDEX statement. See “CLUSTER Option” on page 2-126.

IN Clause

Use the IN clause to specify a storage space to hold the entire index. The storage space that you specify must already exist.

Storing an Index in a dbspace

Use the IN *dbspace* clause to specify the dbspace where you want your index to reside. When you use this clause with any option except the TABLE keyword, you create a detached index.

The IN *dbspace* clause allows you to isolate an index. For example, if the **customer** table is created in the **custdata** dbspace, but you want to create an index in a separate dbspace called **custind**, use the following statements:

```
CREATE TABLE customer
    .
    .
    .
    IN custdata EXTENT SIZE 16;

CREATE INDEX idx_cust ON customer (customer_num) IN custind;
```

Storing an Index Fragment in a Named Partition (IDS)

Besides the option of storing a fragment of the index in a dbspace, Dynamic Server supports storing fragments of the index in a named subset of the dbspace, called a *partition*. Unless you explicitly declare names for the fragments in the PARTITION BY or FRAGMENT BY clause, each fragment, by default, has the same name as the dbspace where it resides. This includes all fragmented tables and indexes migrated from earlier releases of Dynamic Server.

Storing an Index in a dbslice (XPS)

Using Extended Parallel Server, the IN *dbslice* clause allows you to fragment an index across multiple dbspaces. The database server fragments the table by round-robin in the dbspaces that make up the dbslice when the table is created.

Storing Data in an extspace (IDS)

In general, use the *extspace* storage option in conjunction with the “USING Access-Method Clause (IDS)” on page 2-132. For more information, refer to the user documentation for your custom-access method.

Creating an Attached Index with the IN TABLE Keywords (IDS)

In some earlier releases of Dynamic Server, if you did not use the storage options to specify a distribution scheme, then, by default, the index used the same distribution scheme as the table on which it was built. Such an index is called an *attached index*. If you omit the Storage Options clause, Dynamic Server creates new indexes as detached indexes by default, but supports existing attached indexes created by earlier release versions. (The **DEFAULT_ATTACH** environment variable, however, can override the default behavior so that the new index is attached.)

You can also specify **IN TABLE** as the storage option to create an attached index, even if the **DEFAULT_ATTACH** environment variable is not set. An attached index is created in the same dbspace (or dbspaces, if the table is fragmented) as the table on which it is built. If the distribution scheme of a table changes, all attached indexes start using the new distribution scheme.

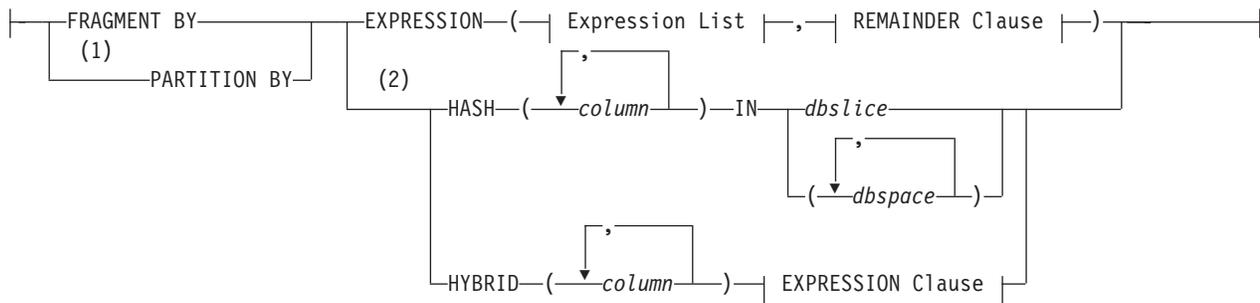
Only B-tree indexes that are nonfragmented and that are on nonfragmented tables can be attached. All other indexes, including extensibility related indexes, such as R-trees and functional indexes, must be detached. You cannot create an attached index using a collating order different from that of the table, nor different from what **DB_LOCALE** specifies. For information about the **DB_LOCALE** and **DEFAULT_ATTACH** environment variables, see the *IBM Informix Guide to SQL: Reference*.

Important: Attached indexes are supported in this release for backward compatibility with Dynamic Server 7.x, but IBM does not recommend use of the **DEFAULT_ATTACH** environment variable nor of the **IN TABLE** storage option in new applications. Attached indexes are a deprecated feature that might not be supported in some future release of Dynamic Server.

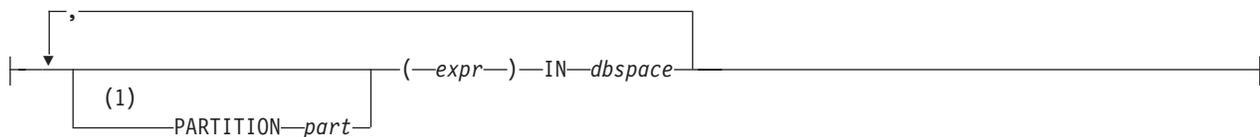
FRAGMENT BY Clause for Indexes

Use the **FRAGMENT BY** clause to create a detached index and to define its fragmentation strategy across dbspaces, partitions (IDS), or dbslices (XPS).

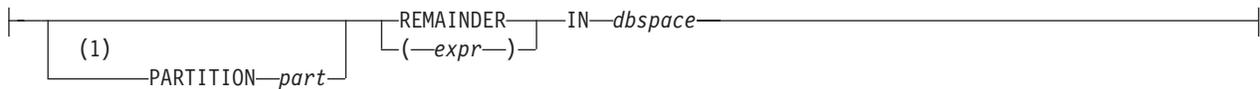
FRAGMENT BY Clause for Indexes:



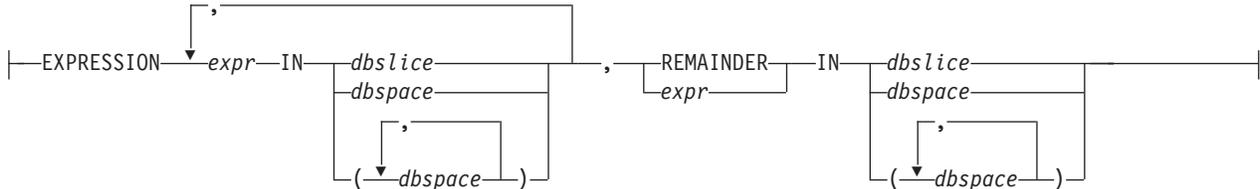
Expression List:



REMAINDER Clause:



EXPRESSION Clause:



Notes:

- 1 Dynamic Server only
- 2 Extended Parallel Server only

Element	Description	Restrictions	Syntax
<i>column</i>	Column on which to fragment the index	Must exist in the current table	"Identifier " on page 5-21
<i>dbslice</i>	Dbslice storing all the index fragments	Must exist	"Identifier " on page 5-21
<i>dbspace</i>	Dbpace in which to store the index fragment that <i>expr</i> defines	Must exist. Include no more than 2,048 <i>dbspace</i> names	"Identifier " on page 5-21
<i>expr</i>	Expression defining which index keys each fragment stores	See "Restrictions on Fragmentation Expressions."	"Expression" on page 4-40; "Condition " on page 4-5
<i>part</i>	Name that you declare here for a partition of a <i>dbspace</i>	Required for any partition in the same <i>dbspace</i> as another partition of the same index. Name must be unique among partitions of the same table.	"Identifier " on page 5-21

Here the IN keyword introduces the name of a storage space where an index fragment is to be stored. If you list multiple *dbspace* names after the IN keyword, use parentheses to delimit the *dbspace* list. All *dbspaces* that store the fragments must have the same page size. The parentheses around the list of fragment definitions that follow the EXPRESSION keyword are optional.

Restrictions on Fragmentation Expressions

The following restrictions apply to the expression:

- Each fragment expression can contain columns only from the current table, with data values only from a single row.
- The columns contained in a fragment expression must be the same as the indexed columns or a subset of the indexed columns.
- The expression must return a BOOLEAN (true or false) value.
- No subqueries, aggregates, user-defined routines, nor references to fields of a ROW type column are valid.

- The built-in CURRENT, DATE, SYSDATE, and TODAY functions are not valid.

In Extended Parallel Server, you can fragment indexes on any column of a table, even if the table spans multiple coservers. The columns that you specify in the FRAGMENT BY clause do not need to be part of the index key.

Detached indexes of Extended Parallel Server can be either locally detached or globally detached. A *locally detached index* is an index in which, for each data tuple in a table, the corresponding index tuple is guaranteed to be on the same coserver. The table and index fragmentation strategies need not be identical if co-locality can be guaranteed. If the data tuple and index tuple co-locality do not exist, then the index is a *globally-detached index*. For performance implications of globally-detached indexes, see your *IBM Informix Performance Guide*. For more information on options for distribution strategies in Extended Parallel Server, see “Fragmenting by EXPRESSION” on page 2-214, “Fragmenting by HASH (XPS)” on page 2-216, and “Fragmenting by HYBRID (XPS)” on page 2-217, respectively, in the description of the CREATE TABLE statement.

Fragmentation of System Indexes

System indexes (such as those that implement referential constraints and unique constraints) utilize user-defined indexes if they exist. If no user-defined indexes can be utilized, system indexes remain nonfragmented, and are moved to the dbspace where the database was created.

To fragment a system index, create the fragmented index on the constraint columns, and then add the constraint using the ALTER TABLE statement.

Fragmentation of Unique Indexes (IDS)

You can fragment unique indexes on a table that uses a round-robin or an expression-based distribution scheme, but any columns referenced in the fragment expression must be indexed columns. If your index fragmentation strategy violates this restriction, the CREATE INDEX statement fails, and work is rolled back.

Fragmentation of Indexes on Temporary Tables

You can fragment a unique index on a temporary table only if the underlying table uses an expression-based distribution scheme. That is, the CREATE Temporary TABLE statement that defines the temporary table must specify an explicit expression-based distribution scheme.

If you try to create a fragmented, unique index on a temporary table for which you did not specify a fragmentation strategy when you created the table, the database server creates the index in the first dbspace that the **DBSPACETEMP** environment variable specifies. For more information on the **DBSPACETEMP** environment variable, see the *IBM Informix Guide to SQL: Reference*.

For more information on the default storage characteristics of temporary tables, see “Where Temporary Tables are Stored” on page 2-237.

Index Modes (IDS)

Use the index mode options to specify the behavior of the index during insert, delete, and update operations.

Index Modes:



The following table explains the index modes

Mode	Effect
------	--------

DISABLED

The database server does not update the index after insert, delete, and update operations that modify the base table. The optimizer does not use the index during the execution of queries.

ENABLED

The database server updates the index after insert, delete, and update operations that modify the base table. The optimizer uses the index during query execution. If an insert or update operation causes a duplicate key value to be added to a unique index, the statement fails.

FILTERING

The database server updates a unique index after insert, delete, and update operations that modify the base table. (This option is not available with duplicate indexes.)

The optimizer uses the index during query execution. If an insert or update operation causes a duplicate key value to be added to a unique index in filtering mode, the statement continues processing, but the bad row is written to the violations table associated with the base table. Diagnostic information about the unique-index violation is written to the diagnostics table associated with the base table.

If you specify filtering for a unique index, you can also specify one of the following error options.

Error Option	Effect
--------------	--------

WITHOUT ERROR

A unique-index violation during an insert or update operation returns no integrity-violation error to the user.

WITH ERROR

Any unique-index violation during an insert or update operation returns an integrity-violation error to the user.

Specifying Modes for Unique Indexes

You must observe the following guidelines when you specify modes for unique indexes in CREATE INDEX statements:

- You can set the mode of a unique index to enabled, disabled, or filtering.
- If you do not specify a mode, then by default the index is enabled.
- For an index set to filtering mode, if you do not specify an error option, the default is WITHOUT ERROR.
- When you add a new unique index to an existing base table and specify the disabled mode for the index, your CREATE INDEX statement succeeds even if duplicate values in the indexed column would cause a unique-index violation.

- When you add a new unique index to an existing base table and specify the enabled or filtering mode for the index, your CREATE INDEX statement succeeds provided that no duplicate values exist in the indexed column that would cause a unique-index violation. However, if any duplicate values exist in the indexed column, your CREATE INDEX statement fails and returns an error.
- When you add a new unique index to an existing base table in the enabled or filtering mode, and duplicate values exist in the indexed column, erroneous rows in the base table are not filtered to the violations table. Thus, you cannot use a violations table to detect the erroneous rows in the base table.

Adding a Unique Index When Duplicate Values Exist in the Column

If you attempt to add a unique index in the enabled mode but receive an error message because duplicate values are in the indexed column, take the following steps to add the index successfully:

1. Add the index in the disabled mode. Issue the CREATE INDEX statement again, but this time specify the DISABLED keyword.
2. Start a violations and diagnostics table for the target table with the START VIOLATIONS TABLE statement.
3. Issue a SET Database Object Mode statement to change the mode of the index to enabled. When you issue this statement, existing rows in the target table that violate the unique-index requirement are duplicated in the violations table. You receive an integrity-violation error message, however, and the index remains disabled.
4. Issue a SELECT statement on the violations table to retrieve the nonconforming rows that are duplicated from the target table. You might need to join the violations and diagnostics tables to get all the necessary information.
5. Take corrective action on the rows in the target table that violate the unique-index requirement.
6. After you fix all the nonconforming rows in the target table, issue the SET Database Object Mode statement again to switch the disabled index to the enabled mode. This time the index is enabled, and no integrity violation error message is returned because all rows in the target table now satisfy the new unique-index requirement.

Specifying Modes for Duplicate Indexes

You must observe the following guidelines when you specify modes for duplicate indexes in CREATE INDEX statements:

- You can set a duplicate index to enabled or disabled mode. Filtering mode is available only for unique indexes.
- If you do not specify the mode of a duplicate index, by default the index is enabled.

How the Database Server Treats Disabled Indexes (IDS)

Whether a disabled index is a unique or duplicate index, the database server effectively ignores the index during data-manipulation (DML) operations.

When an index is disabled, the database server stops updating it and stops using it during queries, but the catalog information about the disabled index is retained. You cannot create a new index on a column or set of columns if a disabled index on that column or set of columns already exists. Similarly, you cannot create an active (enabled) unique, foreign-key, or primary-key constraint on a column or on a set of columns if the indexes on which the active constraint depends are disabled.

LOCK MODE Options (XPS)

The LOCK MODE options specify the locking granularity of the index.

LOCK MODE Options:



In COARSE lock mode, index-level locks are acquired on the index instead of item-level or page-level locks. This mode reduces the number of lock calls on the index. Use the coarse-lock mode when you know the index is not going to change, as when only read-only operations are performed on the index.

If you specify no lock mode, the default is NORMAL. That is, the database server places item-level or page-level locks on the index as necessary.

Generalized-Key Indexes (XPS)

If you are using Extended Parallel Server, you can create *generalized-key* (GK) indexes. Keys in a conventional index consist of one or more columns of the STATIC table that is being indexed. A GK index stores information about the records in a STATIC table based on the results of a query.

GK indexes provide a form of pre-computed index capability that supports faster query processing, especially in data-warehousing environments. The optimizer can use the GK index to improve performance.

A GK index is *defined on* a table when that table is the one being indexed. A GK index *depends on* a table when the table appears in the FROM clause of the index. Before you create a GK index, keep the following issues in mind:

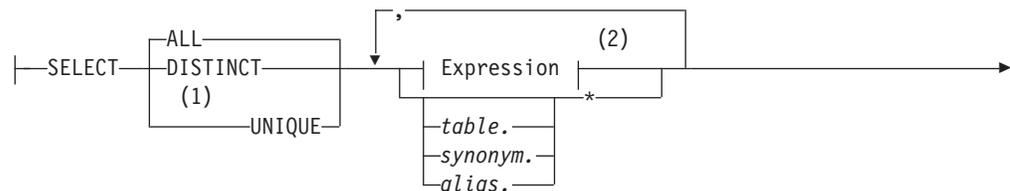
- All tables used in a GK index must be STATIC tables. If you try to change the type of a table to a non-static type while a GK index depends on that table, the database server returns an error.
- Any table involved in a GK index must be a STATIC type. UPDATE, DELETE, INSERT, and LOAD operations are not valid on such a table until you drop the dependent GK index and the table type changes.

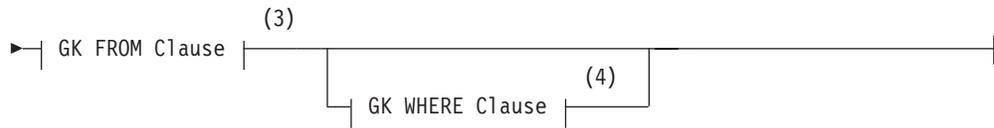
Key-only index scans are not available with GK indexes.

SELECT Clause for Generalized-Key Index

In Extended Parallel Server, the options of the GK SELECT clause are a subset of the options of "SELECT" on page 2-528. The GK SELECT clause has this syntax:

GK SELECT Clause:





Notes:

- 1 Informix extension
- 2 See "Expression" on page 4-40
- 3 See " FROM Clause for Generalized-Key Index"
- 4 See " WHERE Clause for Generalized-Key Index" on page 2-143

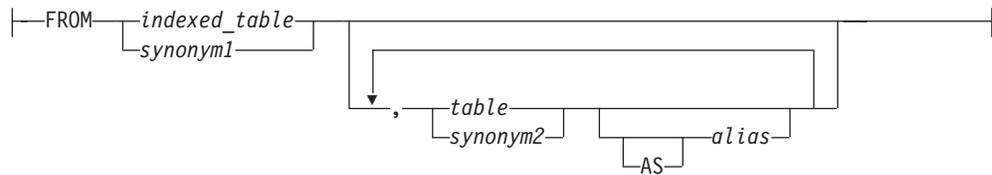
Element	Description	Restrictions	Syntax
<i>alias</i>	Temporary name assigned to the table in the FROM clause	You cannot use an alias for the table on which the index is built	"Identifier " on page 5-21
<i>synonym, table</i>	Synonym or table from which to retrieve data	The synonym and the table to which it points must exist	"Database Object Name " on page 5-16

The following restrictions apply to expressions in the GK SELECT clause:

- It cannot refer to any SPL routine.
- It cannot include the USER, TODAY, CURRENT, or DBINFO built-in functions, nor any function that refers to a point in time or interval.

FROM Clause for Generalized-Key Index

GK FROM Clause:



Element	Description	Restrictions	Syntax
<i>alias</i>	Temporary name for a table	You cannot use an alias with <i>indexed_table</i>	"Identifier " on page 5-21
<i>indexed_table, synonym1</i>	Table on which the index is being built	The FROM clause must include the indexed table	"Database Object Name " on page 5-16
<i>synonym2, table</i>	Synonym or identifier of table from which to retrieve data	The synonym and the table to which it points must exist	"Database Object Name " on page 5-16

All tables that appear in the FROM clause must be local static tables. That is views, non-static tables, and remote tables are not valid in the FROM clause.

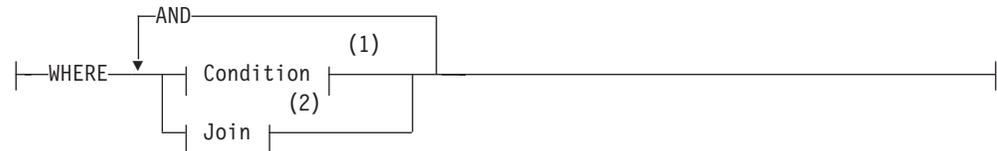
Tables that you specify in the FROM clause must be *transitively joined on key* to the indexed table. Table **A** is transitively joined on key to table **B** if **A** and **B** are joined with equal joins on the unique-key columns of **A**.

Suppose that tables **A**, **B**, and **C** each have **col1** as a primary key. In the following example, **B** is joined on key to **A** and **C** is joined on key to **B**. **C** is transitively joined on key to **A**.

```
CREATE GK INDEX gki
  (SELECT A.col1, A.col2 FROM A, B, C
   WHERE A.col1 = B.col1 AND B.col1 = C.col1);
```

WHERE Clause for Generalized-Key Index

GK WHERE Clause:



Notes:

- 1 See "Condition " on page 4-5
- 2 See "Specifying a Join in the WHERE Clause" on page 2-564

The WHERE clause for a GK index has the following restrictions:

- It cannot include USER, TODAY, CURRENT, or DBINFO built-in functions, nor any function that references a time value or a time-interval value.
- It cannot refer to any SPL routine.
- It cannot include any subquery.
- It cannot include any aggregate function.
- It cannot include any IN, LIKE, or MATCHES expression.

The ONLINE Keyword (IDS)

By default, CREATE INDEX attempts to place an exclusive lock on the indexed table to prevent all other users from accessing the table while the index is being created. The CREATE INDEX statement fails if another user already has a lock on the table, or is currently accessing the table at the Dirty Read isolation level.

The DBA can reduce the risk of non-exclusive access errors, and can increase the availability of the indexed table, by including the ONLINE keyword as the last specification of the CREATE INDEX statement. This instructs the database server to create the index online while concurrent users can continue to access the table.

The database server builds the index, even if other users are performing Dirty Read and DML operations on the indexed table. Immediately after you issue the CREATE INDEX . . . ONLINE statement, the new index is not visible to the query optimizer for use in query plans or cost estimates, and the database server does not support any other DDL operations on the indexed table, until after the specified index has been built without errors. At this time, the database server briefly locks the table while updating the system catalog with information about the new index.

The indexed table in a CREATE INDEX . . . ONLINE statement can be permanent or temporary, logged or unlogged, and fragmented or non-fragmented. You cannot specify the ONLINE keyword, however, when you create a functional index, a clustered index, a virtual index, or an R-tree index.

The following statement instructs the database server to create a unique online index called **idx_1** on the **lname** column of the **customer** table:

```
CREATE UNIQUE INDEX idx_1 ON customer(lname) ONLINE;
```

If, while this index is being constructed, other users insert into the **customer** table new rows in which **lname** is not unique, the database server issues an error after it has created the new **idx_1** index and registered it in the system catalog.

The term *online index* refers to the locking strategy that the database follows in creating or dropping an index with the **ONLINE** keyword, rather than to properties of the index that persist after its creation (or its destruction) has completed. This term appears in some error messages, however, and in recovery or restore operations, the database server re-creates as an online index any index that you created as an online index.

No more than one **CREATE INDEX . . . ONLINE** or **DROP INDEX . . . ONLINE** statement can concurrently reference online indexes on the same table, or online indexes that have the same identifier.

Automatic Calculation of Distribution Statistics

When the **CREATE INDEX** statement runs successfully, with or without the **ONLINE** keyword, Dynamic Server automatically gathers statistics for the newly created index, and updates the **sysdistrib** system catalog table with values that are equivalent to an **UPDATE STATISTICS** operation in a mode that depends on the type of index:

- Index level statistics, equivalent to the statistics gathered by **UPDATE STATISTICS** in the **LOW** mode, are calculated for most types of indexes, including B-tree, Virtual Index Interface, and functional indexes.
- Column distribution statistics, equivalent to the distribution generated in the **HIGH** mode, for a non-opaque leading indexed column of an ordinary B-tree index. The resolution percentage is 1.0 if the table has fewer than a million rows, and 0.5 for larger table sizes.

These distribution statistics are available to the query optimizer when it designs query plans for the table on which the new index was created.

For composite key indexes, only distributions of the leading column are created implicitly by the **CREATE INDEX** statement.

The implicit creation of distribution statistics is not supported for the following types of indexes:

- Indexes on columns of user-defined data types
- Indexes on columns of the built-in opaque data types (including **BOOLEAN** and **LVARCHAR**)
- R-tree indexes
- Attached indexes.

If the calculation of distribution statistics fails during the **CREATE INDEX** operation, the database server reports that failure in the error log, but continues to create the index.

When distributions are successfully created by an explicit or implicit CREATE INDEX operation, explain information (similar to one generated by UPDATE STATISTICS) such as following is generated if the SET EXPLAIN facility is set to ON.

```

Index:          idx_01 on nita.foo
STATISTICS CREATED AUTOMATICALLY:
Column Distribution for:          nita.foo.a
Mode:          MEDIUM
Number of Bins:    101 Bin size:  100.0
Sort data:        0.3 MB
Completed building distribution in:  0 minutes 33 seconds

```

See the description of the “UPDATE STATISTICS ” on page 2-705 statement for information about distribution statistics and about the difference between LOW mode and MEDIUM mode distributions.

Related Information

Related statements: “ALTER INDEX” on page 2-30, “CREATE OPCLASS” on page 2-149, “CREATE TABLE ” on page 2-191, “DROP INDEX ” on page 2-329, “RENAME INDEX ” on page 2-489, and “SET Database Object Mode” on page 2-591

For a discussion of the structure of indexes, see your *IBM Informix Administrator’s Reference*.

For a discussion of the different types of indexes and information about performance issues with indexes, see your *IBM Informix Performance Guide*.

For a discussion of the GLS aspects of the CREATE INDEX statement, see the *IBM Informix GLS User’s Guide*.

For information about operator classes, refer to the “CREATE OPCLASS” on page 2-149 statement and *IBM Informix User-Defined Routines and Data Types Developer’s Guide*.

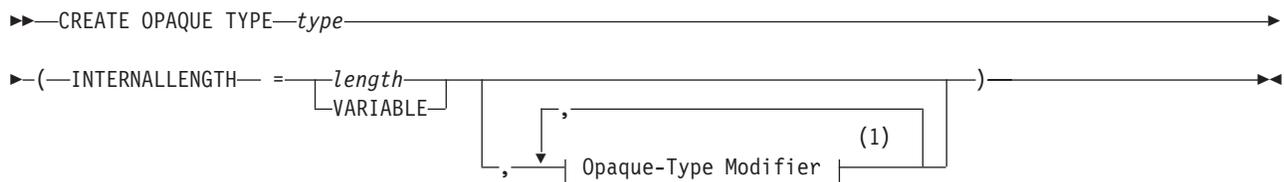
For information about the indexes that DataBlade modules provide, refer to your DataBlade module documentation.

CREATE OPAQUE TYPE

Use the CREATE OPAQUE TYPE statement to create an opaque data type.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

1 See “ Opaque-Type Modifier” on page 2-147

Element	Description	Restrictions	Syntax
<i>length</i>	Number of bytes needed to store a value of this data type	Positive integer returned when <code>sizeof()</code> directive is applied to the type structure	“Literal Number ” on page 4-185
<i>type</i>	Name that you declare here for the new opaque data type	Must be unique among data type names in the database	“Identifier ” on page 5-21

Usage

The CREATE OPAQUE TYPE statement registers a new opaque data type in the `sysxdtypes` system catalog table.

To create an opaque type, you must have the Resource privilege on the database. When you create the opaque type, only you, the owner, have the Usage privilege on the new opaque data type. You can use the GRANT or REVOKE statements to grant or revoke the Usage privilege of other users of the database.

To view the privileges on a data type, check the `sysxdtypes` system catalog table for the owner name, and check the `sysxdtypeauth` system catalog table for additional type privileges that might have been granted.

For details of system catalog tables, see the *IBM Informix Guide to SQL: Reference*.

The DB-Access utility can also display privileges on opaque data types.

Declaring a Name for an Opaque Type

The name that you declare for an opaque data type is an SQL identifier. When you create an opaque type in a database that is not ANSI-compliant, the name must be unique among the names of data types within the database.

When you create an opaque type in an ANSI-compliant database, *owner.type* combination must be unique within the database. The owner name is case sensitive. If you do not put quotation marks around the owner name, the name of the opaque-type *owner* is stored in uppercase letters.

INTERNALLENGTH Modifier

The INTERNALLENGTH modifier specifies the storage size that is required for the opaque data type as fixed length or varying length.

Fixed-Length Opaque Types

A fixed-length opaque type has an internal structure of fixed size. To create a fixed-length opaque type, specify the size of the internal structure, in bytes, for the INTERNALLENGTH modifier. The next example creates a fixed-length opaque type called `fixlen_typ` and allocates 8 bytes for storing this data type.

```
CREATE OPAQUE TYPE fixlen_typ(INTERNALLENGTH=8, CANNOTHASH)
```

Varying-Length Opaque Types

A varying-length opaque type has an internal structure whose size might vary from one value to another. For example, the internal structure of an opaque data type might hold the actual value of a string up to a certain size, but beyond this size it might use an LO-pointer to a CLOB to hold the value.

To create a varying-length opaque data type, use the VARIABLE keyword with the INTERNALLENGTH modifier. The following statement creates a variable-length opaque data type called **varlen_typ**:

```
CREATE OPAQUE TYPE varlen_typ
(INTERNALLENGTH=VARIABLE, MAXLEN=1024)
```

Opaque-Type Modifier

Opaque-Type Modifier:



Element	Description	Restrictions	Syntax
<i>align_value</i>	Byte boundary on which to align an opaque type that is passed to a user-defined routine. Default is 4 bytes.	Must be 1, 2, 4, or 8, depending on the C definition of the opaque data type and hardware and compiler used to build the object file for the data type	“Literal Number ” on page 4-185
<i>length</i>	Maximum length to allocate for instances of varying-length opaque types. Default is 2 kilobytes.	Must be a positive integer ≤ 32 kilobytes. Do not specify for fixed-length data types. Values that exceed this length return errors.	“Literal Number ” on page 4-185

Modifiers can specify the following optional information for opaque types:

- MAXLEN specifies the maximum length for varying-length types.
- CANNOTHASH specifies that the database server cannot use the built-in hash function on the opaque type.
- ALIGNMENT specifies the byte boundary on which the database server aligns the opaque type.
- PASSEDBYVALUE specifies that an opaque type that requires 4 bytes or fewer of storage is passed by value.

By default, opaque types are passed to user-defined routines by reference.

Defining an Opaque Type

To define a new opaque data type to the database server, you must provide the following information in the C or Java language:

- A data structure that serves as the internal storage of the opaque data type
The internal storage details of the type are hidden, or opaque. Once you define a new opaque data type, the database server can manipulate it without knowledge of the C or Java structure in which it is stored.
- Support functions that allow the database server to interact with this internal structure.
The support functions tell the database server how to interact with the internal structure of the data type. These support functions must be written in the C or Java programming language.
- Additional user-defined functions that other support functions or end users can invoke to operate on the opaque type (optional)

Possible support functions include operator functions and cast functions. Before you can use these functions in SQL statements, they must be registered with the appropriate CREATE CAST, CREATE PROCEDURE, or CREATE FUNCTION statement.

The following table summarizes the support functions for an opaque data type.

Function	Description	Invoked
input()	Converts the opaque type from its external LVARCHAR representation to its internal representation	When a client application sends a character representation of the opaque type in an INSERT, UPDATE, or LOAD statement
output()	Converts the opaque type from its internal representation to its external LVARCHAR representation	When the database server sends a character representation of the opaque type as a result of a SELECT or FETCH statement
receive()	Converts the opaque type from its internal representation on the client computer to its internal representation on the server computer Provides platform-independent results regardless of differences between client and server computer types	When a client application sends an internal representation of the opaque type in an INSERT, UPDATE, or LOAD statement
send()	Converts the opaque type from its internal representation on the server computer to its internal representation on the client computer Provides platform-independent results regardless of differences between client and database server computer types	When the database server sends an internal representation of the opaque type as a result of a SELECT or FETCH statement
db_receive()	Converts the opaque type from its internal representation on the local database to the dbsendrcv type for transfer to an external database on the local server	When a local database receives a dbsendrcv type from an external database on the local database server
db_send()	Converts the opaque type from its internal representation on the local database to the dbsendrcv type for transfer to an external database on the local server	When a local database sends a dbsendrcv type to an external database on the local database server
server_receive()	Converts the opaque type from its internal representation on the local server computer to the srvsendrcv type for transfer to a remote database server Use any name for this function.	When the local database server receives a srvsendrcv type from a remote database server
server_send()	Converts the opaque type from its internal representation on the local server computer to the srvsendrcv type for transfer to a remote database server Use any name for this function.	When the local database server sends a srvsendrcv type to a remote database server
import()	Performs any tasks needed to convert from the external (character) representation of an opaque type to the internal format for a bulk copy	When DB–Access (LOAD) or the High Performance Loader (HPL) initiates a bulk copy from a text file to a database
export()	Performs any tasks needed to convert from the internal representation of an opaque type to the external (character) format for a bulk copy	When DB–Access (UNLOAD) or the High Performance Loader initiates a bulk copy from a database to a text file
importbinary()	Performs any tasks needed to convert from the internal representation of an opaque type on the client computer to the internal representation on the server computer for a bulk copy	When DB–Access (LOAD) or the High Performance Loader initiates a bulk copy from a binary file to a database

Function	Description	Invoked
<code>exportbinary()</code>	Performs any tasks needed to convert from the internal representation of an opaque type on the server computer to the internal representation on the client computer for a bulk copy	When DB–Access (UNLOAD) or the High Performance Loader initiates a bulk copy from a database to a binary file
<code>assign()</code>	Performs any processing required before storing the opaque type to disk. This support function must be named <code>assign()</code> .	When the database server executes INSERT, UPDATE, or LOAD, before it stores the opaque type to disk
<code>destroy()</code>	Performs any processing necessary before removing a row that contains the opaque type. This support function must be named <code>destroy()</code> .	When the database server executes the DELETE or DROP TABLE, before it removes the opaque type from disk
<code>lohandles()</code>	Returns a list of the LO-pointer structures (pointers to smart large objects) in an opaque type	When the database server must search opaque types for references to smart large objects; when <code>oncheck</code> runs, or an archive is performed
<code>compare()</code>	Compares two values of the opaque type and returns an integer value to indicate whether the first value is less than, equal to, or greater than the second value	When the database server encounters an ORDER BY, UNIQUE, DISTINCT, or UNION clause in a SELECT statement, or when CREATE INDEX creates a B-tree index

After you write the necessary support functions for the opaque type, use the CREATE FUNCTION statement to register these support functions in the same database as the opaque type. Certain support functions convert other data types to or from the new opaque type. After you create and register these support functions, use the CREATE CAST statement to associate each function with a particular cast. The cast must be registered in the same database as the support function.

After you have written the necessary C language or Java language source code to define an opaque data type, you then use the CREATE OPAQUE TYPE statement to register the opaque data type in the database.

Related Information

Related statements: “CREATE CAST ” on page 2-95, “CREATE DISTINCT TYPE ” on page 2-100, “CREATE FUNCTION ” on page 2-113, “CREATE ROW TYPE ” on page 2-168, “CREATE TABLE ” on page 2-191, and “DROP TYPE ” on page 2-345

For a description of an opaque type, see the *IBM Informix Guide to SQL: Reference*.

For information on how to define an opaque type, see *IBM Informix User-Defined Routines and Data Types Developer’s Guide*.

For information on how to use the Java language to define an opaque type, see the *J/Foundation Developer’s Guide*.

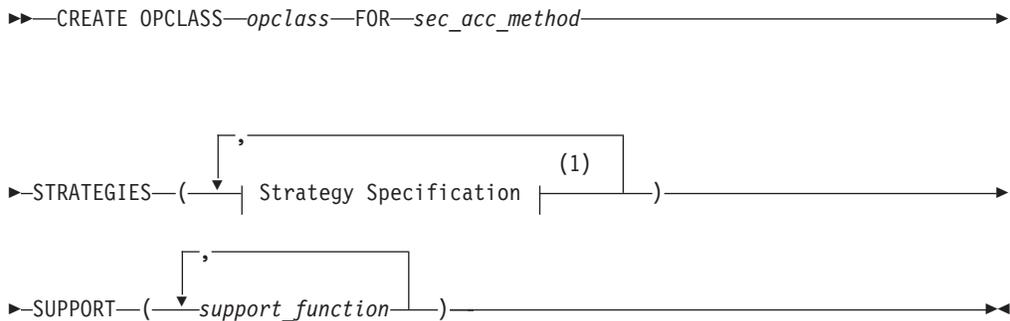
For information about the GLS aspects of the CREATE OPAQUE TYPE statement, refer to the *IBM Informix GLS User’s Guide*.

CREATE OPCLASS

Use the CREATE OPCLASS statement to create an *operator class* for a *secondary-access method*.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 See "CREATE OPCLASS" on page 2-149

Element	Description	Restrictions	Syntax
<i>opclass</i>	Name that you declare here for a new operator class	Must be unique among operator classes within the database	"Identifier " on page 5-21
<i>sec_acc_method</i>	Secondary-access method with which the new operator class is associated	Must already exist and must be registered in the sysams table	"Identifier " on page 5-21
<i>support_function</i>	Support function that the secondary-access method requires	Must be listed in the order expected by the access method	"Identifier " on page 5-21

Usage

An *operator class* is the set of operators that support a secondary-access method for query optimization and building the index. A *secondary-access method* (sometimes referred to as an *index access method*) is a set of database server functions that build, access, and manipulate an index structure such as a B-tree, R-tree, or an index structure that a DataBlade module provides.

The database server provides the B-tree and R-tree secondary-access methods. For more information on the **btree** secondary-access method, see "Default Operator Classes" on page 2-153.

Define a new operator class when you want one of the following:

- An index to use a different order for the data than the sequence that the default operator class provides
- A set of operators that is different from any existing operator classes that are associated with a particular secondary-access method

You must have the Resource privilege or be the DBA to create an operator class. The actual name of an operator class is an SQL identifier. When you create an operator class, the *opclass* name must be unique within the database.

When you create an operator class in an ANSI-compliant database, the *owner.opclass* combination must be unique within the database. The owner name is case sensitive. If you do not put quotation marks around the *owner* name (or else set the **ANSIOWNER** environment variable), the name of the operator-class owner is stored in uppercase letters.

The following CREATE OPCLASS statement creates a new operator class called **abs_btree_ops** for the **btree** secondary-access method:

```
CREATE OPCLASS abs_btree_ops FOR btree
  STRATEGIES (abs_lt, abs_lte, abs_eq, abs_gte, abs_gt)
  SUPPORT (abs_cmp);
```

An operator class has two kinds of operator-class functions:

- Strategy functions
Specify strategy functions of an operator class in the STRATEGY clause of the CREATE OPCLASS statement. In the preceding CREATE OPCLASS code example, the **abs_btree_ops** operator class has five strategy functions.
- Support functions
Specify support functions of an operator class in the SUPPORT clause. In the preceding CREATE OPCLASS code example, the **abs_btree_ops** operator class has one support function.

STRATEGIES Clause

Strategy functions are functions that users can invoke within a DML statement to operate on a specific data type. The query optimizer uses the strategy functions to determine whether a given index can be used to process a query.

If a query includes a UDF or a column on which an index exists, and if the qualifying operator in the query matches any function in the STRATEGIES clause, then the query optimizer considers using the index for the query. For more information on query plans, see your *IBM Informix Performance Guide*.

When you create a new operator class, the STRATEGIES clause identifies the strategy functions for the secondary-access method. Each strategy specification lists the name of a strategy function (and optionally, the data types of its parameters). You must list these functions in the order that the secondary-access method expects. For the specific order of strategy operators for the default operator classes for a B-tree index and for an R-tree index, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

Strategy Specification

The STRATEGIES keyword introduces a comma-separated list of function names or function signatures for the new operator class. Each element of this list is called a *strategy specification* and has the following syntax:

Strategy Specification:

```
|—strategy_function—|
| (—input_type—,—input_type—|
| |,—output_type—| ) |
```

Element	Description	Restrictions	Syntax
<i>input_type</i>	Data type of an input parameter to the strategy function for which you intend to use a specific secondary-access method	A <i>strategy function</i> accepts two input parameter and can have one optional output parameter	"Data Type " on page 4-21
<i>output_type</i>	Data type of the optional output parameter of the strategy function	Optional output parameter for side-effect indexes	"Data Type " on page 4-21
<i>strategy_function</i>	Strategy function to associate with the specified operator class	Must be listed in the order that the specified secondary-access method expects	"Identifier " on page 5-21

Each *strategy_function* is an external function. The CREATE OPCLASS statement does not verify that a user-defined function of the name you specify exists. However, for the secondary-access method to use the strategy function, the external function must be:

- Compiled in a shared library
- Registered in the database with the CREATE FUNCTION statement

Optionally, you can specify the signature of a strategy function in addition to its name. A strategy function requires two input parameters and an optional output parameter. To specify the function signature, specify:

- An *input data type* for each of the two input parameters of the strategy function, in the order that the strategy function uses them
- Optionally, one *output data type* for an output parameter of the strategy function

You can specify UDTs as well as built-in data types. If you do not specify the function signature, the database server assumes that each strategy function takes two arguments of the same data type and returns a BOOLEAN value.

Indexes on Side-Effect Data

Side-effect data are additional values that a strategy function returns after a query that contains the strategy function. For example, an image DataBlade module might use a *fuzzy* index to search image data. The index ranks the images according to how closely they match the search criteria. The database server returns the rank values as side-effect data with the qualifying images.

SUPPORT Clause

Support functions are functions that the secondary-access method uses internally to build and search the index. Specify these functions for the secondary-access method in the SUPPORT clause of the CREATE OPCLASS statement.

You must list the names of the support functions in the order that the secondary-access method expects. For the specific order of support operators for the default operator classes for a B-tree index and an R-tree index, refer to "Default Operator Classes" on page 2-153.

The support function is an external function. CREATE OPCLASS does not verify that a specified support function exists. For the secondary-access method to use a support function, however, the support function must meet these criteria:

- Be compiled in a shared library
- Be registered in the database with the CREATE FUNCTION statement

Default Operator Classes

Each secondary-access method has a default operator class that is associated with it. By default, the CREATE INDEX statement associates the default operator class with an index. For example, the following CREATE INDEX statement creates a B-tree index on the **zipcode** column and automatically associates the default B-tree operator class with this column:

```
CREATE INDEX zip_ix ON customer(zipcode)
```

For each of the secondary-access methods that Dynamic Server provides, it provides a *default operator class*, as follows:

- The default B-tree operator class is a built-in operator class.
The database server implements the operator-class functions for this operator class and registers it as **btree_ops** in the system catalog tables of a database.
- The default R-tree operator class is a registered operator class.
The database server registers this operator class as **rtree_ops** in the system catalog tables. The database server does *not* implement the operator-class functions for the default R-tree operator class.

Important: To use an R-tree index, you must install a spatial DataBlade module such as the Geodetic DataBlade module or any other third-party DataBlade module that implements the R-tree index. These implement the R-tree operator-class functions.

DataBlade modules can provide other types of secondary-access methods. If a DataBlade module provides a secondary-access method, it might also provide a default operator class. For more information, refer to your DataBlade module user's guide.

Related Information

Related statements: "CREATE FUNCTION " on page 2-113, "CREATE INDEX " on page 2-124, and "DROP OPCLASS " on page 2-331

For information on support functions and how to create and extend an operator class, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

For more about R-tree indexes, see the *IBM Informix R-Tree Index User's Guide*.

For information about the GLS aspects of the CREATE OPCLASS statement, refer to the *IBM Informix GLS User's Guide*.

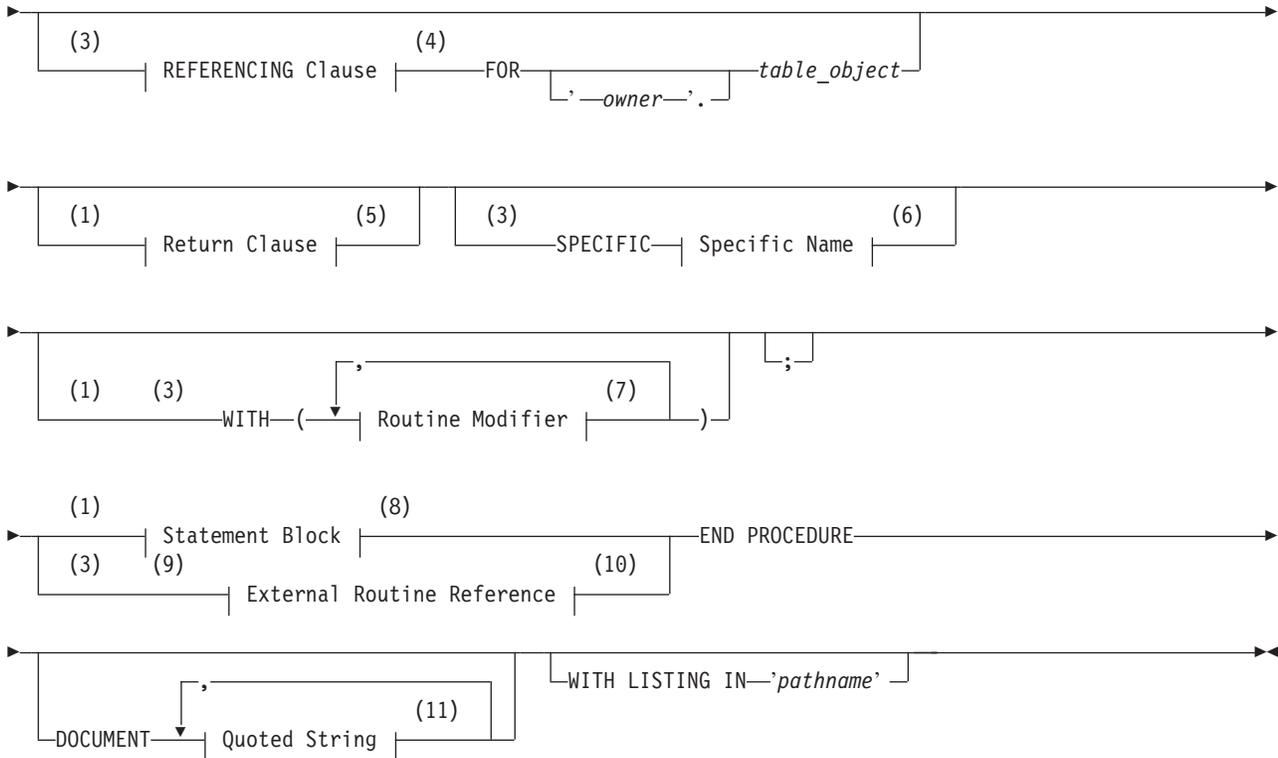
CREATE PROCEDURE

Use the CREATE PROCEDURE statement to create a user-defined procedure. (To create a procedure from text of source code that is in a separate file, use the CREATE PROCEDURE FROM statement.)

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax

```
►► CREATE DBA PROCEDURE procedure  
(1)  
function ( Routine Parameter List  
(2) )
```



Notes:

- 1 Stored Procedure Language only
- 2 See "Routine Parameter List" on page 5-67
- 3 Dynamic Server only
- 4 See "The REFERENCING and FOR Clauses" on page 2-157
- 5 See "Return Clause " on page 5-55
- 6 See "Specific Name " on page 5-73
- 7 See "Routine Modifier" on page 5-59
- 8 See "Statement Block " on page 5-74
- 9 External routines only
- 10 See " External Routine Reference" on page 5-19
- 11 See "Quoted String" on page 4-189

Element	Description	Restrictions	Syntax
<i>function, procedure</i>	Name declared here for a new SPL routine	(XPS) The name must be unique among all SPL routines in the database. (IDS) See "Procedure Names in Dynamic Server" on page 2-159.	"Identifier " on page 5-21
<i>owner</i>	Owner of <i>table_object</i>	Must own <i>table_object</i>	" Owner Name " on page 5-45
<i>pathname</i>	File to store compile-time warnings	Must exist on the computer where the database resides	Operating system specific

Element	Description	Restrictions	Syntax
<i>table_object</i>	Name or synonym of a table or view whose triggers can call this UDR	Must exist in the local database	"Identifier " on page 5-21

Usage

The entire length of a CREATE PROCEDURE statement must be less than 64 kilobytes. This length is the literal length of the CREATE PROCEDURE statement text, including blank spaces, tabs, and other white space characters.

In Informix ESQL/C, you can use CREATE PROCEDURE only as text within a PREPARE statement. If you want to create a procedure for which the text is known at compile time, you must use a CREATE PROCEDURE FROM statement.

Routines use the collating order that was in effect when they were created. See "SET COLLATION" on page 2-584 statement of Dynamic Server for information about using non-default collation.

Using CREATE PROCEDURE Versus CREATE FUNCTION

In Extended Parallel Server, besides using this statement to create SPL procedures, you must use CREATE PROCEDURE to write and register an SPL routine that returns one or more values (that is, an SPL function). Extended Parallel Server does not support the CREATE FUNCTION statement.

In Dynamic Server, although you can use CREATE PROCEDURE to write and register an SPL routine that returns one or more values (that is, an SPL function), it is recommended that you use CREATE FUNCTION instead. To register an external function, you must use CREATE FUNCTION.

Use the CREATE PROCEDURE statement to write and register an SPL procedure or to register an external procedure.

For information on how terms such as user-defined procedures and user-defined functions are used in this document, see "Relationship Between Routines, Functions, and Procedures."

Relationship Between Routines, Functions, and Procedures

A *procedure* is a routine that can accept arguments but does not return any values. A *function* is a routine that can accept arguments and returns one or more values. *User-defined routine* (UDR) is a generic term that includes both user-defined procedures and user-defined functions. For information about named and unnamed returned values, see "Return Clause " on page 5-55.

You can write a UDR in SPL (a *SPL routine*) or in an external language (an *external routine*) that the database server supports. Where the term UDR appears in this document, it can refer to both SPL routines and external routines.

The term *user-defined procedure* refers to SPL procedures and external procedures. *User-defined function* refers to SPL functions and external functions.

In the documentation of earlier releases, the term *stored procedure* was used for both SPL procedures and SPL functions. In this document, the term *SPL routine* replaces

the term stored procedure. When it is necessary to distinguish between an SPL function and an SPL procedure, this document does so.

The term *external routine* applies to an external procedure or (for Dynamic Server) an external function, both constructs designating UDRs that are written in a programming language other than SPL. When it is necessary to distinguish between an external function and an external procedure, this document does so.

Extended Parallel Server does not support external routines, but the term user-defined routine (UDR) encompasses both *SPL routines* and *external routines*. Wherever the term UDR appears, it is applicable to SPL routines.

Privileges Necessary for Using CREATE PROCEDURE

You must have the Resource privilege on a database to create a user-defined procedure within that database.

In Dynamic Server, before you can create an SPL procedure, you must also have the Usage privilege on the SPL language. For more information, see “Usage Privilege in Stored Procedure Language” on page 2-407.

By default, the Usage privilege on SPL is granted to PUBLIC. You must also have at least the Resource privilege on a database to create an SPL procedure within that database.

DBA Keyword and Privileges on the Procedure

If you create a UDR with the DBA keyword, it is known as a *DBA-privileged UDR*. You need the DBA privilege to create a DBA-privileged UDR.

Among users who do not hold the DBA privilege, only those to whom the DBA grants the Execute privilege can invoke the DBA-privileged UDR. If the DBA grants the Execute privilege to PUBLIC, however, then all users can use the DBA-privileged UDR. For additional information about DBA-privileged UDRs, see “Ownership of Created Database Objects” on page 2-162.

If you omit the DBA keyword, the UDR is known as an *owner-privileged UDR*.

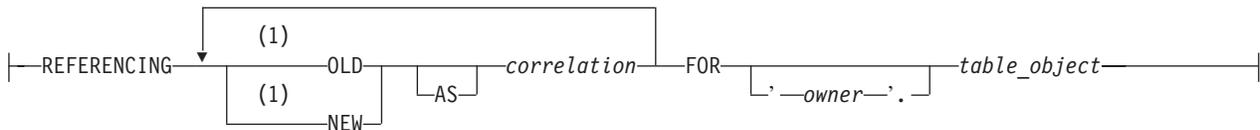
When you create an owner-privileged UDR in an ANSI-compliant database, only you can execute the UDR. Before other users can execute an owner-privileged UDR, its owner must grant the Execute privilege, either to individual users, or to roles, or to PUBLIC.

If you create an owner-privileged UDR in a database that is not ANSI compliant, anyone can execute the UDR because PUBLIC is granted the Execute privilege by default. To restrict access to an owner-privileged UDR to specific users, the owner must revoke the Execute privilege on the UDR from PUBLIC, and then grant it to specified users or roles. Setting the **NODEFDAC** environment variable to yes prevents privileges on any UDR from being granted to PUBLIC. If this environment variable is set to yes, no one besides the owner of the UDR can invoke it unless the owner grants the Execute privilege for that UDR to other users.

The REFERENCING and FOR Clauses

The REFERENCING clause can declare correlation names for the original value and for the updated value in columns of the *table_object* that the FOR clause specifies.

REFERENCING and FOR Clauses:



Notes:

- 1 Use path no more than once

Element	Description	Restrictions	Syntax
<i>correlation</i>	Name that you declare here to qualify an old or new column value (as <i>correlation.column</i>) in a trigger routine	Must not be <i>table_object</i>	"Identifier " on page 5-21
<i>owner</i>	Owner of <i>table_object</i>	Must own <i>table_object</i>	" Owner Name " on page 5-45
<i>table_object</i>	Name or synonym for the table or view whose triggers can call this procedure	Must exist in the local database	"Identifier " on page 5-21

If you include the REFERENCING and FOR *table_object* clauses immediately after the parameter list of the CREATE PROCEDURE statement, the routine that you create is known as a *trigger UDR* or *trigger routine*. The FOR clause specifies the table or view whose triggers can invoke the routine from the FOR EACH ROW section of their Triggered Action list.

In the REFERENCING clause, the OLD *correlation* specifies a prefix by which the trigger routine can reference the value that a column of *table_object* had before the trigger routine modifies that column value. The NEW *correlation* specifies a prefix for referencing the new value that the trigger routine assigns to the column. Whether the trigger routine can use correlation names to reference the OLD column value, the NEW column value, or both values depends on the type of triggering event:

- A trigger routine invoked by an Insert trigger can reference only the NEW correlation name.
- A trigger routine invoked by a Delete trigger or by a Select trigger can reference only the OLD correlation name.
- A trigger routine invoked by an Update trigger can reference both the OLD and the NEW correlation names.

For information about how to use the *correlation.column* notation in triggered actions, see "REFERENCING Clauses" on page 2-254.

If you include the REFERENCING and FOR *table_object* clauses immediately after the parameter list of the CREATE PROCEDURE statement, the procedure that you

create is known as a *trigger UDR* or *trigger routine*. The FOR clause specifies the table or view whose triggers can invoke this procedure from the FOR EACH ROW section of their Triggered Action list.

Besides the general requirements for any Dynamic Server UDR that is written in the SPL language, trigger routines can support certain additional syntax features, and are subject to certain restrictions, that are not features (or that are not restrictions) for ordinary UDRs that are not trigger routines:

- A trigger routine must include the FOR *table_object* clause that specifies the name of the table or view in the local database whose triggers can invoke this routine.
- A trigger routine can also include the REFERENCING clause to declare correlation names for OLD and NEW values that SPL statements in the UDR can reference.
- Trigger routines can be invoked only in the FOR EACH ROW section of the Triggered Action list in the trigger definition.
- Correlated variables for OLD or NEW values can appear in the IF statement of SPL and in CASE expressions.
- Correlated variables for OLD values cannot be on the left-hand side of a LET expression
- Correlated variables for NEW values cannot be on the left-hand side of a LET expression if the FOR clause specifies a view whose INSTEAD OF trigger action list invokes the trigger routine.
- Only correlated variables for NEW values can be on the left-hand side of a LET expression that references correlated variables. In this case, the FOR clause must specify a table, rather than a view, and the trigger whose action invokes the SPL routine cannot be an INSTEAD OF trigger.
- Both OLD and NEW values can be on the right-hand side of a LET expression.
- The Boolean operators SELECTING, INSERTING, DELETING, and UPDATING are valid in trigger routines (and only in trigger routines and in other UDRs that are invoked in triggered action statements) in contexts where Boolean expressions are valid. These operators return TRUE ('t') if the triggering event matches the DML operation referenced by the name of the operator, and they return FALSE ('f') otherwise.
- If a single triggering event activates multiple triggers on the same table or view, then all of the BEFORE actions take place before any of the FOR EACH ROW actions, and all of the AFTER actions follow the FOR EACH ROW actions. The order of execution of different triggers on the same event is not guaranteed.
- Trigger routines must be written in the SPL language. They cannot be written in an external language, like the C or Java language, but they can include calls to external language routines, such as the **mi_trigger** application programming interface for trigger introspection.
- Trigger routines cannot reference savepoints. Any changes to the data values or to the schema of the database by a triggered action must be committed or rolled back in their entirety. Partial rollback of a triggered action is not supported.

For more information about the **mi_trigger** API, refer to the *IBM Informix Datablade API Programmer's Guide* and to the *IBM Informix Datablade API Function Reference*.

If you include the REFERENCING clause but omit the FOR clause, or if you include the FOR clause but omit the REFERENCING clause, the CREATE PROCEDURE statement fails with an error.

If you omit the REFERENCING and FOR clauses, the UDR cannot use the SELECTING, INSERTING, DELETING, and UPDATING operators, and cannot declare variables that can represent and manipulate column values in triggered actions on the table or view that the trigger definition specifies.

See the “REFERENCING Clauses” on page 2-254 section in the CREATE TRIGGER statement description for the syntax of the REFERENCING clause for Delete, Insert, Select, and Update triggers on tables, and for Delete, Insert, and Update INSTEAD OF triggers on views.

Procedure Names in Extended Parallel Server

In Extended Parallel Server, the name for any SPL routine that you create must be unique among the names of all SPL routines in the database.

Procedure Names in Dynamic Server

Because Dynamic Server offers *routine overloading*, you can define more than one user-defined routine (UDR) with the same name, but different parameter lists. You might want to overload UDRs in the following situations:

- You create a UDR with the same name as a built-in routine (such as `equal()`) to process a new user-defined data type.
- You create *type hierarchies* in which subtypes inherit data representation and UDRs from supertypes.
- You create *distinct types*, which are data types that have the same internal storage representation as an existing data type, but have different names and cannot be compared to the source type without casting. Distinct types inherit UDRs from their source types.

For a brief description of the routine signature that uniquely identifies each UDR, see “Routine Overloading and Routine Signatures (IDS)” on page 5-19.

Using the SPECIFIC Clause to Specify a Specific Name

You can declare a *specific name* that is unique in the database for a user-defined procedure. A specific name is useful when you are overloading a procedure.

DOCUMENT Clause

The quoted string in the DOCUMENT clause provides a synopsis and description of a UDR. The string is stored in the `sysprocbody` system catalog table and is intended for the user of the UDR.

Anyone with access to the database can query the `sysprocbody` system catalog table to obtain a description of one or all the UDRs stored in the database. A UDR or application program can query the system catalog tables to fetch the DOCUMENT clause and display it for a user.

For example, to find the description of the SPL procedure `raise_prices`, shown in “SPL Procedures” on page 2-160, enter a query such as this example:

```
SELECT data FROM sysprocbody b, sysprocedures p
WHERE b.procid = p.procid
      --join between the two catalog tables
      AND p.procname = 'raise_prices'
      -- look for procedure named raise_prices
      AND b.datakey = 'D';-- want user document
```

The preceding query returns the following text:

```
USAGE: EXECUTE PROCEDURE raise_prices( xxx )
xxx = percentage from 1 - 100
```

For external procedures, you can use a DOCUMENT clause at the end of the CREATE PROCEDURE statement, whether or not you use the END PROCEDURE keywords.

Using the WITH LISTING IN Option

The WITH LISTING IN clause specifies a filename where compile time warnings are sent. After you compile a UDR, this file holds one or more warning messages. This listing file is created on the computer where the database resides.

If you do not use the WITH LISTING IN clause, the compiler does not generate a list of warnings.

On UNIX, if you specify a filename but not a directory, this listing file is created in your home directory on the computer where the database resides. If you do not have a home directory on this computer, the file is created in the root directory (the directory named "/").

On Windows, if you specify a filename but not a directory, this listing file is created in your current working directory if the database is on the local computer. Otherwise, the default directory is %INFORMIXDIR%\bin.

SPL Procedures

SPL procedures are UDRs written in Stored Procedure Language (SPL) that do not return a value. To write and register an SPL routine, use the CREATE PROCEDURE statement. Embed appropriate SQL and SPL statements between the CREATE PROCEDURE and END PROCEDURE keywords. You can also follow the UDR definition with the DOCUMENT and WITH FILE IN options.

SPL routines are parsed, optimized (as far as possible), and stored in the system catalog tables in executable format. The body of an SPL routine is stored in the **sysprocbody** system catalog table. Other information about the routine is stored in other system catalog tables, including **sysprocedures**, **sysprocplan**, and **sysprocauth**.

If the Statement Block portion of the CREATE PROCEDURE statement is empty, no operation takes place when you call the procedure. You might use such a "dummy" procedure in the development stage when you intend to establish the existence of a procedure but have not yet coded it.

If you specify an optional clause after the parameter list, you must place a semicolon after the clause that immediately precedes the Statement Block.

The following example creates an SPL procedure:

```
CREATE PROCEDURE raise_prices ( per_cent INT )
  UPDATE stock SET unit_price =
    unit_price + (unit_price * (per_cent/100));
END PROCEDURE
DOCUMENT "USAGE: EXECUTE PROCEDURE raise_prices( xxx )",
"xxx = percentage from 1 - 100 "
WITH LISTING IN '/tmp/warn_file';
```

External Procedures (IDS)

External procedures are procedures you write in an external language that the database server supports.

To create a C user-defined procedure

1. Write a C function that does not return a value.
2. Compile the C function and store the compiled code in a shared library (the shared-object file for C).
3. Register the C function in the database server with the CREATE PROCEDURE statement.

To create a user-defined procedure written in the Java language

1. Write a Java static method, which can use the JDBC functions to interact with the database server.
2. Compile the Java source and create a JAR file (the shared-object file).
3. Execute the **install_jar()** procedure with the EXECUTE PROCEDURE statement to install the JAR file in the current database.
4. If the UDR uses user-defined types, create a mapping between SQL data types and Java classes, using the **setUDTextName()** procedure that is explained in "EXECUTE PROCEDURE " on page 2-364.
5. Register the UDR with the CREATE PROCEDURE statement. (If an external routine returns a value, you must register it with the CREATE FUNCTION statement, rather than CREATE PROCEDURE.)

Rather than storing the body of an external routine directly in the database, the database server stores only the pathname of the shared-object file that contains the compiled version of the routine. The database server executes an external routine by invoking the external object code.

When the IFX_EXTEND_ROLE configuration parameter is set to ON, only users who have the built-in EXTEND role can create external procedures.

Registering a User-Defined Procedure

This example registers a C user-defined procedure named **check_owner()** that takes one argument of the type LVARCHAR. The external routine reference specifies the path to the C shared library where the procedure object code is stored. This library contains a C function **unix_owner()**, which is invoked during execution of the **check_owner()** procedure.

```
CREATE PROCEDURE check_owner ( owner lvarchar )
  EXTERNAL NAME "/usr/lib/ext_lib/genlib.so(unix_owner)"
  LANGUAGE C
END PROCEDURE;
```

This example registers a user-defined procedure named **showusers()** that is written in the Java language:

```
CREATE PROCEDURE showusers()
  WITH (CLASS = "jvp") EXTERNAL NAME 'admin_jar:admin.showusers'
  LANGUAGE JAVA;
```

The EXTERNAL NAME clause specifies that the Java implementation of the **showusers()** procedure is a method called **showusers()**, which resides in the **admin** Java class that resides in the **admin_jar** JAR file.

Ownership of Created Database Objects

The user who creates an owner-privileged UDR owns any database objects that the UDR creates when it executes, unless some other *owner* is specified for the object. In other words, the UDR owner, not the user who executes the owner-privileged UDR, is the owner of any database objects created by the UDR unless another owner is specified in the DDL statement that creates the database object.

In the case of a DBA-privileged UDR, however, the user who executes the UDR, not the UDR owner, owns any database objects that the UDR creates, unless some other owner is specified for the database object within the UDR.

For examples, see “Ownership of Created Database Objects” on page 2-121 in the description of the CREATE FUNCTION statement.

Examples

For this example, assume that you have two overloaded procedures that are defined as follows:

```
CREATE PROCEDURE raise_prices ( per_cent INT)
  UPDATE stock SET unit_price = unit_price + (unit_price * (per_cent/100) );
END PROCEDURE
```

```
CREATE PROCEDURE raise_prices ( per_cent INT, selected_unit CHAR )
  UPDATE stock SET unit_price = unit_price + (unit_price * (per_cent/100) )
  where unit=selected_unit;
END PROCEDURE
```

In order to refer to the above procedures, you would need to provide the procedure name followed by the parameter list, as in the following examples:

```
DROP PROCEDURE raise_prices(INT);
DROP PROCEDURE raise_prices(INT, CHAR);
```

A more convenient way is to use the specific name to identify each of them. The following example will create the procedure using the specific name:

```
CREATE PROCEDURE raise_prices ( per_cent INT ) SPECIFIC
  raise_prices_all
  UPDATE stock SET unit_price = unit_price + (unit_price * (per_cent/100) );
END PROCEDURE
```

```
DROP SPECIFIC PROCEDURE raise_prices_all;
```

```
CREATE PROCEDURE raise_prices ( per_cent INT, selected_unit CHAR )
  SPECIFIC raise_prices_by_unit
  UPDATE stock SET unit_price = unit_price + (unit_price * (per_cent/100) )
  where unit=selected_unit;
END PROCEDURE
```

We can simply drop them using their specific names:

```
DROP SPECIFIC PROCEDURE raise_prices_by_all;
DROP SPECIFIC PROCEDURE raise_prices_by_unit;
```

Related Information

Related statements: “ALTER FUNCTION ” on page 2-28, “ALTER PROCEDURE” on page 2-32, “ALTER ROUTINE” on page 2-33, “CREATE FUNCTION ” on page 2-113, “CREATE FUNCTION FROM ” on page 2-122, “CREATE PROCEDURE FROM ” on page 2-163, “DROP FUNCTION ” on page 2-327, “DROP PROCEDURE ” on page 2-332, “DROP ROUTINE ” on page 2-335, “EXECUTE FUNCTION ” on page 2-356,

“EXECUTE PROCEDURE ” on page 2-364, “GRANT ” on page 2-395, “PREPARE ” on page 2-468, “REVOKE” on page 2-494, and “UPDATE STATISTICS ” on page 2-705.

For a discussion of how to create and use SPL routines, see the *IBM Informix Guide to SQL: Tutorial*. For a discussion of external routines, see *IBM Informix User-Defined Routines and Data Types Developer’s Guide*.

For information about how to create C UDRs, see the *IBM Informix Datablade API Programmer’s Guide*. For more information on the **NODEFDAC** environment variable and the related system catalog tables (**sysprocedures**, **sysprocplan**, **sysprocbody** and **sysprocauth**), see the *IBM Informix Guide to SQL: Reference*.

CREATE PROCEDURE FROM

Use the CREATE PROCEDURE FROM statement to access a user-defined procedure. The actual text of the CREATE PROCEDURE statement resides in a separate file.

This statement is an extension to the ANSI/ISO standard for SQL. You can use this statement with Informix ESQL/C.

In Extended Parallel Server, which does not support the CREATE FUNCTION FROM statement, you can use the CREATE PROCEDURE FROM statement to create a SPL routine from a file, whether or not the SPL routine returns a value.

Syntax

►► CREATE PROCEDURE FROM 'file' file_var ◀◀

Element	Description	Restrictions	Syntax
<i>file</i>	Pathname and filename of file that contains full text of a CREATE PROCEDURE statement. Default pathname is the current directory.	Must exist, and can contain only one CREATE PROCEDURE statement. See also “Default Directory That Holds the File” on page 2-164.	Operating-system specific
<i>file_var</i>	Name of a program variable that contains <i>file</i> specification	Must be of a character data type; its contents have same restrictions as <i>file</i>	Language specific

Usage

You cannot create a user-defined procedure directly in an Informix ESQL/C program. That is, the program cannot contain the CREATE PROCEDURE statement.

To use a user-defined procedure in an ESQL/C program

1. Create a source file with the CREATE PROCEDURE statement.
2. Use the CREATE PROCEDURE FROM statement to send the contents of this source file to the database server for execution.

The file can contain only one CREATE PROCEDURE statement.

For example, suppose that the following CREATE PROCEDURE statement is in a separate file, called **raise_pr.sql**:

```

CREATE PROCEDURE raise_prices( per_cent INT )
  UPDATE stock -- increase by percentage;
  SET unit_price = unit_price +
    ( unit_price * (per_cent / 100) );
END PROCEDURE;

```

In the Informix ESQL/C program, you can access the **raise_prices()** SPL procedure with the following CREATE PROCEDURE FROM statement:

```
EXEC SQL create procedure from 'raise_pr.sql';
```

In Dynamic Server, if you are not sure whether the UDR in the file returns a value, use the CREATE ROUTINE FROM statement.

When the IFX_EXTEND_ROLE configuration parameter is set to ON, only users who have the built-in EXTEND role can create external routines.

User-defined procedures, like user-defined functions, use the collating order that was in effect when they were created. See “SET COLLATION” on page 2-584 for information about using non-default collation.

Default Directory That Holds the File

The database server treats the specified filename (and any pathname) as relative.

On UNIX, if you specify a simple filename instead of a full pathname as the *file* parameter, the client application looks for the file in your home directory on the computer where the database resides. If you do not have a home directory on this computer, the default directory is the root directory.

On Windows, if you specify a filename but no directory as the *file* parameter, the client application looks for the file in your current working directory if the database is on the local computer. Otherwise, the default directory is **%INFORMIXDIR%\bin**.

Important: The Informix ESQL/C preprocessor does not process the contents of the file that you specify. It only sends the contents to the database server for execution. Therefore, there is no syntactic check that the file that you specify in CREATE PROCEDURE FROM actually contains a CREATE PROCEDURE statement. To improve readability of the code, however, it is recommended that you match these two statements.

Related Information

Related statements: “CREATE PROCEDURE ” on page 2-153, “CREATE FUNCTION FROM ” on page 2-122, and “CREATE ROUTINE FROM ” on page 2-166

CREATE ROLE

Use the CREATE ROLE statement to declare and register a new role.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax

► CREATE ROLE *role* [*'role'*] ◄

Element	Description	Restrictions	Syntax
<i>role</i>	Name declared here for a role that the DBA creates	Must be unique among <i>role</i> and user names in the database. Maximum number of bytes is 32.	" Owner Name " on page 5-45

Usage

CREATE ROLE declares a new role and registers it in the system catalog. A role can associate a set of authorization identifiers with a set of access privileges on database objects. The system catalog maintains information about the roles (and their corresponding privileges) that are granted to users or to other roles.

Only the database administrator (DBA) can use CREATE ROLE to create a new role. The DBA can assign the privileges required for some work task to a role, such as **engineer**, and then use the GRANT statement to assign that role to specific users, instead of granting that set of privileges to each user individually.

The *role* name is an *authorization identifier*. It cannot be a user name that is known to the database server or to the operating system of the database server. The *role* name cannot already be listed in the **username** column of the **sysusers** system catalog table, nor in the **grantor** or **grantee** columns of the **systabauth**, **syscolauth**, **sysfragauth**, **sysprocauth**, or **sysroleauth** system catalog tables.

In Dynamic Server, the *role* name also cannot match the name of any user or role that is already listed in the **grantor** or **grantee** columns of the **sysxdtypeauth** system catalog table, nor any built-in role, such as EXTEND or DBSECADM.

After a role is created, the DBA can use the GRANT statement to assign the role to PUBLIC, to users, or to other roles, and to grant specific privileges to the role. (A role cannot, however, hold database-level privileges.) After a role is granted successfully to a user or to PUBLIC, the user must use the SET ROLE statement to enable the role. Only then can the user exercise the privileges of the role.

To create the role **engineer**, for example, enter the following statement:

```
CREATE ROLE engineer;
```

To grant access privileges to the role **engineer**, the DBA can issue GRANT statements that include **engineer** in the list of grantees:

```
GRANT USAGE ON LANGUAGE SPL TO engineer;
```

To assign the role **engineer** to user **kaycee**, the DBA could issue this statement:

```
GRANT engineer TO kaycee;
```

To activate the role **engineer**, user **kaycee** must issue the following statement:

```
SET ROLE engineer;
```

If this SET ROLE statement is successful, user **kaycee** acquires whatever privileges have been granted to the role **engineer**, in addition to any other privileges that **kaycee** already holds as an individual or as PUBLIC.

A user can be granted several roles, but no more than one non-default role, as specified by SET ROLE, can be enabled for any user at a given time.

An exception to requiring SET ROLE to explicitly enable a role is any default role that the DBA specifies in the GRANT DEFAULT ROLE *role* TO *user* statement. If that statement succeeds, the default *role* is automatically enabled when *user* connects to the database. Any role can be a default role. (Similarly, users to whom the Dynamic Server DBSA grants the EXTEND role need not execute SET ROLE before they can create and drop external routines and shared libraries.)

CREATE ROLE, when used with the GRANT and SET ROLE statements, enables a DBA to create one set of privileges for a role and then grant the role to many users, instead of granting the same set of privileges individually to many users.

With the GRANT DEFAULT ROLE and SET ROLE DEFAULT statements, *default roles* enable a DBA to assign privileges to a role that is activated automatically when any user who holds that default role connects to the database. This feature is useful when an application performs operations that require specific access privileges, but the application does not include SET ROLE statements.

The REVOKE statement can cancel access privileges of a role, remove users from a role, or cancel the default status of a role for one or more users. A role exists until either the DBA or a user to whom the role was granted with the WITH GRANT OPTION keywords uses the DROP ROLE statement to drop the role.

Related Information

Related statements: “DROP ROLE ” on page 2-334, “GRANT ” on page 2-395, “REVOKE” on page 2-494, and “SET ROLE ” on page 2-647

For a discussion of how to use roles, see the *IBM Informix Database Design and Implementation Guide*.

CREATE ROUTINE FROM

Use the CREATE ROUTINE FROM statement to register a UDR by referencing the text of a CREATE FUNCTION or CREATE PROCEDURE statement that resides in a separate file. Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL. You can use this statement with ESQL/C.

Syntax

```

▶▶ CREATE ROUTINE FROM 'file'
                       └─ file_var ─┘

```

Element	Description	Restrictions	Syntax
<i>file</i>	Pathname and filename for the text of a CREATE PROCEDURE or CREATE FUNCTION statement. Default path is the current directory.	Must exist and can contain only one CREATE FUNCTION or CREATE PROCEDURE statement.	Operating-system dependent
<i>file_var</i>	Name of a program variable that contains <i>file</i> specification	Must be a character data type; contents must satisfy <i>file</i> restrictions	Language specific

Usage

ESQL/C programs cannot use the CREATE FUNCTION or CREATE PROCEDURE statement directly to define a UDR. You must instead do this:

1. Create a source file with the CREATE FUNCTION or CREATE PROCEDURE statement.
2. Execute the CREATE ROUTINE FROM statement from an ESQL/C program to send the contents of this source file to the database server for execution. The file that you specify can contain only one CREATE FUNCTION or CREATE PROCEDURE statement.

The file specification that you provide is relative. If you include no pathname, the client application looks for the file in the current directory.

If you do not know at compile time whether the UDR in the file is a function or a procedure, use the CREATE ROUTINE FROM statement in the Informix ESQL/C program. If you know whether the UDR is a function or a procedure, you can improve the readability of your code by using the matching SQL statement to access the source file:

- To access user-defined functions, use CREATE FUNCTION FROM.
- To access user-defined procedures, use CREATE PROCEDURE FROM.

When the IFX_EXTEND_ROLE configuration parameter is set to 0N, only users to whom the Database Server Administrator (DBSA) has granted the built-in EXTEND role can create external routines.

Routines use the collating order that was in effect when they were created. See "SET COLLATION" on page 2-584 for information about using non-default collation.

Examples

```
EXEC SQL CREATE ROUTINE FROM 'del_ord.sql';
```

ESQL/C source code example:

```
#include <stdio.h>

main()
{
    printf( "CREATE ROUTINE FROM ESQL Program running.\n\n");
    EXEC SQL WHENEVER ERROR STOP;
    EXEC SQL connect to 'stores_demo';

    EXEC SQL CREATE ROUTINE FROM 'del_ord.sql';

    EXEC SQL disconnect current;
    printf("\nCREATE ROUTINE Sample Program over.\n\n");

    exit(0);
}

del_ord.sql

CREATE FUNCTION delete_order( p_order_num int) RETURNING int, int;
DEFINE item_count int;
SELECT count(*) INTO item_count FROM items
```

```

        WHERE order_num = p_order_num;
DELETE FROM orders WHERE order_num = p_order_num;
RETURN p_order_num, item_count;
END FUNCTION;

```

Related Information

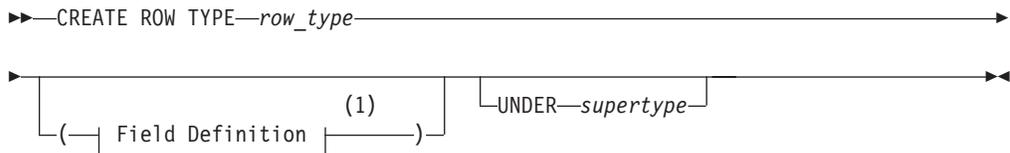
Related statements: “CREATE FUNCTION ” on page 2-113, “CREATE FUNCTION FROM ” on page 2-122, “CREATE PROCEDURE ” on page 2-153, and “CREATE PROCEDURE FROM ” on page 2-163

CREATE ROW TYPE

Use the CREATE ROW TYPE statement to create a named ROW type.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 See “Field Definition” on page 2-170

Element	Description	Restrictions	Syntax
<i>row_type</i>	Name that you declare here for a new named ROW data type	See “Procedure for Creating a Subtype” on page 2-170.	“Identifier ” on page 5-21
<i>supertype</i>	Name of the supertype within a data type inheritance hierarchy	Must already exist in the database as a named ROW type	“Data Type ” on page 4-21

Usage

The CREATE ROW TYPE statement declares a named ROW data type and registers it in the system catalog. You can assign a named ROW data type to a table or view to create a *typed table* or *typed view*. You can also define a column as a named ROW type. Although you can assign a ROW type to a table to define the schema of the table, ROW data types are not the same as table rows. Table rows consist of one or more *columns*; ROW data types consist of one or more *fields*, defined using the Field Definition syntax.

A named ROW data type is valid in most contexts where you can specify a data type. Named ROW types are said to be *strongly typed*. No two named ROW types are equivalent, even if they are structurally equivalent.

ROW types without identifiers are called *unnamed ROW types*. Any two unnamed ROW types are considered equivalent if they are structurally equivalent. For more information, see “ROW Data Types” on page 4-34.

Privileges on named ROW type columns are the same as privileges on any column. For more information, see “Table-Level Privileges” on page 2-399. (To see what privileges you have on a column, check the **syscolauth** system catalog table, which is described in the *IBM Informix Guide to SQL: Reference*.)

Privileges on Named Row Data Types

Privileges for operations on a typed table (a table that is assigned a named ROW data type) are the same as privileges on any table. For more information, see “Table-Level Privileges” on page 2-399. This table shows which access privileges you need to create a named ROW type.

Task	Privileges Required
Create a named ROW type	Resource privilege on the database
Create a named ROW type as a subtype under a supertype	Under privilege on the supertype, as well as the Resource privilege

For information about Resource and Under privileges and the ALL keyword in the context of privileges, see the “GRANT ” on page 2-395 statement.

To find out what privileges exist on a ROW type, check the **sysxdtypes** system catalog table for the *owner* name and the **sysxdttypeauth** system catalog table for privileges on the ROW type that might have been granted to users or to roles.

To find out what privileges you have on a given table, check the **systabauth** system catalog table. For more information on system catalog tables, see the *IBM Informix Guide to SQL: Reference*.

Inheritance and Named ROW Types

A named ROW type can belong to an inheritance hierarchy, as either a subtype or a supertype. Use the UNDER clause in the CREATE ROW TYPE statement to create a named ROW data type as a subtype of an existing ROW data type.

The supertype must also be a named ROW data type. If you create a named ROW data type under an existing supertype, then the new type name *row_type* becomes the name of the subtype.

When you create a named ROW type as a subtype, the subtype inherits all fields of the supertype. In addition, you can add new fields to the subtype when you create it. The new fields are specific to the subtype alone.

You cannot substitute a ROW type in an inheritance hierarchy for its supertype or for its subtype. For example, consider a type hierarchy in which **person_t** is the supertype and **employee_t** is the subtype. If a column is of type **person_t**, the column can only contain **person_t** data. It cannot contain **employee_t** data. Likewise, if a column is of type **employee_t**, the column can only contain **employee_t** data. It cannot contain **person_t** data.

Creating a Subtype

In most cases, you add new fields when you create a named ROW type as a subtype of another named ROW type (its supertype). To create the fields of a named ROW type, use the field definition clause, as described in “Field Definition” on page 2-170. When you create a subtype, you must use the UNDER keyword to

associate the supertype with the named ROW type that you want to create. The next example creates the **employee_t** type under the **person_t** type:

```
CREATE ROW TYPE employee_t (salary NUMERIC(10,2),
    bonus NUMERIC(10,2)) UNDER person_t;
```

The **employee_t** type inherits all the fields of **person_t** and has two additional fields: **salary** and **bonus**; but the **person_t** type is not altered.

Type Hierarchies

When you create a subtype, you create a *type hierarchy*. In a type hierarchy, each subtype that you create inherits its properties from a single supertype. If you create a named ROW type **customer_t** under **person_t**, **customer_t** inherits all the fields of **person_t**. If you create another named ROW type, **salesrep_t** under **customer_t**, **salesrep_t** inherits all the fields of **customer_t**.

Thus, **salesrep_t** inherits all the fields that **customer_t** inherited from **person_t** as well as all the fields defined specifically for **customer_t**. For a discussion of type inheritance, refer to the *IBM Informix Guide to SQL: Tutorial*.

Procedure for Creating a Subtype

Before you create a named ROW type as a subtype in an inheritance hierarchy, check the following information:

- Verify that you are authorized to create new data types. You must have the Resource privilege on the database. You can find this information in the **sysusers** system catalog table.
- Verify that the supertype exists. You can find this information in the **sysxtotypes** system catalog table.
- Verify that you are authorized to create subtypes to that supertype. You must have the Under privilege on the supertype. You can find this information in the **sysusers** system catalog table.
- Verify that the name that you declare for the named ROW type is unique. In an ANSI-compliant database, the *owner.type* combination must be unique within the database. In a database that is not ANSI-compliant, the name must be unique among data type names in the database. To verify whether the name for a new data type is unique, check the **sysxtotypes** system catalog table. The name must not be the name of an existing data type.
- If you are defining fields for the ROW type, check that no duplicate field names exist in both new and inherited fields.

Important: When you create a subtype, you cannot redefine fields that it inherited for its supertype. If you attempt to redefine these fields, the database server returns an error.

You cannot apply constraints to named ROW data types, but you can specify constraints when you create or alter a table that uses the named ROW types. You can also specify NOT NULL constraints on individual fields of a ROW type.

Field Definition

Use the Field Definition clause to define a new field in a named ROW type.

Field Definition:

|—*field*—*data_type*—|
└─ NOT NULL ─┘

Element	Description	Restrictions	Syntax
<i>data_type</i>	Data type of the field	See “Restrictions on Serial and Simple-Large-Object Data Types.”	“Identifier ” on page 5-21
<i>field</i>	Name of a field in <i>row_type</i>	Must be unique among field names of this ROW type and of its supertype	“Identifier ” on page 5-21

The NOT NULL constraint on the named ROW type field applies to the corresponding columns when a typed table of the named ROW type is created.

Restrictions on Serial and Simple-Large-Object Data Types

Serial and simple-large-object data types cannot be nested within a table. Therefore, if a ROW type contains a BYTE, TEXT, SERIAL, BIGSERIAL, or SERIAL8 field, you cannot use the ROW type to define a column in a table that is not based on a ROW type. For example, the following code example produces an error:

```
CREATE ROW TYPE serialtype (s serial, s8 serial8);  
CREATE TABLE tab1 (col1 serialtype); --INVALID CODE
```

You cannot create a ROW type that has a BYTE or TEXT value that is stored in a separate storage space. That is, you cannot use the IN clause to specify the storage location. For example, the following example produces an error:

```
CREATE ROW TYPE row1 (field1 byte IN blobspace1); --INVALID CODE
```

A table hierarchy can include no more than one SERIAL, BIGSERIAL, or SERIAL8 column. If a supertable has a SERIAL column, none of its subtables can contain a SERIAL column (but a subtable can have a BIGSERIAL or SERIAL8 column if no other subtable contains a BIGSERIAL or SERIAL8 column, respectively). Consequently, when you create the named ROW types on which the table hierarchy is to be based, they can contain at most one SERIAL and one BIGSERIAL or SERIAL8 field among them.

You cannot set the starting SERIAL, BIGSERIAL, or SERIAL8 value in the CREATE ROW TYPE statement. To modify the value for a serial field, you must use either the MODIFY clause of the ALTER TABLE statement, or else use the INSERT statement to insert a value that is larger than the current maximum (or default) serial value.

Serial fields in ROW types have performance implications across a table hierarchy. To insert data into a subtable whose supertable (or its supertable) contains the serial counter, the database server must also open the supertable, update the serial value, and close the supertable, thus adding extra overhead.

In contexts where these restrictions or performance issues for SERIAL, BIGSERIAL, or SERIAL8 data types conflict with your design goals, you might consider using sequence objects to emulate the functionality of serial fields or serial columns.

Related Information

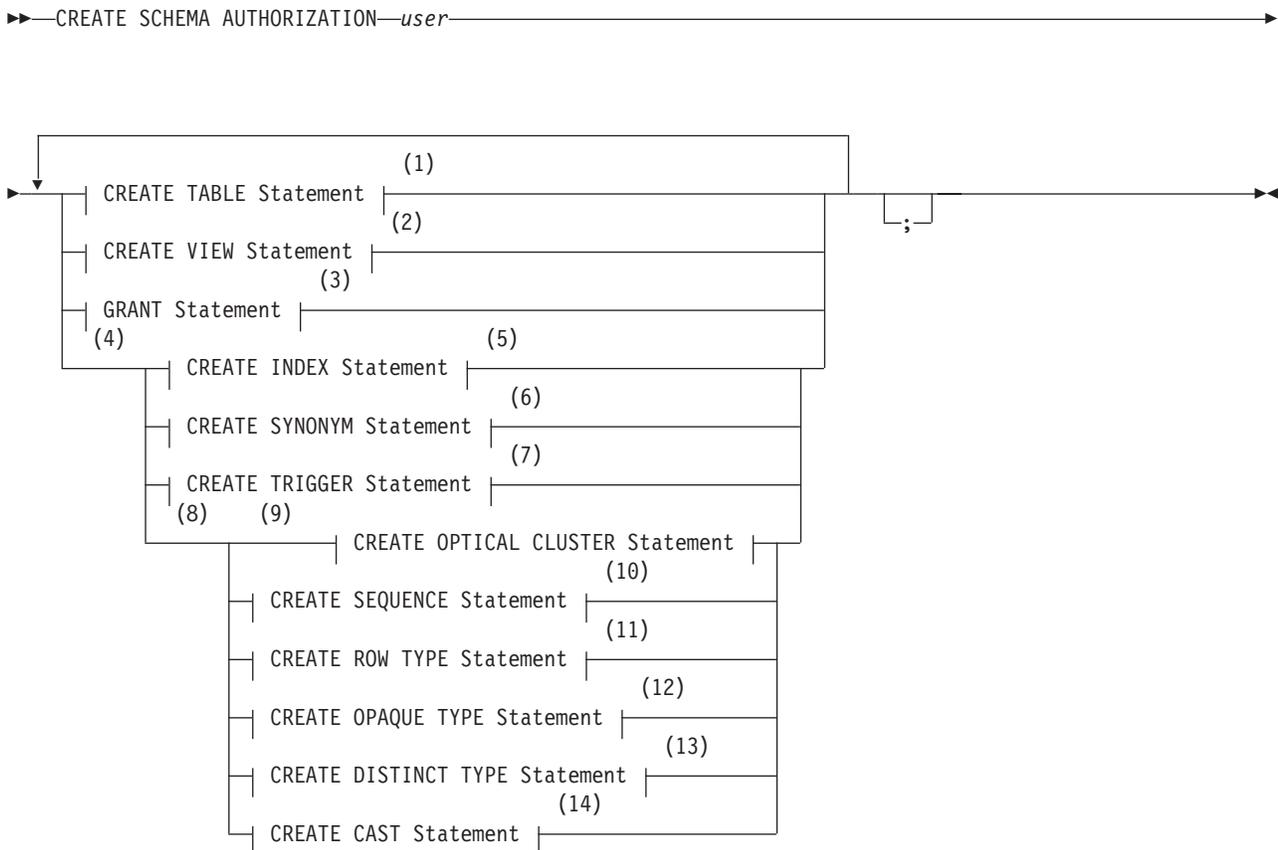
Related statements: "DROP ROW TYPE " on page 2-336, "CREATE TABLE " on page 2-191, "CREATE CAST " on page 2-95, "GRANT " on page 2-395, and "REVOKE" on page 2-494

For a discussion of named ROW types, see the *IBM Informix Database Design and Implementation Guide* and the *IBM Informix Guide to SQL: Reference*.

CREATE SCHEMA

Use the CREATE SCHEMA statement to issue a block of data definition language (DDL) and GRANT statements as a unit. Use this statement with DB–Access.

Syntax



Notes:

- 1 See "CREATE TABLE " on page 2-191
- 2 See "CREATE VIEW " on page 2-274
- 3 See "GRANT " on page 2-395
- 4 Informix extension
- 5 See "CREATE INDEX " on page 2-124
- 6 See "CREATE SYNONYM " on page 2-187
- 7 See "CREATE TRIGGER" on page 2-239

- 8 Dynamic Server only
- 9 Optical Subsystem only. See the *IBM Informix: Optical Subsystem Guide*.
- 10 See "CREATE SEQUENCE" on page 2-184
- 11 See "CREATE ROW TYPE " on page 2-168
- 12 See "CREATE OPAQUE TYPE " on page 2-145
- 13 See "CREATE DISTINCT TYPE " on page 2-100
- 14 See "CREATE CAST " on page 2-95

Element	Description	Restrictions	Syntax
<i>user</i>	User who owns the database objects that this statement creates	If you have DBA privileges, you can specify the name of any user. Otherwise, you must have the Resource privilege, and you must specify your own user name.	" Owner Name " on page 5-45

Usage

The CREATE SCHEMA statement allows the DBA to specify an owner for all database objects that the CREATE SCHEMA statement creates. You cannot issue CREATE SCHEMA until you have created the database that stores the objects.

Users with the Resource privilege can create a schema for themselves. In this case, *user* must be the name of the person with the Resource privilege who is running the CREATE SCHEMA statement. Anyone with the DBA privilege can also create a schema for someone else. In this case, *user* can specify a user other than the person who is running the CREATE SCHEMA statement.

You can put CREATE and GRANT statements in any logical order, as the following example shows. Statements are considered part of the CREATE SCHEMA statement until a semicolon (;) or an end-of-file symbol is reached.

```
CREATE SCHEMA AUTHORIZATION sarah
  CREATE TABLE mytable (mytime DATE, mytext TEXT)
  GRANT SELECT, UPDATE, DELETE ON mytable TO rick
  CREATE VIEW myview AS
    SELECT * FROM mytable WHERE mytime > '12/31/2004'
  CREATE INDEX idxtime ON mytable (mytime);
```

Creating Database Objects Within CREATE SCHEMA

All database objects that a CREATE SCHEMA statement creates are owned by *user*, even if you do not explicitly name each database object. If you are the DBA, you can create database objects for another user. If you are not the DBA, specifying an owner other than yourself results in an error message.

You can only grant privileges with the CREATE SCHEMA statement; you cannot use CREATE SCHEMA to revoke or to drop privileges.

If you create a database object or use the GRANT statement outside a CREATE SCHEMA statement, you receive warnings if you use the **-ansi** flag or set **DBANSIWARN**.

Related Information

Related statements: "CREATE CAST " on page 2-95, "CREATE DISTINCT TYPE " on page 2-100, "CREATE INDEX " on page 2-124, "CREATE OPAQUE TYPE " on page 2-145

page 2-145, "CREATE OPCLASS" on page 2-149, "CREATE ROW TYPE " on page 2-168, "CREATE SEQUENCE" on page 2-184, "CREATE SYNONYM " on page 2-187, "CREATE TABLE " on page 2-191, "CREATE VIEW " on page 2-274, and "GRANT " on page 2-395

For a discussion of how to create a database, see the *IBM Informix Database Design and Implementation Guide*.

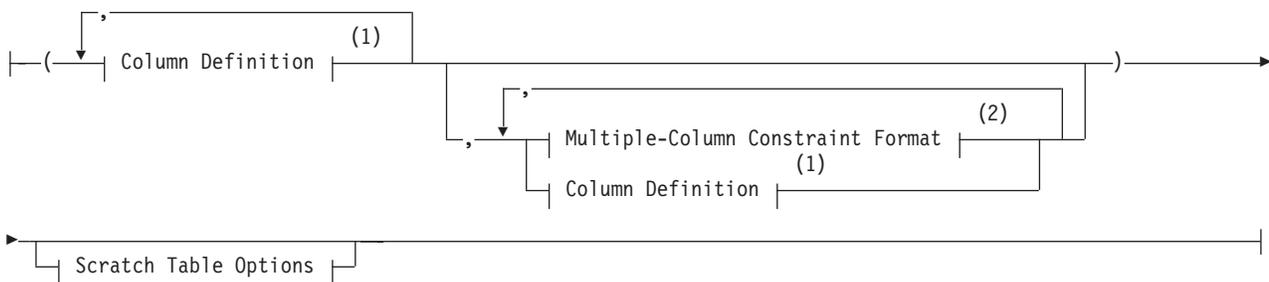
CREATE SCRATCH TABLE

Use the CREATE SCRATCH TABLE statement to create a non-logging temporary table in the current database. Only Extended Parallel Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

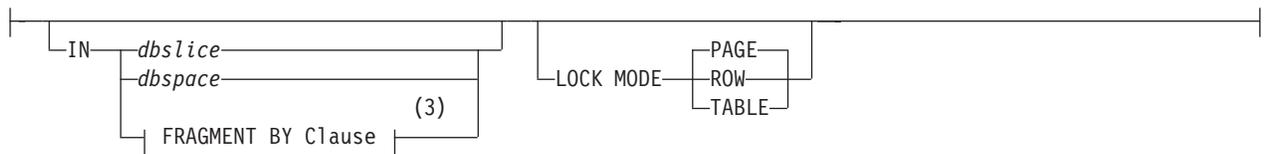
Syntax

►► CREATE SCRATCH TABLE *table* | Scratch Table Definition |

Scratch Table Definition:



Scratch Table Options:



Notes:

- 1 See "Column Definition" on page 2-193
- 2 See "Multiple-Column Constraint Format" on page 2-234
- 3 See "FRAGMENT BY Clause" on page 2-212

Element	Description	Restrictions	Syntax
<i>dbslice</i>	Name of dbslice to store <i>table</i>	Must already exist	"Identifier " on page 5-21
<i>dbspace</i>	Name of dbspace to store <i>table</i> . Default is the dbspace that stores the current database.	Must already exist	"Identifier " on page 5-21
<i>table</i>	Name that you declare here for a nonlogging temporary table	Must be unique in the current session	"Database Object Name " on page 5-16

Examples

```
create scratch table scratch1 (c1 integer default 27 not null );
create scratch table scratch2 (c1 integer not null) in datadbs;
create scratch table scratch3 (c1 integer not null)
  fragment by round robin in rootdbs.1, datadbs;
```

Usage

The CREATE SCRATCH TABLE statement is a special case of the CREATE Temporary TABLE statement. For the complete syntax and semantics of the CREATE SCRATCH TABLE statement, see “CREATE Temporary TABLE ” on page 2-231.

CREATE SECURITY LABEL

Use the CREATE SECURITY LABEL statement to define a new security label for a specified security policy in the current database and to identify its components and the elements of its components.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Element	Description	Restrictions	Syntax
<i>component</i>	A security label component	Must already exist in the database as a component of the specified <i>policy</i> , and be unique among names of components of this <i>label</i> .	“Identifier ” on page 5-21
<i>label</i>	Name you declare here for this label	Must be unique among security label names for this security <i>policy</i>	“Identifier ” on page 5-21
<i>element</i>	An element of the specified <i>component</i>	Must have been defined when its <i>component</i> was defined or was last altered. If <i>component</i> is an array, only a single <i>element</i> can be specified.	“Quoted String” on page 4-189
<i>policy</i>	The security policy of this <i>label</i>	Must already exist in the database	“Identifier ” on page 5-21

Usage

A *security label* is a named database object that supports a specified security policy. A security label can be applied to a user, or to a row or to a column (or to both a row and a column) of a table in the database. When a user who holds a security label attempts to access data that has a security label, the database server takes into account the security label of the column or row and the security label of the user in determining whether to allow the user to access the data.

Every security label stores the following categories of information:

- It identifies an existing security policy that the label supports.

- It identifies at least one, but no more than 16 existing components of the security policy that the label supports.
- It identifies one or more existing elements of each component of the security label. (Only security label components of type SET or TREE can include more than one element in the same security label.)

Only DBSECADM can issue this statement. When the CREATE SECURITY LABEL statement executes successfully, it registers the specified *label* name, the numeric identifier of the associated security *policy*, and the cardinality of its security label *components* in the `sysseclabels` system catalog table.

Components and Elements of a Security Label

Like a security policy, a security label must have at least one component, but no more than 16. The CREATE SECURITY LABEL statement cannot list security label components that are not components of the specified security policy. The same *component* name cannot be specified more than once in the same CREATE SECURITY LABEL statement. These components must already exist in the database, where DBSECADM can register them with the CREATE SECURITY LABEL COMPONENT statement.

Security label components can be of type ARRAY, SET, or TREE, as described in CREATE SECURITY LABEL COMPONENT. For a *component* of type ARRAY, the *element* list can identify only a single element. For components of type SET or TREE, the *element* list can identify multiple component elements that were defined when the component was created (or when it was last altered). See the CREATE SECURITY LABEL COMPONENT statement for more information about the structure and semantics of security label components.

The following example creates a security label called **label1** for a security policy called **MegaCorp**. The label uses two security label components, called **levels** and **compartments**, each with one element, called **VP** and **Marketing** respectively:

```
CREATE SECURITY LABEL MegaCorp.label1
  COMPONENT levels 'VP',
  COMPONENT compartments 'Marketing';
```

For this example to be valid, the **levels** and **compartments** components, and their security label components, **VP** and **Marketing** elements, must have been defined in previously executed CREATE SECURITY LABEL COMPONENT statements.

In the next example, DBSECADM creates a security label called **label2** for the same **MegaCorp** security policy. This label uses three security label components, called **levels**, **compartments**, and **groups**, where two of these components have one element, and another has two:

```
CREATE SECURITY LABEL MegaCorp.label2
  COMPONENT level 'Director',
  COMPONENT compartments 'HR', 'Finance',
  COMPONENT groups 'EntireRegion';
```

These examples illustrate that the components of a security label can be a subset of the components of the security policy that the label supports, and that more than one security label can support the same security policy.

Related Information

Related statements: “ALTER SECURITY LABEL COMPONENT ” on page 2-36, ALTER TABLE, “CREATE SECURITY LABEL COMPONENT ” on page 2-177, “CREATE

SECURITY POLICY " on page 2-181, CREATE TABLE, "DROP SECURITY " on page 2-338, GRANT EXEMPTION, GRANT SECURITY LABEL, "RENAME SECURITY " on page 2-490, REVOKE EXEMPTION, and REVOKE SECURITY LABEL

For a discussion of LBAC security objects, see the Informix Dynamic Server Security Guide.

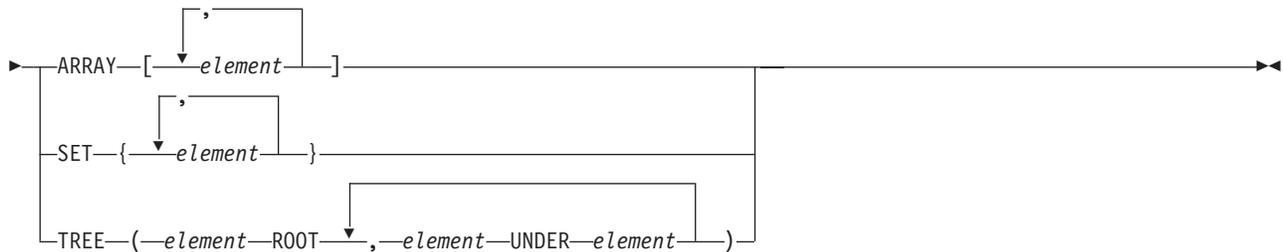
CREATE SECURITY LABEL COMPONENT

Use the CREATE SECURITY LABEL COMPONENT statement to define a new security label component in the current database and to define the elements the component.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

►► CREATE SECURITY LABEL COMPONENT *component* ►►



Element	Description	Restrictions	Syntax
<i>component</i>	Name declared here for this component	Must be unique among the names of security label components in the database.	"Identifier " on page 5-21
<i>element</i>	Component element that is defined here	Must be unique among elements of this <i>component</i> , and no longer than 32 bytes. The left () and right () parentheses, comma (,), and colon (:) symbols are not valid characters.	"Quoted String" on page 4-189

Usage

Only the DBSECADM can issue the CREATE SECURITY LABEL COMPONENT statement, which defines a *security label component*. This is a database object that defines one or more logical categories whose values can be used in a security policy to determine whether a user's request to read or write data is accepted or rejected. The set of all the valid individual values that the security component can have is defined by the set of *security label elements* that this statement specifies for the component.

The logical categories that security label components implement are identified by DBSECADM in the process of designing a *security policy*, which is the core

construct of label-based access control (LBAC). To implement this security feature in the database, however, DBSECADM must create security objects in the following sequence:

1. A set of one or more *security components*, each of which can be defined by the CREATE SECURITY LABEL COMPONENT statement. This statement specifies the name of a security component, the structure of its range of values, and the possible values for this component that can be assigned to a security label that applies a security policy to data or to a user.
2. One or more *security policies*, each of which can be defined by the CREATE SECURITY POLICY statement, which specifies a list of one or more components and a set of rules that the security policy applies to data and to users who attempt read or write operations on data that the security policy protects in the database. A security policy always includes all the elements of a component that CREATE SECURITY POLICY specifies.
3. A *security label* can be defined by the CREATE SECURITY LABEL statement, which specifies one or more values for each of one or more components of the security policy that the label supports. The security label can be applied to data and to users. All the components of a security label must be components of the same security policy, but multiple security policies and multiple security labels can share the same component. A security label typically includes only a subset of the elements of a security component that CREATE SECURITY LABEL COMPONENT defines.

See the GRANT Security and REVOKE Security statements for information on how security labels and exemptions from the rules of a security policy define the LBAC credentials of a user or of a role.

See the CREATE TABLE and ALTER TABLE statements for information on how security labels can be associated with a database table or with an individual data row in a table.

Types and Elements of Security Components

A security label component itself consists of one or more *elements* that the CREATE SECURITY LABEL COMPONENT statement declares as string constants. These elements define the set of values that are valid for the component,

When the CREATE SECURITY LABEL statement executes successfully, Dynamic Server updates the system catalog of the database with the following new entries:

- It creates a new row in the **sysseclabelcomponents** table to register the new component.
- For each *element* of the new component, it creates a new row in the **sysseclabelcomponentelements** table.

The security label component must be defined as one of the three component types. The ARRAY, SET, or TREE keyword that immediately follows the declaration of the *component* name specifies the component type, which must be followed by a list of the *elements* of the security component. These elements define the set of values that the component can have within a security policy. For all three types of security components, the set of elements are under the following restrictions:

- The security component can have no more than 64 elements.
- Each element of a security component is a quoted string constant of no more than 32 bytes.

- Characters in the quoted string constant cannot include the left (() or right ()) parentheses, comma (,), or colon (:) symbols, but other symbols that the DBLOCAL setting supports are valid, including the blank space (ASCII 32) character.
- Each element must be unique among elements of the same security component, but the same quoted string constant value can also be an element of other security components.

The definition of each element within the component implies a level of data sensitivity that a security label associates with a database table or with an individual data row, and also affects the security credentials of users who hold a security label to read or write data that is protected by the same label or by a different label that specifies one or more elements of the component.

Like other database Data Definition Language statements of SQL that can define database objects, CREATE SECURITY LABEL COMPONENT must specify a literal value for each component element, rather than a placeholder. To change the definition of an existing security label component, DBSECADM can use the ALTER SECURITY LABEL COMPONENT to insert a new element into an ARRAY, SET, or TREE component. To drop or rename one or more individual elements of a component, however, DBSECADM must use the DROP SECURITY LABEL COMPONENT statement to destroy the existing component, and then reissue the CREATE SECURITY LABEL COMPONENT statement to create a new component that defines the required set of element values within an ARRAY, SET, or TREE component structure.

ARRAY Components

A security label component of type ARRAY is an ordered set of no more than 64 elements. Each element defines a value that is valid for that component within a security policy. The order in which elements are declared is significant, because it defines a descending order of data sensitivity, with each successive element ranking lower in data sensitivity than the preceding element. The set of label elements of the array and their comma separators must be enclosed between a pair of bracket ([...]) symbols.

When an ARRAY component is specified in the definition of a security label, the label can specify no more than one element of that component as the value of the component.

The following example defines a security label component of type ARRAY called **aquilae** that is an ordered set of five elements called **imperator**, **tribunus**, **centurio**, **miles**, and **asinus**:

```
CREATE SECURITY LABEL COMPONENT aquilae
  ARRAY [ "imperator", "tribunus", "centurio", "miles", "asinus" ];
```

Here the component element with the highest data sensitivity is **imperator** and **asinus** has the lowest data sensitivity, with the data sensitivity of **tribunus** ranking above that of **centurio** but below that of **imperator**.

A component of type ARRAY can be appropriate in contexts where some dimension of a multidimensional security policy can be mapped onto a single scale that is monotonically descending.

SET Components

A security label component of type SET is an unordered set of no more than 64 elements. Each element of the SET is a string constant of no more than 32 bytes, and must be unique within the component, but the same value can be used in other components. The order in which the elements of a SET component are declared is not significant in regard to the data sensitivity of the categories that these elements identify. The elements and their comma separators must be enclosed between a pair of braces ({ ... }) symbols.

When a SET component is specified in the definition of a security label, the label can specify one or multiple elements of that component as valid values for the component.

In the following example, DBSECADM defines a security label component called **departments** that is an unordered set of three elements, called **Marketing**, **HR**, and **finance**:

```
CREATE SECURITY LABEL COMPONENT departments
  SET { 'Marketing', 'HR', 'Finance' };
```

Like all components of type SET, the order in which these elements are declared implies no relative rank in data sensitivity.

A component of type SET can be appropriate in contexts where some dimension of a multidimensional security policy can be represented as nominal categories, without any logical basis for ordering them on a monotonic scale, nor for arranging them in a hierarchy.

TREE Components

A security label component of type TREE has the logical topology of a hierarchy (that is, a simple graph with no loops) that has a single root node and no more than 63 additional nodes. The string constant for the root node must be listed first and must be followed by the ROOT keyword. The string constant for each subsequently declared node must be followed by the keyword UNDER and by the string constant for some previously declared node. The set of elements of the TREE component, including their ROOT and UNDER keywords and comma separators, must be enclosed between a pair of parenthesis ((...)) symbols.

The label element specified after the UNDER keyword is called the *parent* of the label element the precedes the same UNDER keyword (which is called the *child* of that parent element). The CREATE SECURITY LABEL COMPONENT statement fails with an error if a node name that follows the UNDER keyword has not already been declared in the same statement.

The string constant that designates the root node of a tree component has the highest data sensitivity. For a user to read or write protected data, each tree component of the user security label must include at least one of the elements in the tree component of the data row security label, or the ancestor of one such element. For example, if "Beta" is declared UNDER "Alpha" and "Gamma" is declared UNDER "Beta" then "Gamma" also ranks below "Alpha" in data sensitivity. Only elements that are in the same chain of parent-child relationships can be compared in their data sensitivity.

The next example defines a security label component called **Oakland** as a tree structure with six nodes:

```
CREATE SECURITY LABEL COMPONENT Oakland
TREE ( 'Port' ROOT,
      'Downtown' UNDER 'Port',
      'Airport' UNDER 'Port',
      'Estuary' UNDER 'Airport',
      'Avenues' UNDER 'Downtown',
      'Hills' UNDER 'Avenues');
```

Here the root node is **Port**, which has the highest data sensitivity. Within this hierarchy, the **Downtown**, **Avenues**, and **Hills** elements represent descending levels of data sensitivity, and the **Airport** element has a higher data sensitivity than the **Estuary** element. In this example, the four component elements that the UNDER keyword designates as parent nodes are each declared before being included in UNDER specifications. A modified version of this example would also be valid if the **Avenues** node declaration preceded the **Airport** node declaration, but an error would result if the **Hills** node declaration had preceded the **Avenues** node declaration.

A component of type TREE can be appropriate in contexts where some dimension of a multidimensional security policy can be mapped to a single logical hierarchy, or to a group of hierarchies that share a common root.

Related Information

Related statements: “ALTER SECURITY LABEL COMPONENT ” on page 2-36, ALTER TABLE, “CREATE SECURITY LABEL ” on page 2-175, “CREATE SECURITY POLICY ,” CREATE TABLE, “DROP SECURITY ” on page 2-338, GRANT EXEMPTION, GRANT SECURITY LABEL, “RENAME SECURITY ” on page 2-490, REVOKE EXEMPTION, and REVOKE SECURITY LABEL

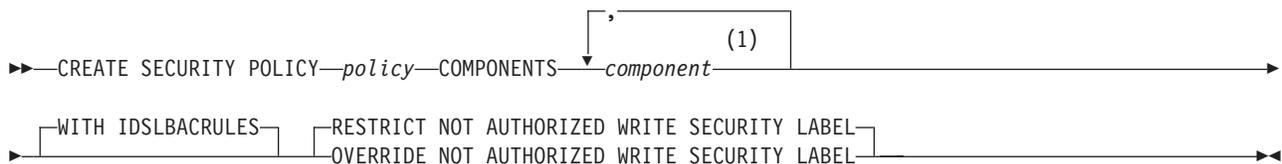
For a discussion of LBAC security objects, see the Informix Dynamic Server Security Guide.

CREATE SECURITY POLICY

Use the CREATE SECURITY POLICY statement to define a new security policy in the current database and to identify its security label components and access rules.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 You can specify no more than 16 components.

Element	Description	Restrictions	Syntax
<i>component</i>	A security label component	Must already exist in the database, and be unique among the names of components for this <i>policy</i>	"Identifier " on page 5-21
<i>policy</i>	Name declared here for a security policy	Must be unique among the names of security policies in the database	"Identifier " on page 5-21

Usage

A *security policy* is a named database object that stores the following information:

- It defines a set of security label components that comprise a security label.
- It associates that security label with a set of access rules.

For tables that are protected by a security policy, the access rules enable Dynamic Server to compare the security credentials of a user with the security label of a row or column. The security policy is applied to determine whether a user who holds a given security label can read or write data in a row or column that is labeled with a security label. A security policy has no effect on data that has no security label.

No more than one security policy can be attached to a table at any point in time, and a security policy can include no more than 16 security label components.

Only DBSECADM can issue this statement. When the CREATE SECURITY POLICY statement executes successfully, Dynamic Server makes the following updates to the system catalog of the current database:

- Registers the specified *policy* name and the cardinality of its security label components in the **syssecpolicies** table
- Creates for each *component* a new row in **syssecpolicycomponentrules** table.

Security Label Components of a Security Policy

The CREATE SECURITY POLICY statement must specify at least one (but no more than 16) security label components. These components must already exist in the database, where DBSECADM can register them with the CREATE SECURITY LABEL COMPONENT statement. The same *component* name cannot be specified more than once in the same CREATE SECURITY POLICY statement.

See the section CREATE SECURITY LABEL COMPONENT for more information about the structure and semantics of security label components.

Rules Associated with a Security Policy

The WITH IDSLBACRULES keywords specify the read access rules and write access rules that the new security policy enforces. If you do not specify them, these keywords are in effect by default, because the IDSLBACRULES access rules are the only access rules that the a security policy can support.

The following IDSLBACRULES access rules for read access, called IDSLBACREAD, apply when data values are read from labeled rows or columns in SELECT, UPDATE, or DELETE operations:

- **IDSLBACREADARRAY**: Each array component of the user security label must be greater than or equal to the array component of the data row security label. That is, only data at or below the level of the user can be read.

- **IDSLBACREADTREE**: Each tree component of the user security label must include at least one of the elements in the tree component of the data row security label (or the ancestor of one such element).
- **IDSLBACREADSET**: Each set component of the user security label must include the set component of the data row security label.

The following **IDSLBACRULES** access rules for write access, called **IDSLBACWRITE**, apply when data values are written to labeled rows or columns in INSERT, UPDATE, or DELETE operations:

- **IDSLBACWRITEARRAY**: Each array component of the user security label must be equal to the array component of the data row security label. That is, only data at the same level as the user can be written.
- **IDSLBACWRITETREE**: Each tree component of the user security label must include at least one of the elements in the tree component of the data row security label (or the ancestor of one such element).
- **IDSLBACWRITESET**: Each set component of the user security label must include the set component of the data row security label.

If DBSECADM omits the WITH IDSLBACRULES keywords, then those rules are in effect by default. If any specification except IDSLBACRULES follows the WITH keyword, however, the CREATE SECURITY POLICY statement fails with an error, and no security policy is created.

Besides the explicit or default WITH IDSLBACRULES keywords, the CREATE SECURITY POLICY statement must also specify the write access rule to enforce when a user is not authorized to write the explicitly specified security label provided in the DELETE, INSERT, or UPDATE statement for a table protected with this security policy. The security label of the user and the exemption credentials that the user holds determine whether the user has write access to an explicitly provided security label.

- If the CREATE SECURITY POLICY statement specifies **OVERRIDE NOT AUTHORIZED WRITE SECURITY LABEL**, then Dynamic Server uses the value of the user security label, rather than the security label that is explicitly specified in the DELETE, INSERT, or UPDATE statement, to determine whether the user has write-access to data values that are protected by a security label in the DELETE, INSERT, or UPDATE operation.
- The default is **RESTRICT NOT AUTHORIZED WRITE SECURITY LABEL**. If you specify these keywords explicitly, or if they are in effect by default, then DELETE, INSERT, or UPDATE statements fail with an error if the user is not authorized to write data in a row or column that has the explicitly specified security label.

The following example creates a security policy called **MegaCorp** that uses three security label components, with no **OVERRIDE** provision for the user security label to provide write access in DELETE, INSERT, or UPDATE operations on data whose explicitly specified security label does not authorize write access for that user:

```
CREATE SECURITY POLICY MegaCorp
  COMPONENTS levels, compartments, groups
  WITH IDSLBACRULES;
```

For this example to be valid, the **levels**, **compartments**, and **groups** security label components (or components that have been renamed to these identifiers) must have been previously defined by the CREATE SECURITY LABEL COMPONENT statement.

Related Information

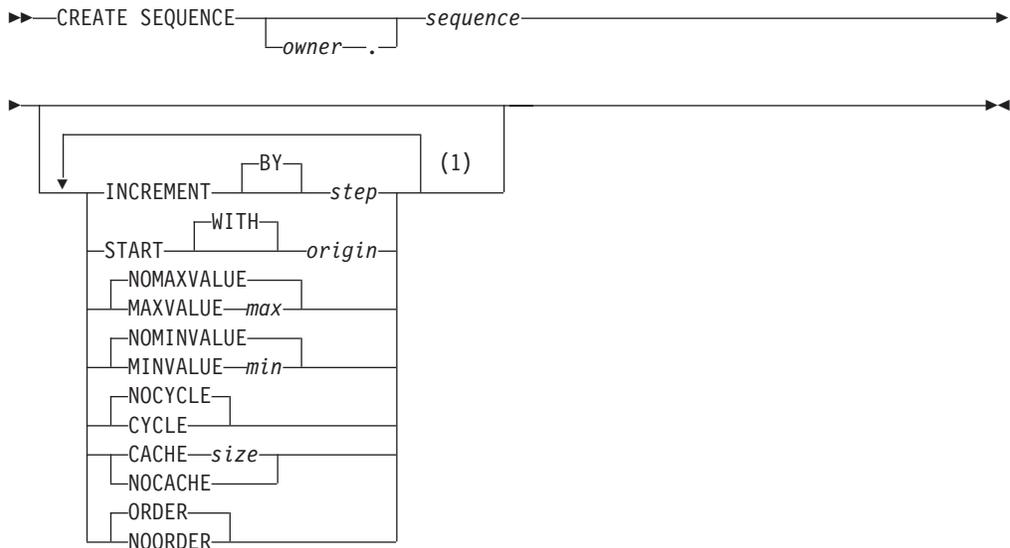
Related statements: “ALTER SECURITY LABEL COMPONENT ” on page 2-36, ALTER TABLE, “CREATE SECURITY LABEL ” on page 2-175, “CREATE SECURITY LABEL COMPONENT ” on page 2-177, CREATE TABLE, “DROP SECURITY ” on page 2-338, GRANT EXEMPTION, GRANT SECURITY LABEL, “RENAME SECURITY ” on page 2-490, REVOKE EXEMPTION, and REVOKE SECURITY LABEL

For a discussion of LBAC security objects, see the Informix Dynamic Server Security Guide.

CREATE SEQUENCE

Use the CREATE SEQUENCE statement to create a sequence database object from which multiple users can generate unique integers. Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 Each keyword option can appear no more than once.

Element	Description	Restrictions	Syntax
<i>max</i>	Upper limit of values	Must be an integer > <i>origin</i>	“Literal Number ” on page 4-185
<i>min</i>	Lower limit of values	Must be an integer less than <i>origin</i>	“Literal Number ” on page 4-185
<i>origin</i>	First number in the sequence	Must be an integer in INT8 range	“Literal Number ” on page 4-185
<i>owner</i>	Owner of <i>sequence</i>	Must be an authorization identifier	“ Owner Name ” on page 5-45
<i>sequence</i>	Name that you declare here for the new sequence	Must be unique among sequence, table, view, and synonym names	“Identifier ” on page 5-21

Element	Description	Restrictions	Syntax
<i>size</i>	Number of values that are preallocated in memory	Integer > 1, but < cardinality of a cycle (= $\lceil (max - min) / step \rceil$)	"Literal Number " on page 4-185
<i>step</i>	Interval between successive values	Nonzero integer in INT range	"Literal Number " on page 4-185

Usage

A sequence (sometimes called a *sequence generator* or *sequence object*) returns a monotonically ascending or descending series of unique integers, one at a time. The CREATE SEQUENCE statement defines a new sequence object, declares its identifier, and registers it in the **syssequences** system catalog table.

Authorized users of a sequence can request a new value by including the *sequence.NEXTVAL* expression in DML statements. The *sequence.CURRVAL* expression returns the current value of the specified *sequence*. **NEXTVAL** and **CURRVAL** expressions are valid only within SELECT, DELETE, INSERT, and UPDATE statements; Dynamic Server returns an error if you attempt to invoke the built-in **NEXTVAL** or **CURRVAL** functions in any other context.

Generated values logically resemble the SERIAL8 data type, but can be negative, and are unique within the sequence. Because the database server generates the values, sequences support a much higher level of concurrency than a serial column can. The values are independent of transactions; a generated value cannot be rolled back, even if the transaction in which it was generated fails.

You can use a sequence to generate primary key values automatically, using one sequence for many tables, or each table can have its own sequence.

CREATE SEQUENCE can specify the following characteristics of a sequence:

- Initial value
- Size and sign of the increment between values.
- Maximum and minimum values
- Whether the sequence recycles values after reaching its limit
- How many values are preallocated in memory for rapid access

A database can support multiple sequences concurrently, but the name of a sequence (or in an ANSI-compliant database, the *owner.sequence* combination) must be unique within the current database among the names of tables, temporary tables, views, synonyms, and sequences.

An error occurs if you include contradictory options, such as specifying both the MINVALUE and NOMINVALUE options, or both CACHE and NOCACHE.

INCREMENT BY Option

Use the INCREMENT BY option to specify the interval between successive numbers in the sequence. The BY keyword is optional. The interval, or *step* value, can be a positive whole number (for an *ascending* sequence) or a negative whole number (for a *descending* sequence) in the INT8 range. If you do not specify any *step* value, the default interval between successive generated values is 1, and the sequence is an ascending sequence.

START WITH Option

Use the START WITH option to specify the first number of the sequence. This *origin* value must be an integer within the INT8 range that is greater than or equal to the *min* value (for an ascending sequence) or that is less than or equal to the *max* value (for a descending sequence), if *min* or *max* is specified in the CREATE SEQUENCE statement. The WITH keyword is optional.

If you do not specify an *origin* value, the default initial value is *min* for an ascending sequence or *max* for a descending sequence. (The “MAXVALUE or NOMAXVALUE Option” and “MINVALUE or NOMINVALUE Option” sections that follow describe the *max* and *min* specifications respectively.)

MAXVALUE or NOMAXVALUE Option

Use the MAXVALUE option to specify the upper limit of values in a sequence. The maximum value, or *max*, must be an integer in the INT8 range that is greater than the value of the *origin*.

If you do not specify a *max* value, the default is NOMAXVALUE. This default setting supports values that are less than or equal to 2e64 for ascending sequences, or less than or equal to -1 for descending sequences.

MINVALUE or NOMINVALUE Option

Use the MINVALUE option to specify the lower limit of values, or *min*. This integer must be in the INT8 range and be less than the value of *origin*.

If you do not specify a *min* value, the default is NOMINVALUE. This default setting supports values that are greater than or equal to 1 for ascending sequences, or greater than or equal to -(2e64) for descending sequences.

CYCLE or NOCYCLE Option

Use the CYCLE option to continue generating sequence values after the sequence reaches the maximum (ascending) or minimum (descending) limit. After an ascending sequence reaches the *max* value, it generates the *min* value for the next sequence value. After a descending sequence reaches the *min* value, it generates the *max* value for the next sequence value.

The default is NOCYCLE. At this default setting, the sequence cannot generate more values after reaching the declared limit. Once the sequence reaches the limit, the next reference to *sequence.NEXTVAL* returns an error.

CACHE or NOCACHE Option

Use the CACHE option to specify the number of sequence values that are preallocated in memory for rapid access. This feature can enhance the performance of a heavily used sequence. The cache *size* must be a positive whole number in the INT range. If you specify the CYCLE option, then *size* must be less than the number of values in a cycle (or less than $\lfloor (max - min) / step \rfloor$). The minimum is 2 preallocated values. The default is 20 preallocated values.

The NOCACHE keyword specifies that no generated values (that is, zero) are preallocated in memory for this sequence object.

The configuration parameter `SEQ_CACHE_SIZE` specifies the maximum number of sequence objects that can have preallocated values in the sequence cache. If this configuration parameter is not set, then by default no more than 10 different sequence objects can be defined with the `CACHE` option.

ORDER or NOORDER Option

These keywords have no effect on the behavior of the sequence. The sequence always issues values to users in the order of their requests, as if the `ORDER` keyword were always specified. The `ORDER` and `NOORDER` keywords are accepted by the `CREATE SEQUENCE` statement for compatibility with implementations of sequence objects in other dialects of SQL.

Examples

```
CREATE SEQUENCE seq_2
  INCREMENT BY 1 START WITH 1
  MAXVALUE 30 MINVALUE 0
  NOCYCLE CACHE 10 ORDER;

CREATE TABLE tab1 (col1 int, col2 int);
INSERT INTO tab1 VALUES (0, 0);

INSERT INTO tab1 (col1, col2) VALUES (seq_2.NEXTVAL, seq_2.NEXTVAL)

SELECT * FROM tab1;
```

col1	col2
0	0
1	1

Related Information

Related statements: “`ALTER SEQUENCE`” on page 2-40, “`DROP SEQUENCE`” on page 2-340, “`RENAME SEQUENCE`” on page 2-491, “`CREATE SYNONYM`,” “`DROP SYNONYM`” on page 2-341, “`GRANT`” on page 2-395, and “`REVOKE`” on page 2-494

For information about the `syssequences` system catalog table in which sequence objects are registered, see the *IBM Informix Guide to SQL: Reference*.

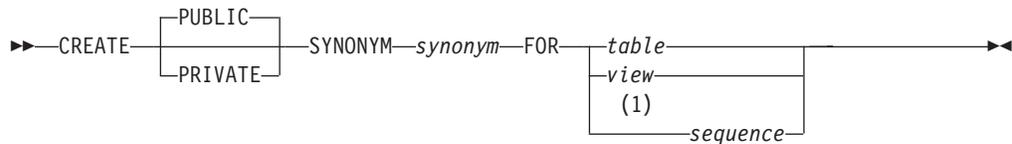
For information about initializing a sequence and generating or reading values from a sequence, see “`NEXTVAL` and `CURRVAL` Operators (IDS)” on page 4-79.

CREATE SYNONYM

Use the `CREATE SYNONYM` statement to declare and register an alternative name for an existing table, view, or sequence object.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>sequence</i>	Name of a local sequence	Must exist in the current database	"Identifier " on page 5-21
<i>table, view</i>	Name of table or view for which <i>synonym</i> is being created	Must exist in the current database, or in a database specified in a qualifier	"Database Object Name " on page 5-16
<i>synonym</i>	Synonym declared here for the name of a table, view, or sequence	Must be unique among table object names; see also Usage notes	"Database Object Name " on page 5-16

Usage

Users have the same privileges for a synonym that they have for the database object that the synonym references. The **syssynonyms**, **syssyntable**, and **systables** system catalog tables maintain information about synonyms.

You cannot create a synonym for a synonym in the same database.

The identifier of the synonym must be unique among the names of tables, temporary tables, views, and sequence objects in the same database. (See, however, the section "Synonyms with the Same Name" on page 2-189.)

Once a synonym is created, it persists until the owner executes the DROP SYNONYM statement. (This persistence distinguishes a synonym from an alias that you can declare in the FROM clause of a SELECT statement; the alias is in scope only during execution of that SELECT statement.) If a synonym refers to a table, view, or sequence in the same database, the synonym is automatically dropped if the referenced table, view, or sequence object is dropped.

Synonyms for Remote and External Tables and Views

A synonym can be created for any table or view in any database on your database server. This example declares a synonym for a table outside your current database, in the **payables** database of your current database server.

```
CREATE SYNONYM mysum FOR payables:jean.summary
```

You can also create a synonym for a table or view that exists in a database of a database server that is not your current database server. Both database servers must be online when you create the synonym. In a network, the remote database server verifies that the table or view referenced by the synonym exists when you create the synonym. The next example creates a synonym for a table supported by a remote database server:

```
CREATE SYNONYM mysum FOR payables@phoenix:jean.summary
```

The identifier **mysum** now refers to the table **jean.summary**, which is in the **payables** database on the **phoenix** database server. If the **summary** table is dropped from the **payables** database, the **mysum** synonym is left intact. Subsequent attempts to use **mysum** return the error: Table not found.

Dynamic Server, however, does not support synonyms for these external objects:

- Typed tables (including any table that is part of a table hierarchy)
- Tables or views that contain any extended data types
- Sequence objects outside the local database

PUBLIC and PRIVATE Synonyms

If you use the PUBLIC keyword (or no keyword at all), anyone who has access to the database can use your synonym. If the database is not ANSI-compliant, a user does not need to know the name of the owner of a public synonym. Any synonym in a database that is not ANSI-compliant *and* was created in an Informix database server earlier than Version 5.0 is a public synonym.

In an ANSI-compliant database, all synonyms are private. If you use the PUBLIC or PRIVATE keywords, the database server issues a syntax error.

If you use the PRIVATE keyword to declare a synonym in a database that is not ANSI-compliant, the unqualified synonym can be used by its owner. Other users must qualify the synonym with the name of the owner.

Synonyms with the Same Name

In an ANSI-compliant database, the *owner.synonym* combination must be unique among all synonyms, tables, views, and sequences. You must specify *owner* when you refer to a synonym that you do not own, as in this example:

```
CREATE SYNONYM emp FOR accting.employee
```

In a database that is not ANSI-compliant, no two public synonyms can have the same identifier, and the identifier of a synonym must also be unique among the names of tables, views, and sequences in the same database.

The *owner.synonym* combination of a private synonym must be unique among all the synonyms in the database. That is, more than one private synonym with the same name can exist in the same database, but a different user must own each of these synonyms. The same user cannot create both a private and a public synonym that have the same name. For example, the following code generates an error:

```
CREATE SYNONYM our_custs FOR customer;  
CREATE PRIVATE SYNONYM our_custs FOR cust_calls;-- ERROR!!!
```

A private synonym can be declared with the same name as a public synonym only if the two synonyms have different owners. If you own a private synonym, and a public synonym exists with the same name, the database server resolves the unqualified name as the private synonym. (In this case, you must specify *owner.synonym* to reference the public synonym.) If you use DROP SYNONYM with the unqualified synonym identifier when your private synonym and the public synonym of another user both have the same identifier, only your private synonym is dropped. If you repeat the same DROP SYNONYM statement, the database server drops the public synonym.

Chaining Synonyms

If you create a synonym for a table or view that is not in the current database, and this table or view is dropped, the synonym stays in place. You can create a new synonym for the dropped table or view with the name of the dropped table or view as the synonym, which points to another external or remote table or view. (Synonyms for external sequence objects are not supported.)

In this way, you can move a table or view to a new location and chain synonyms together so that the original synonyms remain valid. (You can chain up to 16 synonyms in this manner.)

The following steps chain two synonyms together for the **customer** table, which will ultimately reside on the **zoo** database server (the CREATE TABLE statements are not complete):

1. In the **stores_demo** database on the database server that is called **training**, issue the following statement:

```
CREATE TABLE customer (lname CHAR(15)...) 
```

2. On the database server called **acctng**, issue the following statement:

```
CREATE SYNONYM cust FOR stores_demo@training:customer 
```

3. On the database server called **zoo**, issue the following statement:

```
CREATE TABLE customer (lname CHAR(15)...) 
```

4. On the database server called **training**, issue the following statement:

```
DROP TABLE customer  
CREATE SYNONYM customer FOR stores_demo@zoo:customer 
```

The synonym **cust** on the **acctng** database server now points to the **customer** table on the **zoo** database server.

The following steps show an example of chaining two synonyms together and changing the table to which a synonym points:

1. On the database server called **training**, issue the following statement:

```
CREATE TABLE customer (lname CHAR(15)...) 
```

2. On the database server called **acctng**, issue the following statement:

```
CREATE SYNONYM cust FOR stores_demo@training:customer 
```

3. On the database server called **training**, issue the following statement:

```
DROP TABLE customer  
CREATE TABLE customer (lastname CHAR(20)...) 
```

The synonym **cust** on the **acctng** database server now points to a new version of the **customer** table on the **training** database server.

Related Information

Related statement: "DROP SYNONYM " on page 2-341.

For a discussion of concepts related to synonyms, see the *IBM Informix Database Design and Implementation Guide*.

CREATE TABLE

Use the CREATE TABLE statement to create a new permanent table in the current database, to place data-integrity constraints on columns, to designate where the table is stored, to define a fragmentation strategy, to indicate the size of its initial and subsequent extents, and to specify how to lock the new table.

You can use the CREATE TABLE statement to create relational-database tables or to create typed tables (object-relational tables). For information on how to create temporary tables, see “CREATE Temporary TABLE ” on page 2-231.

Syntax

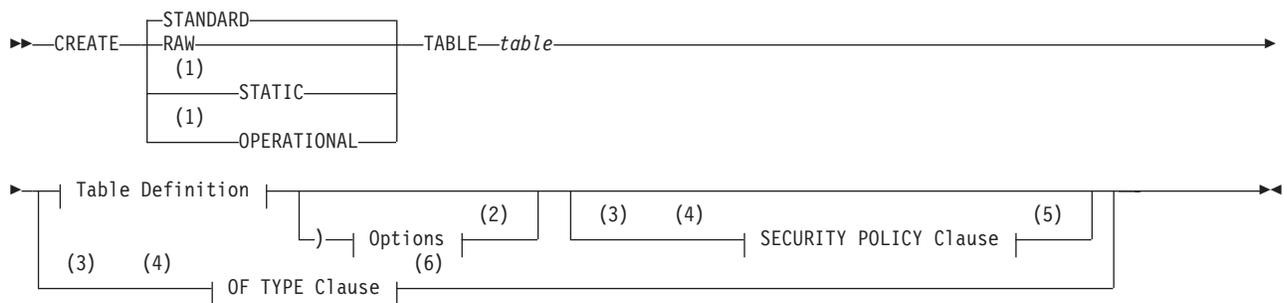
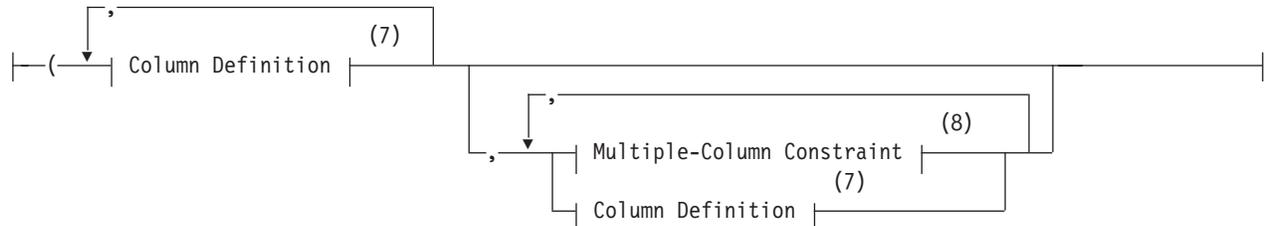


Table Definition:



Notes:

- 1 Extended Parallel Server only
- 2 See “Options Clauses” on page 2-209
- 3 Informix extension
- 4 Dynamic Server only
- 5 See “The SECURITY POLICY Clause (IDS)” on page 2-229
- 6 See “OF TYPE Clause (IDS)” on page 2-227
- 7 See “Column Definition” on page 2-193
- 8 See “Multiple-Column Constraint Format” on page 2-205

Element	Description	Restrictions	Syntax
<i>table</i>	Name that you declare here for the new table	Must be unique among the names of tables, synonyms, views, and sequences in the database	“Identifier ” on page 5-21

Usage

The CREATE TABLE statement can include various clauses, some of which are identified in the following list. Clauses whose names are followed by "(IDS)" are not supported by Extended Parallel Server. Similarly, clauses whose names are followed by "(XPS)" are not supported by Dynamic Server.

Clause	Page	What the Clause Specifies
Logging Options	"Logging Options" on page 2-193	Logging characteristics of the new table
Column Definition	"Column Definition" on page 2-193	Name and other attributes of an individual column
DEFAULT	"DEFAULT Clause" on page 2-195	Default value for an individual column
Single-Column Constraint	"Single-Column Constraint Format" on page 2-197	Data-integrity constraints on an individual column
REFERENCES	"REFERENCES Clause" on page 2-199	Referential-integrity constraints with other columns
CHECK	"CHECK Clause" on page 2-202	Check constraints with other columns
Constraint Definition (IDS)	"Constraint Definition" on page 2-203	Name and other attributes of a constraint
Multiple-Column Constraint	"Multiple-Column Constraint Format" on page 2-205	Data-integrity constraints on a set of columns
WITH CRCOLS (IDS)	"Using the WITH CRCOLS Option (IDS)" on page 2-209	Two shadow columns for Enterprise Replication
Storage Options	"Storage Options" on page 2-210	Attributes of where the table is physically stored
IN	"Using the IN Clause" on page 2-211	Storage object to hold the new table (or part of it)
FRAGMENT BY <i>or</i> PARTITION BY	"FRAGMENT BY Clause" on page 2-212	Distribution scheme of a fragmented table
WITH ROWIDS (IDS)	"Using the WITH ROWIDS Option (IDS)" on page 2-214	Hidden column in a fragmented table
RANGE METHOD (XPS)	"RANGE Method Clause (XPS)" on page 2-218	Fragmentation strategy based on column values
PUT (IDS)	"PUT Clause (IDS)" on page 2-222	Storage location for BLOB or CLOB column values
EXTENT SIZE	"EXTENT SIZE Options" on page 2-224	Size of the first and subsequent extents
USING ACCESS METHOD (IDS)	"USING Access-Method Clause (IDS)" on page 2-224	How to access the new table
LOCK MODE	"LOCK MODE Options" on page 2-225	Locking granularity of the new table
OF TYPE (IDS)	"OF TYPE Clause (IDS)" on page 2-227	Named ROW type of a new typed table
UNDER (IDS)	"Using the UNDER Clause" on page 2-228	Supertable of a new subtable in a table hierarchy

When you create a new table, every column must have a data type associated with it. The *table* name must be unique among all the names of tables, views, sequences, and synonyms within the same database, but the names of columns need only be unique among the column names of the same table.

In an ANSI-compliant database, the combination *owner.table* must be unique among tables, synonyms, views, and sequence objects within the database.

In DB–Access, using CREATE TABLE outside the CREATE SCHEMA statement generates warnings if you use the **-ansi** flag or if you set **DBANSIWARN**.

In Informix ESQL/C, using CREATE TABLE generates warnings if you compile with the **-ansi** flag or set **DBANSIWARN**. For information about the **DBANSIWARN** environment variable, see the *IBM Informix Guide to SQL: Reference*.

Logging Options

Use the Logging Type options to specify logging characteristics that can improve performance in various bulk operations on the table. Other than the default option (STANDARD) that is used for OLTP databases, these logging options are used primarily to improve performance in data warehousing databases.

A table can have either of the following logging characteristics.

Logging Type Effect

STANDARD

Logging tables that allow rollback, recovery, and restoration from archives. This type is the default. Use this type of table for all the recovery and constraints functionality that OLTP databases require.

RAW Nonlogging tables that do not support primary key constraints or unique constraints, but that support referential constraints, and can be indexed and updated. Use this type of table for quickly loading data.

With Extended Parallel Server, you can take advantage of light appends and avoid the overhead of logging and checking constraints by using raw tables.

Warning: Use raw tables for fast loading of data, but set the logging type to STANDARD and perform a level-0 backup before you use the table in a transaction or modify the data within the table. If you must use a raw table within a transaction, either set the isolation level to Repeatable Read or lock the table in exclusive mode to prevent concurrency problems.

Extended Parallel Server supports two additional logging type options.

Option Effect

OPERATIONAL

Logging table that uses light appends; it cannot be restored from archive. Use this type on tables that are refreshed frequently, because light appends allow the quick addition of many rows.

STATIC

Nonlogging table that can contain index and referential constraints but cannot be updated. Use this type for read-only operations, because no logging or locking overhead occurs.

For more information on these logging types of tables, refer to your *IBM Informix Administrator's Guide*.

Column Definition

Use the column definition portion of CREATE TABLE to list the name, data type, default values, and constraints of a *single column*.

As with any SQL identifier, syntactic ambiguities (and sometimes error messages or unexpected behavior) can occur if the *column* name is a keyword, or if it is the same as the *table* name, or the name of another table that you subsequently join with the *table*). For information about the keywords of Dynamic Server, see Appendix A, "Reserved Words for IBM Informix Dynamic Server," on page A-1.

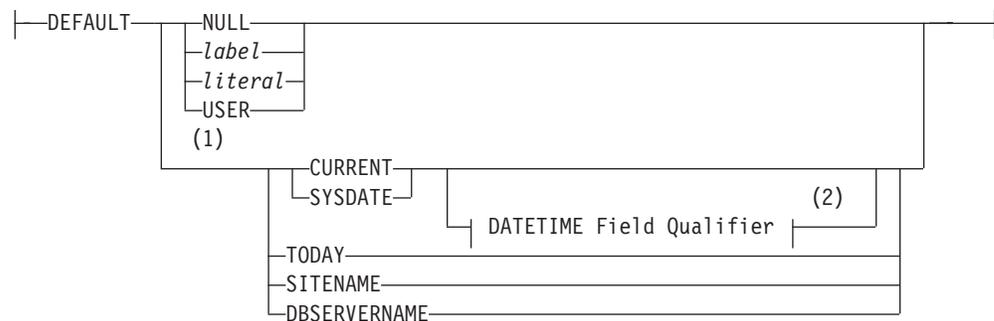
For more information on the keywords of Extended Parallel Server, see Appendix B, "Reserved Words for IBM Informix Extended Parallel Server," on page B-1. For more information on the ambiguities that can occur, see "Use of Keywords as Identifiers" on page 5-23.

In Dynamic Server, if you define a column of a table to be of a named ROW type, the table does not adopt any constraints of the named ROW.

DEFAULT Clause

Use the DEFAULT clause to specify the default value for the database server to insert into a column when no explicit value for the column is specified. You cannot specify default values for SERIAL, BIGSERIAL, or SERIAL8 columns.

DEFAULT Clause:



Notes:

- 1 Informix extension
- 2 See " DATETIME Field Qualifier" on page 4-38

Element	Description	Restrictions	Syntax
<i>label</i>	Name of a security label	Must exist and must belong to the security policy that protects the table. The column must be of type IDSSECURITYLABEL.	"Identifier " on page 5-21
<i>literal</i>	String of alphabetic or numeric characters	Must be an appropriate data type for the column. See "Using a Literal as a Default Value" on page 2-196.	"Expression" on page 4-40

Using NULL as a Default Value

If you specify no default value for a column, the default is NULL unless you place a NOT NULL constraint on the column. In this case, no default exists.

If you specify NULL as the default value for a column, you cannot specify a NOT NULL constraint as part of the column definition. (For details of NOT NULL constraints, see "Using the NOT NULL Constraint" on page 2-198.)

NULL is not a valid default value for a column that is part of a primary key.

If the column is a BYTE or TEXT data type, NULL is the only valid default value.

In Dynamic Server, if the column is a BLOB or CLOB data type, NULL is the only valid default value.

Using a Literal as a Default Value

You can designate a *literal* value as a default value. A literal value is a string of alphabetic or numeric characters. To use a literal value as a default value, you must adhere to the syntax restrictions in the following table.

For Columns of Data Type	Format of Default Value
BOOLEAN	Use 't' or 'f' (respectively for <i>true</i> or <i>false</i>) as a "Quoted String" on page 4-189.
CHAR, CHARACTER VARYING, DATE, VARCHAR, NCHAR, NVARCHAR, LVARCHAR	"Quoted String" on page 4-189. See note that follows for DATE.
DATETIME	"Literal DATETIME " on page 4-181
BIGINT, DECIMAL, FLOAT, INT8, INTEGER, MONEY, SMALLFLOAT, SMALLINT	"Literal Number " on page 4-185
INTERVAL	"Literal INTERVAL" on page 4-183
Opaque data types (IDS)	"Quoted String" on page 4-189 in format of "Single-Column Constraint Format" on page 2-197

DATE literals must be of the format that the **DBDATE** (or else **GL_DATE**) environment variable specifies. In the default locale, if neither **DBDATE** nor **GL_DATE** is set, date literals must be of the *mm/dd/yyyy* format.

Using a Built-in Function as a Default Value

You can specify a built-in function as the default column value. The following table lists built-in functions and operators that you can specify, the data type requirements, and the recommended size (in bytes) for their corresponding columns.

Built-In Function	Data Type Requirement	Recommended Size
CURRENT, SYSDATE	DATETIME column with matching qualifier	Enough bytes to store the longest DATETIME value for the locale
DBSERVERNAME, SITENAME	CHAR, VARCHAR, NCHAR, NVARCHAR, or CHARACTER VARYING column	128 bytes
TODAY	DATE column	Enough bytes to store the longest DATE value for the locale
USER	CHAR, VARCHAR, NCHAR, NVARCHAR, or CHARACTER VARYING column	32 bytes

These column sizes are recommended because, if the column length is too small to store the default value during INSERT or ALTER TABLE operations, the database server returns an error.

In Dynamic Server, you cannot designate a built-in variant function (that is, CURRENT, SYSDATE, USER, TODAY, SITENAME, or DBSERVERNAME) as the default value for a column that holds an opaque or distinct data type. In addition, larger column sizes are required if the data values are encrypted, or if they are encoded in the Unicode character set of the **UTF-8** locale. (See the description of

the SET ENCRYPTION statement later in this chapter for more information about storage size requirements for encrypted data.)

For descriptions of these functions, see “Constant Expressions” on page 4-70.

The following example creates a table called **accounts**. In **accounts**, the **acc_num**, **acc_type**, and **acc_descr** columns have literal default values. The **acc_id** column defaults to the login name of the user.

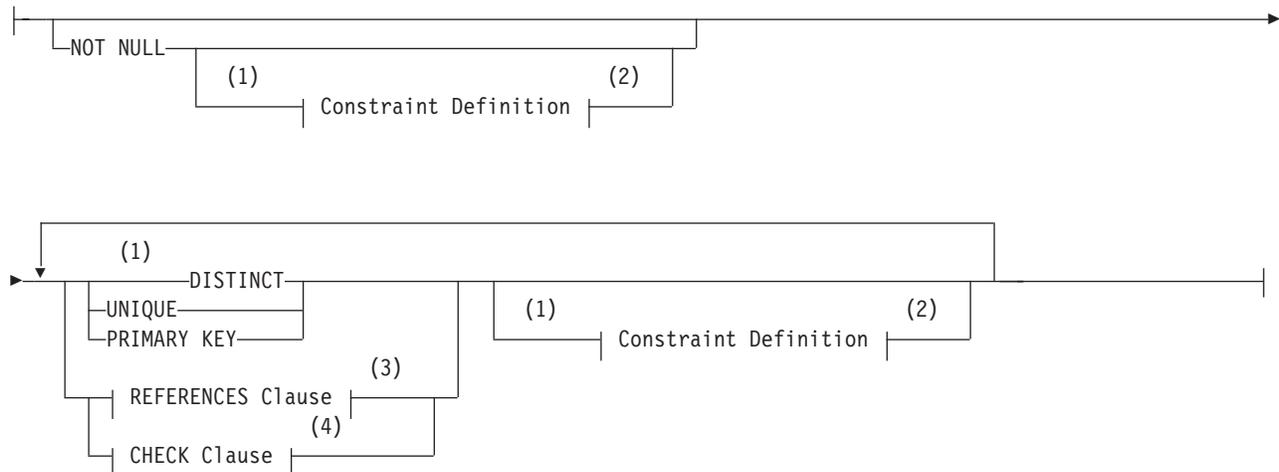
```
CREATE TABLE accounts (
  acc_num INTEGER DEFAULT 1,
  acc_type CHAR(1) DEFAULT 'A',
  acc_descr CHAR(20) DEFAULT 'New Account',
  acc_id CHAR(32) DEFAULT USER)
```

Single-Column Constraint Format

Use the Single-Column Constraint format to associate one or more constraints with a column, in order to perform any of the following tasks:

- Create one or more data-integrity constraints for a column.
- Specify a meaningful name for a constraint.
- Specify the constraint-mode that controls the behavior of a constraint during insert, delete, and update operations.

Single-Column Constraint Format:



Notes:

- 1 Informix extension
- 2 See “Constraint Definition” on page 2-203
- 3 See “REFERENCES Clause” on page 2-199
- 4 See “CHECK Clause” on page 2-202

The following example creates a standard table with two constraints: **num**, a primary-key constraint on the **acc_num** column; and **code**, a unique constraint on the **acc_code** column:

```
CREATE TABLE accounts (
  acc_num INTEGER PRIMARY KEY CONSTRAINT num,
  acc_code INTEGER UNIQUE CONSTRAINT code,
  acc_descr CHAR(30))
```

The types of constraints used in this example are defined in sections that follow.

Restrictions on Using the Single-Column Constraint Format

The single-column constraint format cannot specify a constraint that involves more than one column. Thus, you cannot use the single-column constraint format to define a composite key. For information on multiple-column constraints, see “Multiple-Column Constraint Format” on page 2-205.

You cannot place unique, primary-key, or referential constraints on BYTE or TEXT columns. You can, however, check for NULL or non-NULL values on BYTE or TEXT columns with a check constraint.

You cannot place unique constraints, primary-key constraints, or referential constraints on BLOB or CLOB columns of Dynamic Server. If the constraint is on a column that stores encrypted data, Dynamic Server cannot enforce the constraint.

Using the NOT NULL Constraint

Use the NOT NULL keywords to require that a column receive a value during insert or update operations. If you place a NOT NULL constraint on a column (and no default value is specified), you *must* enter a value into this column when you insert a row or update that column in a row. If you do not enter a value, the database server returns an error, because no default value exists.

The following example creates the **newitems** table. In **newitems**, the column **manucode** does not have a default value nor does it allow NULL values.

```
CREATE TABLE newitems (  
    newitem_num INTEGER,  
    manucode CHAR(3) NOT NULL,  
    promotype INTEGER,  
    descrip CHAR(20))
```

You cannot specify NULL as the explicit default value for a column if you also specify the NOT NULL constraint.

Using UNIQUE or DISTINCT Constraints

Use the UNIQUE or DISTINCT keyword to require that a column or set of columns accepts only unique data values. You cannot insert values that duplicate the values of some other row into a column that has a unique constraint. When you create a UNIQUE or DISTINCT constraint, the database server automatically creates an internal index on the constrained column or columns. (In this context, the keyword DISTINCT is a synonym for UNIQUE.)

You cannot place a unique constraint on a column that already has a primary-key constraint. You cannot place a unique constraint on a BYTE or TEXT column.

As previously noted, you cannot place a unique or primary-key constraint on a BLOB or CLOB column of Dynamic Server.

Opaque data types support a unique constraint only where a secondary-access method supports uniqueness for that type. The default secondary-access method is a generic B-tree, which supports the **equal()** operator function. Therefore, if the definition of the opaque type includes the **equal()** function, a column of that opaque type can have a unique constraint.

The following example creates a simple table that has a unique constraint on one of its columns:

```
CREATE TABLE accounts
  (acc_name CHAR(12),
   acc_num SERIAL UNIQUE CONSTRAINT acc_num);
```

For an explanation of the constraint name, refer to “Declaring a Constraint Name” on page 2-203.

Differences Between a Unique Constraint and a Unique Index

Although a unique index and a unique constraint are functionally similar, besides various differences in the syntax by which you declare, alter, or destroy them, there are additional differences between these two types of database objects:

- In DDL statements, they are registered or dropped in different tables of the system catalog
- In DML statements, enabled unique constraints on a logged table are checked at the end of a statement, but unique indexes are checked on a row-by-row basis, thereby preventing any insert or update of a row that might potentially violate the uniqueness of the specified column (or for a multiple-column column constraint or index, the column list).

For example, if you stored the values 1, 2, and 3 in rows of a logged table that has an INT column, an UPDATE operation on that table that specifies SET c = c + 1 would fail with an error if there were a unique index on the column c, but the statement would succeed if the column had a unique constraint.

Using the PRIMARY KEY Constraint

A *primary key* is a column (or a set of columns, if you use the multiple-column constraint format) that contains a non-NULL, unique value for each row in a table. When you define a PRIMARY KEY constraint, the database server automatically creates an internal index on the column or columns that make up the primary key.

You can designate only one primary key for a table. If you define a single column as the primary key, then it is unique by definition. You cannot explicitly give the same column a unique constraint.

In Dynamic Server, you cannot place a unique or primary-key constraint on a BLOB or CLOB column.

Opaque types of Dynamic Server support a primary key constraint only where a secondary-access method supports the uniqueness for that type. The default secondary-access method is a generic B-tree, which supports the **equal()** function. Therefore, if the definition of the opaque type includes the **equal()** function, a column of that opaque type can have a primary-key constraint.

You cannot place a primary-key constraint on a BYTE or TEXT column.

In the previous two examples, a unique constraint was placed on the column **acc_num**. The following example creates this column as the primary key for the **accounts** table:

```
CREATE TABLE accounts
  (acc_name CHAR(12),
   acc_num SERIAL PRIMARY KEY CONSTRAINT acc_num);
```

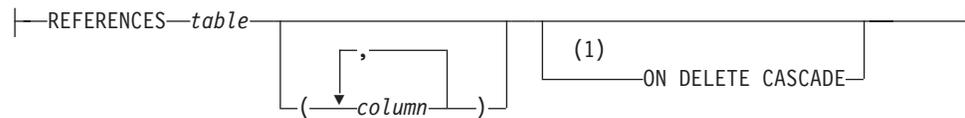
REFERENCES Clause

Use the REFERENCES clause to establish a referential relationship:

- Within a table (that is, between two columns of the same table)

- Between two tables (in other words, create a foreign key)

REFERENCES Clause:



Notes:

- 1 Informix extension

Element	Description	Restrictions	Syntax
<i>column</i>	A referenced column	See "Restrictions on Referential Constraints."	"Identifier " on page 5-21
<i>table</i>	The referenced table	Must reside in the same database as the referencing table	"Identifier " on page 5-21

The *referencing* column (the column being defined) is the column or set of columns that refers to the referenced column or set of columns. The referencing column can contain NULL and duplicate values, but values in the referenced column (or set of columns) must be unique.

The relationship between referenced and referencing columns is called a *parent-child* relationship, where the parent is the referenced column (primary key) and the child is the referencing column (foreign key). The referential constraint establishes this parent-child relationship.

When you create a referential constraint, the database server automatically creates an internal index on the constrained column or columns.

Restrictions on Referential Constraints

You must have the References privilege to create a referential constraint.

When you use the REFERENCES clause, you must observe the following restrictions:

- The referenced and referencing tables must be in the same database.
- The referenced column (or set of columns when you use the multiple-column constraint format) must have a unique or primary-key constraint.
- The data types of the referencing and referenced columns must be identical. The only exception is that a referencing column must be an integer data type if the referenced column is a serial data type.
- You cannot place a referential constraint on a BYTE or TEXT column.
- You cannot place a referential constraint on a BLOB or CLOB column.
- When you use the single-column constraint format, you can reference only one column.
- In Dynamic Server, when you use the multiple-column constraint format, the maximum number of columns in the REFERENCES clause is 16, and the total length of the columns cannot exceed 390 bytes if the page size is 2 kilobytes. (The maximum length increases with the page size.)

- In Extended Parallel Server, when you use the multiple-column constraint format, the maximum number of columns in the REFERENCES clause is 16, and the total length of the columns cannot exceed 380 bytes.

Default Values for the Referenced Column

If the referenced table is different from the referencing table, you do not need to specify the referenced column; the default column is the primary-key column (or columns) of the referenced table. If the referenced table is the same as the referencing table, you must specify the referenced column.

Referential Relationships Within a Table

You can establish a referential relationship between two columns of the same table. In the following example, the **emp_num** column in the **employee** table uniquely identifies every employee through an employee number. The **mgr_num** column in that table contains the employee number of the manager who manages that employee. In this case, **mgr_num** references **emp_num**. Duplicate values appear in the **mgr_num** column because managers manage more than one employee.

```
CREATE TABLE employee
(
  emp_num INTEGER PRIMARY KEY,
  mgr_num INTEGER REFERENCES employee (emp_num)
);
```

Locking Implications of Creating a Referential Constraint

When you create a referential constraint, an exclusive lock is placed on the referenced table. The lock is released when the CREATE TABLE statement is finished. If you are creating a table in a database that supports transaction logging, and you are using transactions, the lock is released at the end of the transaction.

Example That Uses the Single-Column Constraint Format

The following example uses the single-column constraint format to create a referential relationship between the **sub_accounts** and **accounts** tables. The **ref_num** column in the **sub_accounts** table references the **acc_num** column (the primary key) in the **accounts** table.

```
CREATE TABLE accounts (
  acc_num INTEGER PRIMARY KEY,
  acc_type INTEGER,
  acc_descr CHAR(20));
CREATE TABLE sub_accounts (
  sub_acc INTEGER PRIMARY KEY,
  ref_num INTEGER REFERENCES accounts (acc_num),
  sub_descr CHAR(20));
```

When you use the single-column constraint format, you do not explicitly specify the **ref_num** column as a foreign key. To use the FOREIGN KEY keyword, use the “Multiple-Column Constraint Format” on page 2-205.

Using the ON DELETE CASCADE Option

Use the ON DELETE CASCADE option to specify whether you want rows deleted in a child table when corresponding rows are deleted in the parent table. If you do not specify cascading deletes, the default behavior of the database server prevents you from deleting data in a table if other tables reference it.

If you specify this option, later when you delete a row in the parent table, the database server also deletes any rows associated with that row (foreign keys) in a

child table. The principal advantage to the cascading-deletes feature is that it allows you to reduce the quantity of SQL statements you need to perform delete actions.

For example, the **all_candy** table contains the **candy_num** column as a primary key. The **hard_candy** table refers to the **candy_num** column as a foreign key. The following CREATE TABLE statement creates the **hard_candy** table with the cascading-delete option on the foreign key:

```
CREATE TABLE all_candy
  (candy_num SERIAL PRIMARY KEY,
   candy_maker CHAR(25));

CREATE TABLE hard_candy
  (candy_num INT,
   candy_flavor CHAR(20),
   FOREIGN KEY (candy_num) REFERENCES all_candy
   ON DELETE CASCADE);
```

Because ON DELETE CASCADE is specified for the dependent table, when a row of the **all_candy** table is deleted, the corresponding rows of the **hard_candy** table are also deleted. For information about syntax restrictions and locking implications when you delete rows from tables that have cascading deletes, see “Considerations When Tables Have Cascading Deletes” on page 2-305.

CHECK Clause

Use the CHECK clause to designate conditions that must be met *before* data can be assigned to a column during an INSERT or UPDATE statement.

CHECK Clause:

```
|—CHECK—(—| Condition |———)———|
```

Notes:

- 1 See “Condition ” on page 4-5

In Dynamic Server, the *condition* cannot include a user-defined routine.

During an insert or update, if the check constraint of a row evaluates to *false*, the database server returns an error. The database server does not return an error if a row evaluates to NULL for a check constraint. In some cases, you might want to use both a check constraint and a NOT NULL constraint.

Using a Search Condition

The *search condition* that defines a check constraint cannot contain the following elements: user-defined routines, subqueries, aggregates, host variables, or rowids. In addition, the search condition cannot contain the following built-in variant functions: CURRENT, SYSDATE, USER, SITENAME, DBSERVERNAME, or TODAY.

When you specify a date value in a search condition, make sure you specify four digits for the year, so that the **DBCENTURY** environment variable has no effect on the condition. When you specify a two-digit year, the **DBCENTURY** environment

variable can produce unpredictable results if the condition depends on an abbreviated year value. For more information about **DBCENTURY**, see the *IBM Informix Guide to SQL: Reference*.

More generally, the database server saves the settings of environment variables from the time of creation of check constraints. If any of these settings are subsequently changed in a way that can affect the evaluation of a condition in a check constraint, the new settings are disregarded, and the original environment variable settings are used when the condition is evaluated.

With a BYTE or TEXT column, you can check for NULL or not-NULL values. This constraint is the only constraint allowed on a BYTE or TEXT column.

Restrictions When Using the Single-Column Constraint Format

When you use the single-column constraint format to define a check constraint, the check constraint cannot depend on values in other columns of the table. The following example creates the **my_accounts** table that has two columns with check constraints, each in the single-column constraint format:

```
CREATE TABLE my_accounts (
  chk_id SERIAL PRIMARY KEY,
  acct1 MONEY CHECK (acct1 BETWEEN 0 AND 99999),
  acct2 MONEY CHECK (acct2 BETWEEN 0 AND 99999));
```

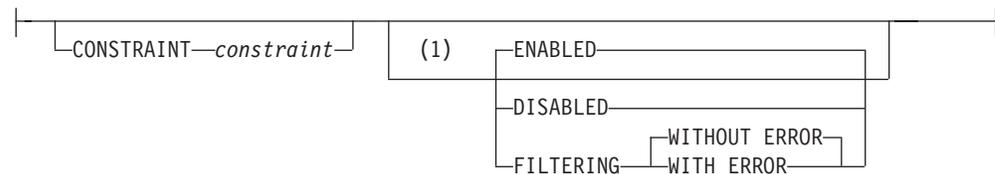
Both **acct1** and **acct2** are columns of MONEY data type whose values must be between 0 and 99999. If, however, you want to test that **acct1** has a larger balance than **acct2**, you cannot use the single-column constraint format. To create a constraint that checks values in more than one column, you must use the "Multiple-Column Constraint Format" on page 2-205.

Constraint Definition

Use the constraint definition portion of CREATE TABLE for these purposes:

- To declare a name for the constraint
- To set a constraint to disabled, enabled, or filtering mode (IDS).

Constraint Definition:



Notes:

- 1 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>constraint</i>	Name of constraint	Must be unique for the table among index and constraint names	"Identifier " on page 5-21

Declaring a Constraint Name

The database server implements the constraint as an index. Whenever you use the single- or multiple-column constraint format to place a data restriction on a

column, but without declaring a *constraint* name, the database server creates a constraint and adds a row for that constraint in the **sysconstraints** system catalog table. The database server also generates an identifier and adds a row to the **sysindexes** system catalog table for each new primary-key, unique, or referential constraint that does not share an index with an existing constraint. Even if you declare a name for a constraint, the database server generates the name that appears in the **sysindexes** table.

If you want, you can specify a meaningful name for the constraint. The name must be unique among the names of constraints and indexes in the database.

Constraint names appear in error messages having to do with constraint violations. You can use this name when you use the DROP CONSTRAINT clause of the ALTER TABLE statement.

In Dynamic Server, you also specify a constraint name when you change the mode of constraint with the SET Database Object Mode statement or the SET Transaction Mode statement, and in the DROP INDEX statement for constraints that are implemented as indexes with user-defined names.

In an ANSI-compliant database, when you declare the name of a constraint of any type, the combination of the *owner* name and *constraint* name must be unique within the database.

In Dynamic Server, the system catalog table that holds information about indexes is the **sysindices** table.

Constraint Names That the Database Server Generates: If you do not specify a constraint name, the database server generates a constraint name using the following template:

```
<constraint_type><tabid>_<constraintid>
```

In this template, *constraint_type* is the letter **u** for unique or primary-key constraints, **r** for referential constraints, **c** for check constraints, and **n** for NOT NULL constraints. In the template, *tabid* and *constraintid* are values from the **tabid** and **constrid** columns of the **sysables** and **sysconstraints** system catalog tables, respectively. For example, the constraint name for a unique constraint might look like " **u111_14** " (with a leading blank space).

If the generated name conflicts with an existing identifier, the database server returns an error, and you must then supply an explicit constraint name.

The generated index name in **sysindexes** (or **sysindices**) has this format:

```
[blankspace]<tabid>_<constraintid>
```

For example, the index name might be something like " **111_14** " (quotation marks used here to show the blank space).

Choosing a Constraint-Mode Option (IDS)

Use the constraint-mode options to control the behavior of constraints in INSERT, DELETE, and UPDATE operations. These are the options.

Mode	Effect
------	--------

DISABLED

Does not enforce the constraint during INSERT, DELETE, and UPDATE operations

ENABLED

Enforces the constraint during INSERT, DELETE, and UPDATE operations. If a target row causes a violation of the constraint, the statement fails. This mode is the default.

FILTERING

Enforces the constraint during INSERT, DELETE, and UPDATE operations. If a target row causes a violation of the constraint, the statement continues processing. The database server writes the row in question to the violations table associated with the target table and writes diagnostic information to the associated diagnostics table.

If you choose filtering mode, you can specify the WITHOUT ERROR or WITH ERROR options. The following list explains these options.

Error Option Effect

WITHOUT ERROR

Does not return an integrity-violation error when a filtering-mode constraint is violated during an insert, delete, or update operation. This is the default error option.

WITH ERROR

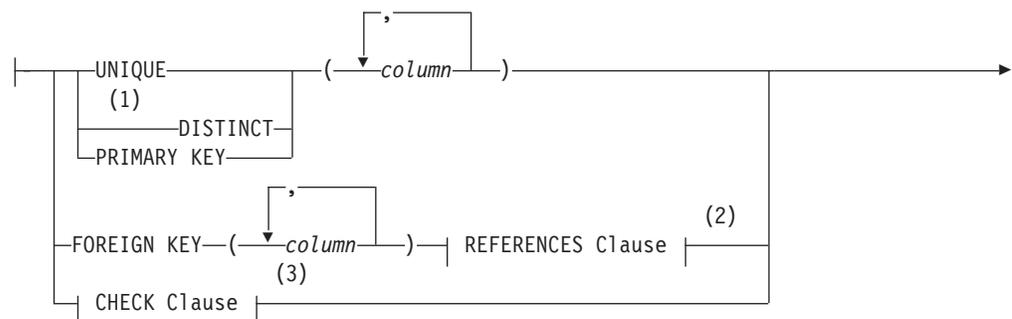
Returns an integrity-violation error when a filtering-mode constraint is violated during an insert, delete, or update operation

To reset the constraint mode of a table, see “SET Database Object Mode” on page 2-591. For information about where the database server stores rows that violate a constraint set to FILTERING, see “START VIOLATIONS TABLE” on page 2-663.

Multiple-Column Constraint Format

Use the multiple-column constraint format to associate one or more columns with a constraint. This alternative to the single-column constraint format allows you to associate multiple columns with a constraint.

Multiple-Column Constraint Format:





Notes:

- 1 Informix extension
- 2 See “REFERENCES Clause” on page 2-199
- 3 See “CHECK Clause” on page 2-202
- 4 See “Constraint Definition” on page 2-203

Element	Description	Restrictions	Syntax
<i>column</i>	Columns on which to place constraint	Not BYTE, TEXT, BLOB, CLOB	“Identifier ” on page 5-21

You can include a maximum of 16 columns in a constraint list. For databases where the page size is two kilobytes, the total length of the list of columns cannot exceed 380 bytes.

When you define a unique constraint (by using the UNIQUE or DISTINCT keyword), a column cannot appear in the constraint list more than once.

Using the multiple-column constraint format, you can perform these tasks:

- Create data-integrity constraints for a set of one or more columns
- Specify a mnemonic name for a constraint
- Specify the constraint-mode option that controls the behavior of a constraint during insert, delete, and update operations.

When you use this format, you can create composite primary and foreign keys, or define check constraints that compare data in different columns.

See also the section “Differences Between a Unique Constraint and a Unique Index” on page 2-199.

Restrictions with the Multiple-Column Constraint Format

Using Large-Object Types in Constraints: You cannot place unique, primary-key, or referential (FOREIGN KEY) constraints on BYTE or TEXT columns. You can, however, check for NULL or non-NULL values with a check constraint.

You cannot place unique or primary-key constraints on BLOB or CLOB columns. If the constraint is on a set of columns that includes a column that stores encrypted data, Dynamic Server cannot enforce the constraint.

You can find detailed discussions of specific constraints in the following sections:

Constraint	For more information, see	For an example, see
CHECK	“CHECK Clause” on page 2-202	“Defining Check Constraints Across Columns” on page 2-207
DISTINCT	“Using UNIQUE or DISTINCT Constraints” on page 2-198	“Examples of the Multiple-Column Constraint Format” on page 2-207

Constraint	For more information, see	For an example, see
FOREIGN KEY	“Using the FOREIGN KEY Constraint”	“Defining Composite Primary and Foreign Keys” on page 2-208
PRIMARY KEY	“Using the PRIMARY KEY Constraint” on page 2-199	“Defining Composite Primary and Foreign Keys” on page 2-208
UNIQUE	“Using UNIQUE or DISTINCT Constraints” on page 2-198	“Examples of the Multiple-Column Constraint Format”

Using the FOREIGN KEY Constraint

A foreign key *joins* and establishes dependencies between tables. That is, it creates a referential constraint. (For more information on referential constraints, see the “REFERENCES Clause” on page 2-199.)

A foreign key references a unique or primary key in a table. For every entry in the foreign-key columns, a matching entry must exist in the unique or primary-key columns if all foreign-key columns contain non-NULL values.

You cannot specify BYTE or TEXT columns as foreign keys.

You cannot specify BLOB or CLOB columns as foreign keys.

Examples of the Multiple-Column Constraint Format

The following example creates a standard table, called **order_items**, with a unique constraint, called **items_constr**, using the multiple-column constraint format:

```
CREATE TABLE order_items
(
  order_id SERIAL,
  line_item_id INT not null,
  unit_price DECIMAL(6,2),
  quantity INT,
  UNIQUE (order_id,line_item_id) CONSTRAINT items_constr
);
```

For constraint names, see “Declaring a Constraint Name” on page 2-203.

Defining Check Constraints Across Columns:

When you use the multiple-column constraint format to define check constraints, a check constraint can apply to more than one column in the same table. (You cannot, however, create a check constraint whose *condition* uses a value from a column in another table.)

This example compares two columns, **acct1** and **acct2**, in the new table:

```
CREATE TABLE my_accounts
(
  chk_id SERIAL PRIMARY KEY,
  acct1 MONEY,
  acct2 MONEY,
  CHECK (0 < acct1 AND acct1 < 99999),
  CHECK (0 < acct2 AND acct2 < 99999),
  CHECK (acct1 > acct2)
);
```

In this example, the **acct1** column must be greater than the **acct2** column, or the insert or update fails.

Defining Composite Primary and Foreign Keys:

When you use the multiple-column constraint format, you can create a composite key. A *composite key* specifies multiple columns for a primary-key or foreign-key constraint.

The next example creates two tables. The first table has a composite key that acts as a primary key, and the second table has a composite key that acts as a foreign key.

```
CREATE TABLE accounts (  
    acc_num INTEGER,  
    acc_type INTEGER,  
    acc_descr CHAR(20),  
    PRIMARY KEY (acc_num, acc_type));  
  
CREATE TABLE sub_accounts (  
    sub_acc INTEGER PRIMARY KEY,  
    ref_num INTEGER NOT NULL,  
    ref_type INTEGER NOT NULL,  
    sub_descr CHAR(20),  
    FOREIGN KEY (ref_num, ref_type) REFERENCES accounts  
        (acc_num, acc_type));
```

In this example, the foreign key of the **sub_accounts** table, **ref_num** and **ref_type**, references the composite key, **acc_num** and **acc_type**, in the **accounts** table. If, during an insert or update, you tried to insert a row into the **sub_accounts** table whose value for **ref_num** and **ref_type** did not exactly correspond to the values for **acc_num** and **acc_type** in an existing row in the **accounts** table, the database server would return an error.

A referential constraint must have a one-to-one relationship between referencing and referenced columns. In other words, if the primary key is a set of columns (a composite key), then the foreign key also must be a set of columns that corresponds to the composite key.

Because of the default behavior of the database server, when you create the foreign-key reference, you do not need to reference the composite-key columns (**acc_num** and **acc_type**) explicitly. You can rewrite the references section of the previous example as follows:

```
FOREIGN KEY (ref_num, ref_type) REFERENCES accounts
```

Default Index Creation Strategy for Constraints

When you create a table with unique or primary-key constraints, the database server creates an internal index that is unique and ascending for each constraint.

When you create a table with a referential constraint, the database server creates an internal index that is ascending and that allows duplicate values for each column that you specify in the referential constraint.

An internal index occupies the same storage location as its table. For fragmented tables, the fragments of an internal index occupy the same dbspace partitions that you specify for the table fragments (or in some cases, the database dbspace).

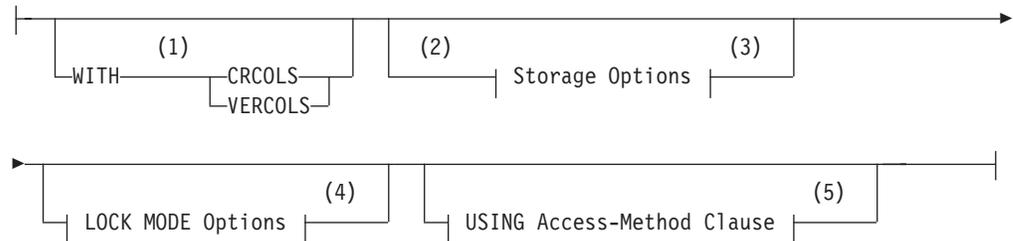
If you require an index fragmentation strategy that is independent of the underlying table fragmentation, do not define the constraint when you create the table. Instead, use the CREATE INDEX statement to create a unique index with the desired fragmentation strategy. Then use the ALTER TABLE statement to add the constraint. The new constraint uses the previously defined index.

Important: In a database without logging, *detached checking* is the only kind of constraint checking available. Detached checking means that constraint checking is performed on a row-by-row basis.

Options Clauses

The Options clauses of the CREATE TABLE statement provides options to create shadow columns, and to specify storage locations, a fragmentation strategy, the extent size, the locking mode, and a user-defined access method for the table.

Options:



Notes:

- 1 Dynamic Server only
- 2 Informix extension
- 3 See "Storage Options" on page 2-210
- 4 See "LOCK MODE Options" on page 2-225
- 5 See "USING Access-Method Clause (IDS)" on page 2-224

Using the WITH CRCOLS Option (IDS)

Use the WITH CRCOLS keywords to create two shadow columns that Enterprise Replication uses for conflict resolution. The first column, **cdrserver**, contains the identity of the database server where the last modification occurred. The second column, **cdrtime**, contains the time stamp of the last modification. You must add these columns before you can use time stamps for UDR conflict resolution.

For most database operations, the **cdrserver** and **cdrtime** columns are hidden. For example, if you include the WITH CRCOLS keywords when you create a table, the **cdrserver** and **cdrtime** columns have the following behavior:

- They are not returned by queries that specify an asterisk (*) as the projection list, as in the statement:

```
SELECT * FROM tablename;
```
- They do not appear in DB-Access when you ask for information about the columns of the table.
- They are not included in the number of columns (**ncols**) in the **systables** system catalog table entry for *tablename*.

To view the contents of **cdrserver** and **cdrtime**, you must explicitly specify the columns in the projection list of a SELECT statement, as the following example shows:

```
SELECT cdrserver, cdrtime FROM tablename;
```

For more information about how to use this option, refer to the *IBM Informix Dynamic Server Enterprise Replication Guide*.

Using the WITH VERCOLS Option (IDS)

Use the WITH VERCOLS keywords to create two shadow columns that Dynamic Server uses to support update operations on secondary servers. The first column, **ifx_insert_checksum**, contains a checksum of the row when it was first created. The second column, **ifx_row_version**, contains a version number of the row. When a row is first inserted, **ifx_insert_checksum** is generated, and **ifx_row_version** will be set to one. Each time the row is updated, **ifx_row_version** is incremented by one, but **ifx_insert_checksum** does not change.

For most database operations, the **ifx_insert_checksum** and **ifx_row_version** columns are hidden. For example, if you include the WITH VERCOLS keywords when you create a table, the **ifx_insert_checksum** and **ifx_row_version** columns have the following behavior:

- They are not returned by queries that specify an asterisk (*) as the projection list, as in the statement:

```
SELECT * FROM tablename;
```
- They appear in DB–Access when you ask for information about the columns of the table.
- They are included in the number of columns (**ncols**) in the **systables** system catalog table entry for *tablename*.

To view the contents of **ifx_insert_checksum** and **ifx_row_version**, you must explicitly specify the column names in the projection list of a SELECT statement, as the following example shows:

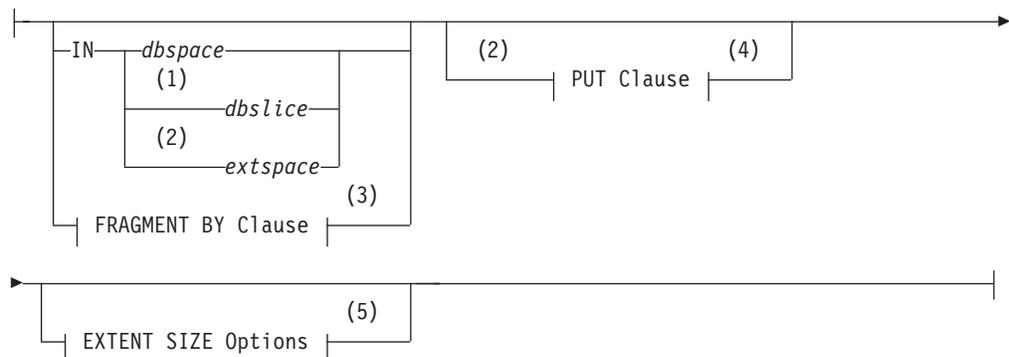
```
SELECT ifx_insert_checksum, ifx_row_version FROM tablename;
```

For more information about how to use this option, refer to the *IBM Informix Administrator's Guide*.

Storage Options

Use these options to specify the storage location, distribution scheme, and extent size for the table. This is an extension to the ANSI/ISO standard for SQL syntax.

Storage Options:



Notes:

- 1 Extended Parallel Server only

- 2 Dynamic Server only
- 3 See “FRAGMENT BY Clause” on page 2-212
- 4 See “PUT Clause (IDS)” on page 2-222
- 5 See “EXTENT SIZE Options” on page 2-224

Element	Description	Restrictions	Syntax
<i>dbslice</i>	Dbslice to store the table	Must already exist	“Identifier ” on page 5-21
<i>dbspace</i>	Dbospace to store the table	Must already exist	“Identifier ” on page 5-21
<i>extspace</i>	Name declared in the onspaces command to a storage area outside the database server	Must already exist	See documentation for your access method.

If you use the “USING Access-Method Clause (IDS)” on page 2-224 to specify an access method, that method must support the storage space.

You can specify a dbospace for the table that is different from the storage location for the database, or fragment the table among dbospaces, or among partitions of one or more dbospaces. If you specify no IN clause nor fragmentation scheme, the new table resides in the same dbospace where the current database resides.

In Dynamic Server, you can use the PUT clause to specify storage options for smart large objects. For more information, see “PUT Clause (IDS)” on page 2-222.

Note: If your table contains simple large objects (TEXT or BYTE), you can specify a separate blobospace for each object. For information on storing simple large objects, refer to “ Large-Object Data Types” on page 4-29.

Using the IN Clause

Use the IN clause to specify a storage space for the table. The storage space that you specify must already exist.

Storing Data in a dbospace: You can use the IN clause to isolate a table. For example, if the **history** database is in the **dbos1** dbospace, but you want the **family** data placed in a separate dbospace called **famdata**, use the following statements:

```
CREATE DATABASE history IN dbos1;

CREATE TABLE family
(
  id_num      SERIAL(101) UNIQUE,
  name        CHAR(40),
  nickname    CHAR(20),
  mother      CHAR(40),
  father      CHAR(40)
)
IN famdata;
```

For more information about how to store and manage your tables in separate dbospaces, see your *IBM Informix Administrator’s Guide*.

Storing Data in a Partition of a dbospace (IDS): Besides the option of storing the table (or a fragment of it) in a dbospace, Dynamic Server supports storing fragments of a table in a subset of a dbospace, called a *partition*. Unless you explicitly declare names for the fragments in the PARTITION BY clause, each fragment, by default, has the same name as the dbospace where it resides. This includes all fragmented tables and indexes migrated from earlier releases of Dynamic Server.

You can store fragments of the same table in multiple partitions of the same dbspace, but each name that you declare after the PARTITION keyword must be unique among partitions of that table. The PARTITION keyword is required when you store more than one fragment of the table in the same dbspace. You can also use the PARTITION keyword to declare a more meaningful name for a dbspace that has only one partition. The number of partitions in a single dbspace cannot exceed 2,147,483,647, which is the maximum value of an INTEGER data type.

Regardless of the length of a row, no more than 16,777,216 pages of the same partition can store rows from the same table, with no more than 256 rows per page. By creating two or more partitions in the same dbspace, however, you can avoid situations in which you cannot insert additional rows of a table into the unused portion of a dbspace after the cardinality of the number of pages storing rows in the dbspace has reached this upper limit.

Storing Data in a dbslice (XPS): If you are using Extended Parallel Server, the IN *dbslice* clause allows you to fragment a table across a group of dbspaces that share the same naming convention. The database server fragments the table by round-robin in the dbspaces that make up the dbslice at the time the table is created.

To fragment a table across a dbslice, you can use either the IN *dbslice* syntax or the FRAGMENT BY ROUND ROBIN IN *dbslice* syntax.

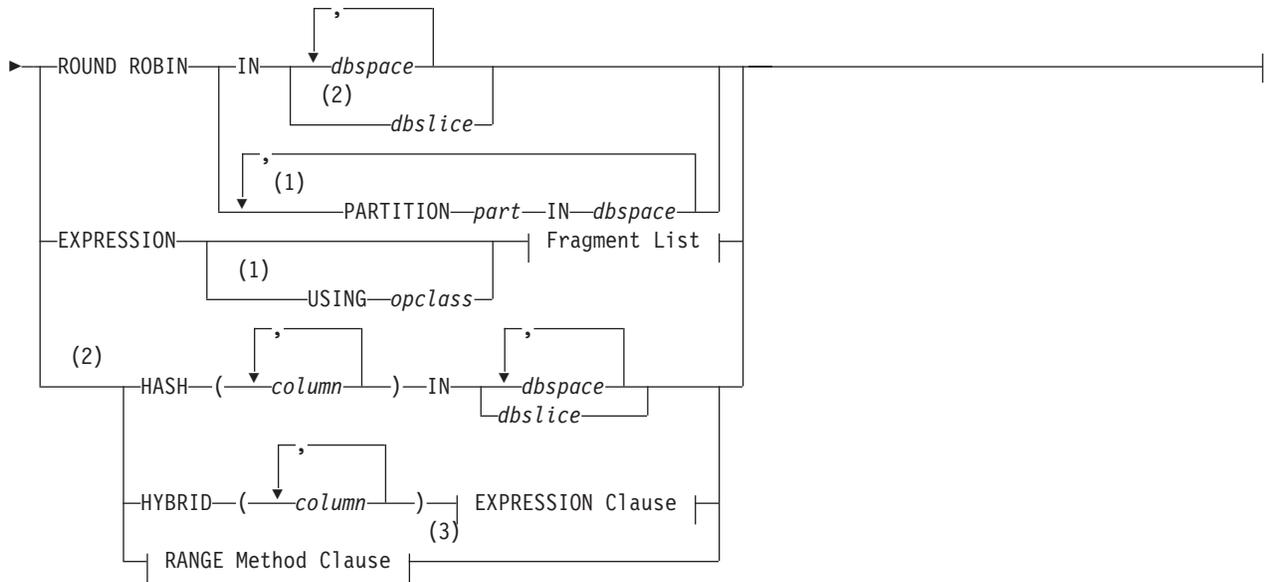
Storing Data in an extspace (IDS): In general, use the *extspace* storage option in conjunction with the "USING Access-Method Clause (IDS)" on page 2-224. For more information, refer to the documentation of your access method.

FRAGMENT BY Clause

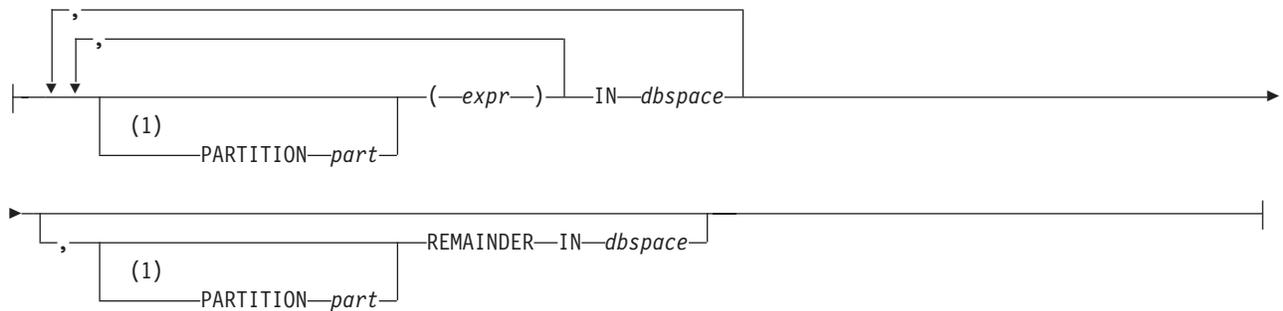
Use the FRAGMENT BY clause to create fragmented tables and to specify their distribution scheme. (The keywords PARTITION BY are a synonym for FRAGMENT BY in Dynamic Server.)

FRAGMENT BY Clause for Tables:

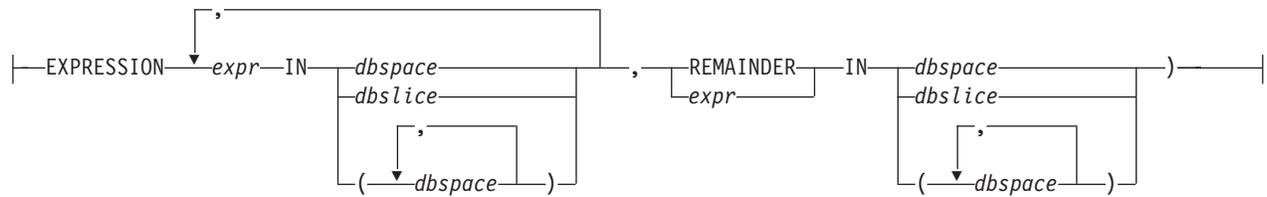




Fragment List:



EXPRESSION Clause:



Notes:

- 1 Dynamic Server only
- 2 Extended Parallel Server only
- 3 See "RANGE Method Clause (XPS)" on page 2-218

Element	Description	Restrictions	Syntax
<i>column</i>	Column to which to apply the fragmentation strategy	Must be a column within the table	"Identifier " on page 5-21
<i>dbslice, dbspace</i>	Dbslice or dbspace to store the table fragment	The <i>dbslice</i> must be defined. You can specify no more than 2,048 <i>dbspaces</i> . All <i>dbspaces</i> that store the fragments must have the same page size.	"Identifier " on page 5-21
<i>expr</i>	Expression that defines a table fragment using a range, hash, or arbitrary rule	Columns can be from the current table only, and data values can be from only a single row. Value returned must be Boolean (true or false).	"Expression" on page 4-40
<i>opclass</i>	No default operator class	Must be defined and must be associated with a B-tree index	"Identifier " on page 5-21
<i>part</i>	Name that you declare here for a partition of a dbspace	Required for any partition in the same dbspace as another partition of the same table. Name must be unique among partitions of the same table.	"Identifier " on page 5-21

When you fragment a table, the IN keyword is followed by the name of the storage space where a table fragment is to be stored.

Using the WITH ROWIDS Option (IDS)

Nonfragmented tables contain a hidden column called **rowid**, but by default, fragmented tables have no **rowid** column. You can use the WITH ROWIDS keywords to add the **rowid** column to a fragmented table. Each row is automatically assigned a unique **rowid** value that remains stable for the life of the row and that the database server can use to find the physical location of the row. Each row requires an additional four bytes to store the **rowid**.

Important: This is a deprecated feature. Use primary keys as an access method rather than the **rowid** column.

You cannot use the WITH ROWIDS clause with typed tables.

Fragmenting by ROUND ROBIN

In a round-robin distribution scheme, specify at least two dbspaces (in Extended Parallel Server) or partitions (in Dynamic Server) where you want the fragments to be placed. As records are inserted into the table, they are placed in the first available fragment. The database server balances the load among the specified fragments as you insert records and distributes the rows in such a way that the fragments always maintain approximately the same number of rows. In this distribution scheme, the database server must scan all fragments when it searches for a row.

With Extended Parallel Server, you can specify a dbslice to fragment a table across a group of dbspaces that share the same naming convention. For a syntax alternative to FRAGMENT BY ROUND ROBIN IN *dbslice* that achieves the same results, see "Storing Data in a dbslice (XPS)" on page 2-212.

With Dynamic Server, you can use the PUT clause to specify round-robin fragmentation for smart large objects. For more information, see the "PUT Clause (IDS)" on page 2-222.

Fragmenting by EXPRESSION

In an *expression-based* distribution scheme, each fragment expression in a rule specifies a storage space. Each fragment expression in the rule isolates data and aids the database server in searching for rows.

To fragment a table by expression, specify one of the following rules:

- Range rule

A range rule specifies fragment expressions that use a range to specify which rows are placed in a fragment, as the next example shows:

```
FRAGMENT BY EXPRESSION c1 < 100 IN dbsp1,  
    c1 >= 100 AND c1 < 200 IN dbsp2, c1 >= 200 IN dbsp3;
```

- Arbitrary rule

An arbitrary rule specifies fragment expressions based on a predefined SQL expression that typically uses OR clauses to group data, as the following example shows:

```
FRAGMENT BY EXPRESSION  
    zip_num = 95228 OR zip_num = 95443 IN dbsp2,  
    zip_num = 91120 OR zip_num = 92310 IN dbsp4,  
    REMAINDER IN dbsp5;
```

Warning: See the note about the **DBCENTURY** environment variable and date values in fragment expressions in the section “Logging Options” on page 2-193.

The USING Oclass Option (IDS):

With the *USING operator class* option, you can specify a nondefault operator class for the fragmentation strategy. The secondary-access method of the chosen operator class must have a B-tree index structure.

In the following example, the **abs_btree_ops** operator class specifies several user-defined strategy functions that order integers based on their absolute values:

```
CREATE OPCLASS abs_btree_ops FOR btree  
    STRATEGIES (abs_lt, abs_lte, abs_eq, abs_gte, abs_gt)  
    SUPPORT (abs_cmp);
```

For the fragmentation strategy, you can specify the **abs_btree_ops** operator class in the *USING* clause and use its strategy functions to fragment the table, as follows:

```
FRAGMENT BY EXPRESSION USING abs_btree_ops  
    (abs_lt(x,345)) IN dbsp1,  
    (abs_gte(x,345) AND abs_lte(x,500)) IN dbsp2,  
    (abs_gt(x,500)) IN dbsp3;
```

For information on how to create and extend an operator class, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

User-Defined Functions in Fragment Expressions (IDS): For rows that include user-defined data types, you can use comparison conditions or user-defined functions to define the range rules. In the following example, comparison conditions define the range rules for the **long1** column, which contains an opaque data type:

```
FRAGMENT BY EXPRESSION  
    long1 < '3001' IN dbsp1,  
    long1 BETWEEN '3001' AND '6000' IN dbsp2,  
    long1 > '6000' IN dbsp3;
```

An implicit, user-defined cast converts 3001 and 6000 to the opaque type.

Alternatively, you can use user-defined functions to define the range rules for the opaque data type of the **long1** column:

```

FRAGMENT BY EXPRESSION
  (lessthan(long1,'3001')) IN dbsp1,
  (greaterthanorequal(long1,'3001') AND
  lessthanorequal(long1,'6000')) IN dbsp2,
  (greaterthan(long1,'6000')) IN dbsp3;

```

Explicit user-defined functions require parentheses around the entire fragment expression before the IN clause, as the previous example shows.

User-defined functions in a fragment expression can be written in SPL or in the C or Java language. These functions must satisfy four requirements:

- They must evaluate to a Boolean value.
- They must be nonvariant.
- They must reside within the same database as the table.
- They must not generate OUT nor INOUT parameters.

For information on how to create UDRs for fragment expressions, refer to *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

Using the REMAINDER Keyword: Use the REMAINDER keyword to specify the storage space in which to store valid values that fall outside the specified expression or expressions. If you do not specify a remainder, and a row is inserted or updated with values that do not correspond to any fragment definition, the database server returns an error.

The following example uses an arbitrary rule to define five fragments for specific values of the **c1** column, and a sixth fragment for all other values:

```

CREATE TABLE T1 (c1 INT) FRAGMENT BY EXPRESSION
  PARTITION PART_1 (c1 = 10) IN dbsp1,
  PARTITION PART_2 (c1 = 20) IN dbsp1,
  PARTITION PART_3 (c1 = 30) IN dbsp1,
  PARTITION PART_4 (c1 = 40) IN dbsp2,
  PARTITION PART_5 (c1 = 50) IN dbsp2,
  PARTITION PART_6 REMAINDER IN dbsp2;

```

Here the first three fragments are stored in partitions of the **dbsp1** dbspace, and the other fragments, including the remainder, are stored in partitions of the **dbsp2** dbspace. Explicit fragment names are required in this example, because each dbspace has multiple partitions.

Fragmenting by HASH (XPS)

A hash-distribution scheme distributes the rows as you insert them, so that the fragments maintain approximately the same number of rows. In this distribution scheme, the database server can eliminate fragments when it searches for a row because the hash is known internally. For example, if you have a large database, as in a data-warehousing environment, you can fragment your tables across disks that belong to different coservers. If you expect to perform many queries that scan most of the data, a system-defined hash-distribution scheme can balance the I/O processing. The following example uses eight coservers with one dbspace defined on each coserver.

```

CREATE TABLE customer
(
  cust_id integer,
  descr char(45),
  level char(15),
  sale_type char(10),
  channel char(30),
  corp char(45),

```

```

    cust char(45),
    vert_mkt char(30),
    state_prov char(20),
    country char(15),
    org_cust_id char(20)
)
FRAGMENT BY HASH (cust_id) IN
    customer1_spc,
    customer2_spc,
    customer3_spc,
    customer4_spc,
    customer5_spc,
    customer6_spc,
    customer7_spc,
    customer8_spc
EXTENT SIZE 20 NEXT SIZE 16;

```

You can also specify a *dbslice*. When you specify a *dbslice*, the database server fragments the table across the *dbspaces* that make up the *dbslice*.

Serial Columns in HASH-Distribution Schemes: If you base table fragmentation on a SERIAL or SERIAL8 column, only a hash-distribution scheme is valid. In addition, the serial column must be the only column in the hashing key. (These restrictions apply only to *table* distributions. Fragmentation schemes for *indexes* that are based on SERIAL or SERIAL8 columns are not subject to these restrictions.)

The following excerpt is from a CREATE TABLE statement:

```

CREATE TABLE customer
(
    cust_id serial,
    . . .
)
FRAGMENT BY HASH (cust_id) IN customer1_spc, customer2_spc;

```

You might notice a difference between serial-column values in fragmented and nonfragmented tables. The database server assigns serial values round-robin across fragments, so a fragment might contain values from noncontiguous ranges. For example, if there are two fragments, the first serial value is placed in the first fragment, the second serial value is placed in the second fragment, the third value is placed in the first fragment, and so on.

Fragmenting by HYBRID (XPS)

The HYBRID clause allows you to apply two distribution schemes to the same table. You can use a combination of hash- and expression-distribution schemes or a combination of range-distribution schemes on a table. This section discusses the hash and expression form of hybrid fragmentation. For details of range fragmentation, see “RANGE Method Clause (XPS)” on page 2-218.

In hybrid fragmentation, the EXPRESSION clause determines the base fragmentation strategy of the table, associating an expression with a set of *dbspaces* (*dbspace*, *dbslice*, or *dbspacelist* format) for data storage. The hash columns determine the *dbspace* within the specified set of *dbspaces*.

When you specify a *dbslice*, the database server fragments the table across the *dbspaces* that make up the *dbslice*. Similarly, if you specify a *dbspace* list, the database server fragments the table across the *dbspaces* in that list. In the next example, **my_hybrid** distributes rows based on two columns of the table. The value of **col1** determines in which *dbslice* the row belongs.

The hash value of **col2** determines which dbspace (within the previously determined dbslice) to insert into.

```
CREATE TABLE my_hybrid
  (col1 INT, col2 DATE, col3 CHAR(10))
  HYBRID (col2) EXPRESSION col1 < 100 IN dbslice1,
  col1 >= 100 and col1 < 200 IN dbslice2,REMAINDER IN dbslice3;
```

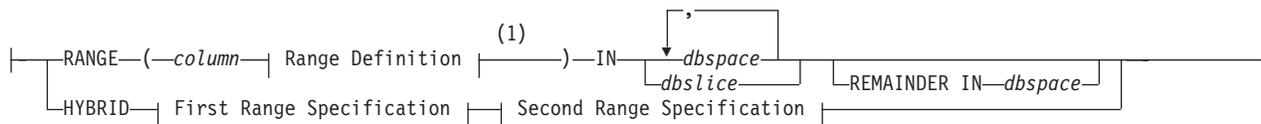
For more information on an expression-based distribution scheme, see “Fragmenting by EXPRESSION” on page 2-214.

RANGE Method Clause (XPS)

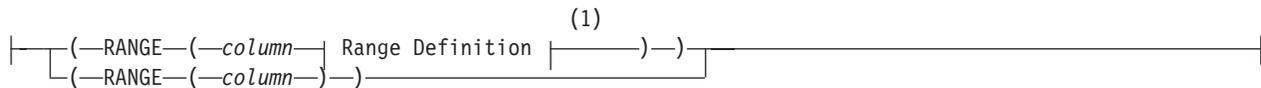
In Extended Parallel Server, you can use a range-fragmentation method as a convenient alternative to fragmenting by the EXPRESSION or HYBRID clauses. This provides a method to implicitly and uniformly distribute data whose fragmentation column values are dense or naturally uniform.

In a range-fragmented table, each dbspace stores a contiguous, completely bound and non-overlapping range of integer values over one or two columns. In other words, the database server implicitly clusters rows within the fragments, based on the range of the values in the fragmentation column.

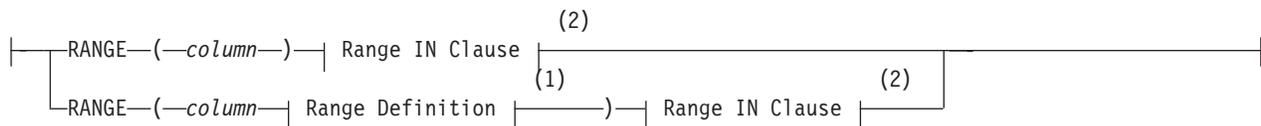
RANGE Method Clause:



First Range Specification:



Second Range Specification:



Notes:

- 1 See “ Range Definition” on page 2-219
- 2 See “ Range IN Clause” on page 2-219

Element	Description	Restrictions	Syntax
<i>column</i>	Column on which to apply the fragmentation strategy	Must be in the current table and must be of data type INT or SMALL INT	“Identifier ” on page 5-21
<i>dbslice</i>	Dbslice that contains the dbspaces where the table fragments reside	Must exist when you execute the statement	“Identifier ” on page 5-21
<i>dbspace</i>	Dbspace that contains the table fragment	Must exist when you execute the statement. The maximum number of dbspaces is 2048.	“Identifier ” on page 5-21

For hybrid strategies with two range definitions, the second *column* must have a different column name from the first. For hybrid strategies with exactly one range definition, both occurrences of *column* must specify the same column.

If you list more than one *dbslice*, including a remainder *dbslice*, each *dbslice* must contain the same number of *dbspaces*. Unless you are specifying the *dbspace* in the REMAINDER option, you must specify at least two *dbspaces*.

Range Definition

Use the range definition to specify the minimum and maximum values of the entire range.

Range Definition:



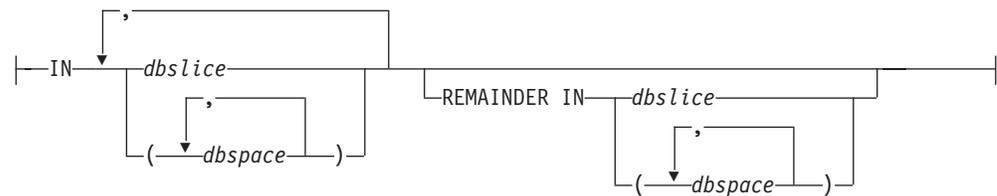
Element	Description	Restrictions	Syntax
<i>max_val</i>	Maximum value in the range	Must be an INT or SMALLINT greater than or equal to <i>min_val</i> , if <i>min_val</i> is supplied	"Literal Number " on page 4-185
<i>min_val</i>	Minimum value in the range; the default is 0.	Must be an INT or SMALLINT less than or equal to <i>max_val</i>	"Literal Number " on page 4-185

You do not need to specify a minimum value. The minimum and maximum values define the exact range of values to allocate for each storage space.

Range IN Clause

Use the IN clause to specify the storage spaces in which to distribute the data.

Range IN Clause:



Element	Description	Restrictions	Syntax
<i>dbslice</i>	<i>Db</i> slice that contains the <i>dbspaces</i> to store table fragments	Must exist	"Identifier " on page 5-21
<i>dbspace</i>	<i>Db</i> space to store the table fragment	Must exist	"Identifier " on page 5-21

If you specify more than one *dbslice*, including a remainder *dbslice*, each *dbslice* must contain the same number of *dbspaces*.

Unless you are specifying the dbspace in the REMAINDER option, the minimum number of dbspaces that you can specify is two. The maximum number of dbspaces that you can specify is 2,048.

When you use a range-fragmentation method, the number of integer values between the minimum and maximum specified values must be equal to or greater than the number of storage spaces specified, so that the database server can allocate non-overlapping contiguous ranges across the dbspaces. For example, the following code returns an error, because the allocations for the range cannot be distributed across all specified dbspaces:

```
CREATE TABLE Tab1 (Col1 INT...)
  FRAGMENT BY RANGE (Col1 MIN 5 MAX 7)
  IN db1, db2, db3, db4, db5, db6; -- returns an error
```

The error for this example occurs because the specified range contains three values (5, 6, and 7), but six dbspaces are specified; three values cannot be distributed across six dbspaces.

Using the REMAINDER Keyword

Use the REMAINDER keyword to specify the storage space in which to store valid values that fall outside the specified expression or expressions.

If you do not specify a remainder and a row is inserted or updated such that it no longer belongs to any storage space, the database server returns an error.

Restrictions

If you fragment a table with range fragmentation, you cannot perform the following operations on the table after it is created:

- You cannot change the fragmentation strategy (ALTER FRAGMENT).
- You cannot rename the columns of the table (RENAME COLUMN).
- You cannot alter the table in any way except to change the table type or to change the lock mode.

That is, the Usage-TYPE options and the Lock Mode clause are the only valid options of ALTER TABLE for a table that has range fragmentation.

Examples

The following examples illustrate range fragmentation in its simple and hybrid forms.

Simple Range-Fragmentation Strategy:

The following example shows a simple range-fragmentation strategy:

```
CREATE TABLE Tab1 (Col1 INT...)
  FRAGMENT BY RANGE (Col1 MIN 100 MAX 200)
  IN db1, db2, db3, db4;
```

In this example, the database server fragments the table according to the following allocations.

Storage Space	Holds Values	Storage Space	Holds Values
db1	100 <= Col1 < 125	db3	150 <= Col1 < 175
db2	125 <= Col1 < 150	db4	175 <= Col1 < 200

The previous table shows allocations that can also be made with an expression-based fragmentation scheme:

```
CREATE TABLE ... FRAGMENT BY EXPRESSION
  Col1 >= 100 AND Col1 < 125 IN db1
  Col1 >= 125 AND Col1 < 150 IN db2
  Col1 >= 150 AND Col1 < 175 IN db3
  Col1 >= 175 AND Col1 < 200 IN db4;
```

As the examples show, the range-fragmentation example requires much less coding to achieve the same results. The same is true for the hybrid-range fragmentation compared to hybrid-expression fragmentation methods.

Column-Major-Range Allocation:

The following example demonstrates the syntax for column-major-range allocation, a hybrid-range fragmentation strategy:

```
CREATE TABLE tab2 (col2 INT, colx char (5))
  FRAGMENT BY HYBRID
  ( RANGE (col2 MIN 100 MAX 220))
  RANGE (col2)
  IN dbs11, dbs12, dbs13;
```

This type of fragmentation creates a distribution across dbslices and provides a further subdivision within each dbslice (across the dbspaces in the dbslice) such that when a query specifies a value for col1 (for example, WHERE col1 = 127), the query uniquely identifies a dbspace. To take advantage of the additional subdivision, you must specify more than one dbslice.

Row-Major-Range Allocation:

The following example demonstrates the syntax for row-major-range allocation, a hybrid-range fragmentation strategy:

```
CREATE TABLE tab3 (col3 INT, colx CHAR (5))
  FRAGMENT BY HYBRID
  ( RANGE (col3) )
  RANGE (col3 MIN 100 MAX 220)
  IN dbs11, dbs12, dbs13;
```

This fragmentation strategy is the counterpart to the column-major-range allocation. The advantages and restrictions are equivalent.

Independent-Range Allocation:

The following example demonstrates the syntax for an independent-range allocation, a hybrid-range fragmentation strategy:

```
CREATE TABLE tab4 (col4 INT, colx CHAR (5), col5 INT)
  FRAGMENT BY HYBRID
  ( RANGE (col4 MIN 100 MAX 200) )
  RANGE (col5 MIN 500 MAX 800)
  IN dbs11, dbs12, dbs13;
```

In this type of range fragmentation, the two columns are independent, and therefore the range allocations are independent. The range allocation for a dbspace on both columns is the conjunctive combination of the range allocation on each of the two independent columns.

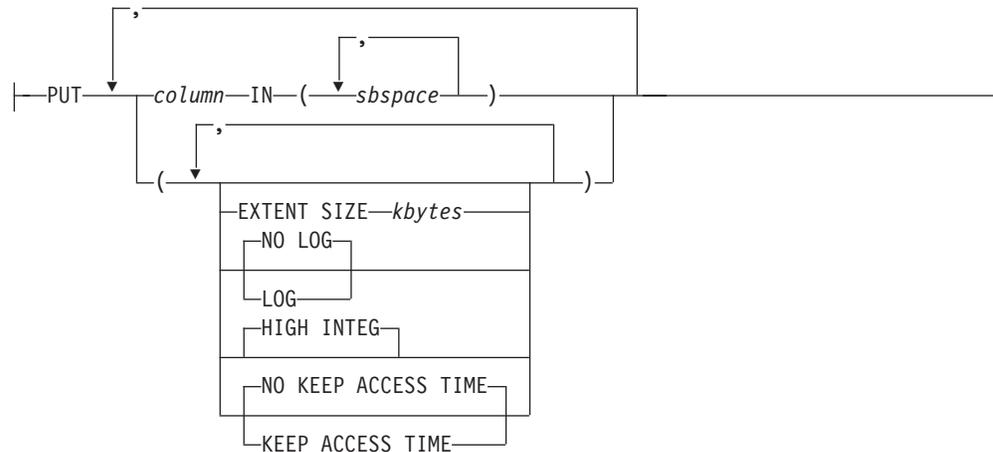
This type of fragmentation does not provide subdivisions within either column. With this type of fragmentation, a query that specifies values for both columns

(such as, WHERE col4 = 128 and col5 = 650) uniquely identifies the dbspace at the intersection of the two dbslices identified by the columns independently.

PUT Clause (IDS)

Use the PUT clause to specify the storage spaces and characteristics for each column that will contain smart large objects.

PUT Clause:



Element	Description	Restrictions	Syntax
<i>column</i>	Column to store in <i>sbspace</i>	Must contain a BLOB, CLOB, user-defined, or complex data type	"Identifier " on page 5-21
<i>kbytes</i>	Number of kilobytes to allocate for the extent size	Must be an integer value	"Literal Number " on page 4-185
<i>sbspace</i>	Name of an area of storage	Must exist	"Identifier " on page 5-21

The column cannot be in the form *column.field*. That is, the smart large object that you are storing cannot be one field of a row type.

A smart large object is contained in a single *sbspace*. The SBSPACENAME configuration parameter specifies the system default *sbspace* in which smart large objects are created, unless you specify another storage area.

Specifying more than one *sbspace* distributes the smart large objects in a round-robin distribution scheme, so that the number of smart large objects in each space is approximately equal. The **syscolattribs** system catalog table contains one row for each *sbspace* that you specify in the PUT clause.

When you fragment smart large objects across different *sbspaces*, you can work with smaller *sbspaces*. If you limit the size of an *sbspace*, backup and archive operations can perform more quickly. For an example that uses the PUT clause, see "Alternative to Full Logging" on page 2-223.

The following storage options are available to store BLOB and CLOB data:

Option Effect

EXTENT SIZE

Specifies how many kilobytes in a smart-large-object extent. The database server might round the EXTENT SIZE up so that the extents are multiples of the sbspace page size.

HIGH INTEG

Produces user-data pages that contain a page header and a page trailer to detect incomplete writes and data corruption. This is the default data-integrity behavior.

KEEP ACCESS TIME

Records, in the smart-large-object metadata, the system time when the smart large object was last read or written.

LOG Follows the logging procedure used with the current database log for the corresponding smart large object. This option can generate large amounts of log traffic and increase the risk of filling the logical log. (See also “Alternative to Full Logging.”)

NO KEEP ACCESS TIME

Does not record the system time when the smart large object was last read or written. This provides better performance than the KEEP ACCESS TIME option and is the default tracking behavior.

NO LOG

Turns off logging. This option is the default behavior.

If a user-defined or complex data type contains more than one large object, the specified large-object storage options apply to all large objects in the type unless the storage options are overridden when the large object is created.

Important: The PUT clause does not affect the storage of simple-large-object data types (BYTE and TEXT). For information on how to store BYTE and TEXT data, see “Large-Object Data Types” on page 4-29.

Alternative to Full Logging

Instead of full logging, you can turn off logging when you load the smart large object initially and then turn logging back on once the object is loaded.

Use the NO LOG option to turn off logging. If you use NO LOG, you can restore the smart-large-object metadata later to a state in which no structural inconsistencies exist. In most cases, no transaction inconsistencies will exist either, but that result is not guaranteed.

The following statement creates the **greek** table. Data values for the table are fragmented into the **db1** and **db2** dbspaces. The PUT clause assigns the smart-large-object data in the **gamma** and **delta** columns to the **sb1** and **sb2** sbspaces, respectively. The TEXT data values in the **eps** column are assigned to the **blb1** blobspace.

```
CREATE TABLE greek
(alpha INTEGER,
 beta VARCHAR(150),
 gamma CLOB,
 delta BLOB,
 eps TEXT IN blb1)
FRAGMENT BY EXPRESSION
alpha <= 5 IN db1, alpha > 5 IN db2
PUT gamma IN (sb1), delta IN (sb2);
```

EXTENT SIZE Options

The EXTENT SIZE options can define the size of extents assigned to the table.

EXTENT SIZE Options:



Element	Description	Restrictions	Syntax
<i>first_kilobytes</i>	Length in kilobytes of the first extent for the table; default is 16.	Must return a positive number; maximum is the chunk size	"Expression" on page 4-40
<i>next_kilobytes</i>	Length in kilobytes of each subsequent extent; default is 16.	Must return a positive number; maximum is the chunk size	"Expression" on page 4-40

The minimum length of *first_kilobytes* (and of *next_kilobytes*) is four times the disk-page size on your system. For example, if you have a 2-kilobyte page system, the minimum length is 8 kilobytes.

The next example specifies a first extent of 20 kilobytes and allows the rest of the extents to use the default size:

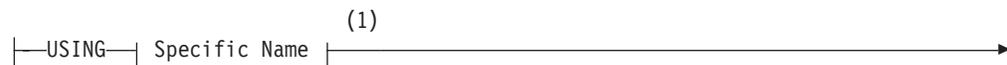
```
CREATE TABLE emp_info
(
  f_name    CHAR(20),
  l_name    CHAR(20),
  position  CHAR(20),
  start_date DATETIME YEAR TO DAY,
  comments  VARCHAR(255)
)
EXTENT SIZE 20;
```

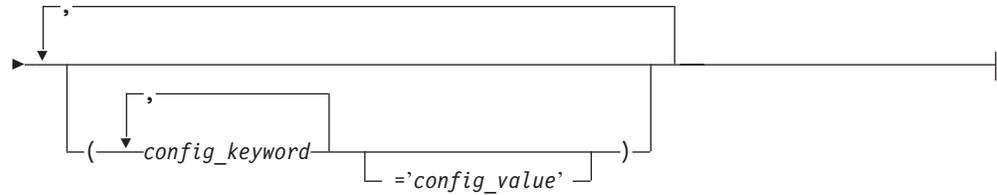
If you need to revise the extent sizes of a table, you can modify the extent and next-extent sizes in the generated schema files of an unloaded table. For example, to make a database more efficient, you might unload a table, modify the extent sizes in the schema files, and then create and load a new table. For information about how to optimize extents, see your *IBM Informix Administrator's Guide*.

USING Access-Method Clause (IDS)

The USING Access Method clause can specify an access method.

USING Access-Method Clause:





Notes:

- 1 See "Specific Name " on page 5-73

Element	Description	Restrictions	Syntax
<i>config_keyword</i>	Configuration keyword associated with the specified access method	No more than 18 bytes. The access method must exist.	Literal keyword
<i>config_value</i>	Value of the specified configuration keyword	No more than 236 bytes. Must be defined by the access method.	"Quoted String" on page 4-189

A *primary-access method* is a set of routines to perform DDL and DML operations, such as create, drop, insert, delete, update, and scan, to make a table available to the database server. Dynamic Server provides a built-in primary-access method.

You store and manage a virtual table either outside of the database server in an extspace or inside the database server in an sbspace. (See "Storage Options" on page 2-210.) You can access a virtual table with SQL statements. Access to a virtual table requires a user-defined primary-access method.

DataBlade modules can provide other primary-access methods to access virtual tables. When you access a virtual table, the database server calls the routines associated with that access method rather than the built-in table routines. For more information on these other primary-access methods, refer to your access-method documentation.

You can retrieve a list of configuration values for an access method from a table descriptor (**mi_am_table_desc**) using the MI_TAB_AMPARAM macro. Not all keywords require configuration values.

The access method must already exist. For example, if an access method called **textfile** exists, you can specify it with the following syntax:

```

CREATE TABLE mybook
  (...)
  IN myextspace
  USING textfile (DELIMITER=':');

```

LOCK MODE Options

Use the LOCK MODE options to specify the locking granularity of the table.

LOCK MODE Options:



Notes:

- 1 Extended Parallel Server only

You can subsequently change the lock mode of the table with the ALTER TABLE statement.

Granularity	Effect
PAGE	Obtains and releases one lock on a whole page of rows This is the default locking granularity. Page-level locking is especially useful when you know that the rows are grouped into pages in the same order that you are using to process all the rows. For example, if you are processing the contents of a table in the same order as its cluster index, page locking is appropriate.
ROW	Obtains and releases one lock per row Row-level locking provides the highest level of concurrency. If you are using many rows at one time, however, the lock-management overhead can become significant. You might also exceed the maximum number of locks available, depending on the configuration of your database server, but Dynamic Server can support up to 18 million locks on 32-bit platforms, or 600 million locks on 64-bit platforms. Only tables with row-level locking can support the LAST COMMITTED isolation level feature.
TABLE (XPS)	Places a lock on the entire table This type of lock reduces update concurrency compared to row and page locks. A table lock reduces the lock-management overhead for the table. With table locking, multiple read-only transactions can still access the table.

The following table describes the locking-granularity options available.

Precedence and Default Behavior

In Dynamic Server, you do not need to specify the lock mode each time you create a new table. You can globally set the locking granularity of all new tables in the following environments:

- Database session of an individual user
You can set the **IFX_DEF_TABLE_LOCKMODE** environment variable to specify the lock mode of new tables during your current session.
- Database server (all sessions on the database server)
If you are a DBA, you can set the **DEF_TABLE_LOCKMODE** configuration parameter in the ONCONFIG file to determine the lock mode of all new tables in the database server.
If you are not a DBA, you can set the **IFX_DEF_TABLE_LOCKMODE** environment variable for the database server, before you run **oninit**, to specify the lock mode of all new tables of the database server.

The LOCK MODE setting in a CREATE TABLE statement takes precedence over the settings of the **IFX_DEF_TABLE_LOCKMODE** environment variable and the **DEF_TABLE_LOCKMODE** configuration parameter.

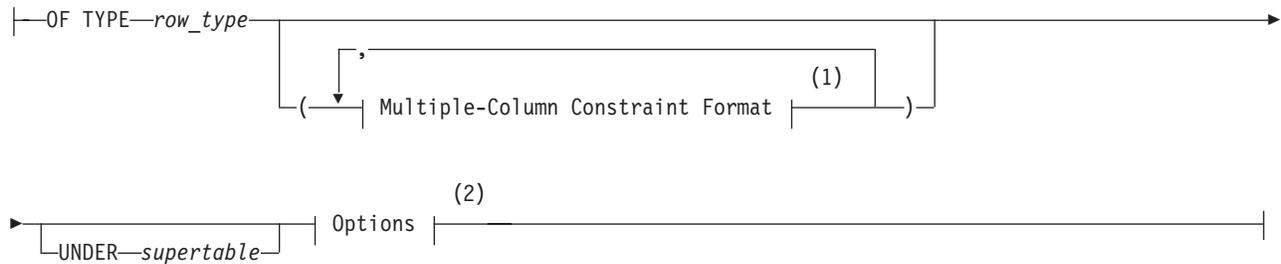
If CREATE TABLE specifies no lock mode setting, the default mode depends on the setting of the **IFX_DEF_TABLE_LOCKMODE** environment variable or the **DEF_TABLE_LOCKMODE** configuration parameter. For information about **IFX_DEF_TABLE_LOCKMODE**, refer to the *IBM Informix Guide to SQL: Reference*.

For information about the DEF_TABLE_LOCKMODE configuration parameter, refer to the *IBM Informix Administrator's Reference*.

OF TYPE Clause (IDS)

Use the OF TYPE clause to create a *typed table* for an object-relational database. A typed table is a table to which you assign a named ROW data type.

OF TYPE Clause:



Notes:

- 1 See "Multiple-Column Constraint Format" on page 2-205
- 2 See "Options Clauses" on page 2-209

Element	Description	Restrictions	Syntax
<i>row_type</i>	Name of the ROW type on which this table is based	Must be a named ROW data type registered in the local database	"Identifier " on page 5-21
<i>supertable</i>	Name of the table from which this table inherits its properties	Must already exist as a typed table	"Identifier " on page 5-21

If you use the UNDER clause, the *row_type* must be derived from the ROW type of the *supertable*. A type hierarchy must already exist in which the named ROW type of the new table is a subtype of the named ROW type of the *supertable*.

Jagged rows are any set rows from a table hierarchy in which the number of columns is not fixed among the typed tables within the hierarchy. Some APIs, such as Informix ESQL/C and Informix JDBC Driver, do not support queries that return jagged rows.

When you create a typed table, CREATE TABLE cannot specify names for its columns, because the column names were declared when you created the ROW type. Columns of a typed table correspond to the fields of the named ROW type. The ALTER TABLE statement cannot add additional columns to a typed table.

For example, suppose you create a named ROW type, **student_t**, as follows:

```

CREATE ROW TYPE student_t
  (name      VARCHAR(30),
   average   REAL,
   birthdate DATETIME YEAR TO DAY)
  
```

If a table is assigned the type **student_t**, the table is a typed table whose columns are of the same name and data type, and in the same order, as the fields of the type **student_t**. For example, the following CREATE TABLE statement creates a typed table named **students** whose type is **student_t**:

```
CREATE TABLE students OF TYPE student_t
```

The **students** table has the following columns:

```
name      VARCHAR(30)
average   REAL
birthdate DATETIME
```

For more information about ROW types, refer to the CREATE ROW TYPE statement on page 1-194.

Using Large-Object Data in Typed Tables

Use the BLOB or CLOB instead of BYTE or TEXT data types when you create a typed table that contains columns for large objects. For backward compatibility, you can create a named-ROW type that contains BYTE or TEXT fields and use that ROW type to re-create an existing (untyped) table as a typed table. Although you can use a ROW type that contains BYTE or TEXT fields to create a typed table, such a ROW type is not valid as a column. You can, however, use a ROW type that contains BLOB or CLOB fields both in typed tables and in columns.

Using the UNDER Clause

Use the UNDER clause to specify inheritance (that is, define the new table as a subtable). The subtable inherits properties from the specified supertable. In addition, you can define new properties specific to the subtable.

Continuing the example shown in “OF TYPE Clause (IDS)” on page 2-227, the following statements create a typed table, **grad_students**, that inherits all of the columns of the **students** table but also has columns for **adviser** and **field_of_study** that correspond to fields in the **grad_student_t** ROW type:

```
CREATE ROW TYPE grad_student_t
(adviser      CHAR(25),
 field_of_study CHAR(40)) UNDER student_t;

CREATE TABLE grad_students OF TYPE grad_student_t UNDER students;
```

When you use the UNDER clause, the subtable inherits these properties:

- All columns in the supertable
- All constraints defined on the supertable
- All indexes defined on the supertable
- All triggers defined on the supertable.
- All referential integrity constraints
- The access method
- The storage option specification (including fragmentation strategy)
If a subtable defines no fragments, but if its supertable has fragments defined, then the subtable inherits the fragments of the supertable.

Tip: Any heritable attributes that are added to a supertable after subtables have been created are automatically inherited by existing subtables. You do not need to add all heritable attributes to a supertable before you create its subtables.

Restrictions on Table Hierarchies: Inheritance occurs in one direction only, namely from supertable to subtable. Properties of subtables are *not* inherited by supertables. The section “System Catalog Information” on page 2-230 lists the inherited database objects for which the system catalog maintains no information regarding subtables.

No two tables in a table hierarchy can have the same data type. For example, the final line of the next code example is invalid, because the tables **tab2** and **tab3** cannot have the same row type (**rowtype2**):

```
create row type rowtype1 (...);
create row type rowtype2 (...) under rowtype1;
create table tab1 of type rowtype1;
create table tab2 of type rowtype2 under tab1;
create table tab3 of type rowtype2 under tab1; -- This is not valid.
```

Access Privileges on Tables

The privileges on a table describe both who can access the information in the table and who can create new tables. For more information about access privileges, see the description of the “GRANT ” on page 2-395 statement.

In an ANSI-compliant database, no default table-level privileges exist. You must grant these privileges explicitly.

Setting the environment variable **NODEFDAC** to yes prevents default privileges from being granted to PUBLIC on new tables in a database that is not ANSI compliant, as described in the *IBM Informix Guide to SQL: Reference*. For more information about privileges, see the *IBM Informix Guide to SQL: Tutorial*.

The SECURITY POLICY Clause (IDS)

The optional Security Policy clause can use the following syntax to specify a security policy that is associated with the table. .

SECURITY POLICY Clause:

```
|-----|
| SECURITY POLICY—policy |
```

Element	Description	Restrictions	Syntax
<i>policy</i>	Name of a security policy	Must exist in the database	“Identifier ” on page 5-21

Only DBSECAM can create a table that includes the Security Policy clause to specify a security policy for the table.

The following guidelines apply to tables that can be protected by including a valid SECURITY POLICY clause in the CREATE TABLE statement:

- A table is not protected unless it has a security policy associated with it and has either rows secured or at least one column secured. The former indicates that the table is a protected table with row level granularity and the latter indicates that the table is a protected table with column level granularity.
- Securing rows with the IDSSECURITYLABEL column clause fails if the table does not have a security policy associated with it.
- Securing a column with the COLUMN SECURED WITH clause fails if the table does not have a security policy associated with it.
- A table can have at most one security policy.
- A table can have any number of protected columns, and each protected column can have a different label, or several protected columns can share the same label.
- A security policy cannot be associated with a temporary table nor with a typed table in a table hierarchy.

System Catalog Information

When you create a table, the database server adds information about the table to the **sysables** system catalog table, and column information to **syscolumns** system catalog table. The **sysfragments** system catalog table contains information about fragmentation strategies and the location of fragments. The **sysblobs** system catalog table contains information about the location of dbspaces and of simple large objects. (The **syschunks** table in the **sysmaster** database contains information about the location of smart large objects.)

The **sysabauth**, **syscolauth**, **sysfragauth**, **sysprocauth**, **sysusers**, and **sysxdtpeauth** tables contain information about the privileges that various CREATE TABLE options require. The **sysables**, **sysxdtypes**, and **sysinherits** system catalog tables provide information about table types.

Constraints, indexes, and triggers are recorded in the system catalog for the supertable, but not for subtables that inherit them. Fragmentation information, however, is recorded for both supertables and subtables. For more information about inheritance, refer to the *IBM Informix Guide to SQL: Tutorial*.

Related Information

Related statements: "ALTER TABLE" on page 2-44, "CREATE INDEX " on page 2-124, "CREATE DATABASE " on page 2-97, "CREATE EXTERNAL TABLE (XPS) " on page 2-104, "CREATE ROW TYPE " on page 2-168, "CREATE Temporary TABLE " on page 2-231, "DROP TABLE " on page 2-342, "SET Database Object Mode" on page 2-591, and "SET Transaction Mode " on page 2-661

For Extended Parallel Server, see also "SET Default Table Type " on page 2-602 and "SET Default Table Space " on page 2-601.

For discussions of database and table creation, including discussions on data types, data-integrity constraints, and tables in hierarchies, see the *IBM Informix Database Design and Implementation Guide*.

For information about the system catalog tables that store information about objects in the database, see the *IBM Informix Guide to SQL: Reference*. For information about the **syschunks** table (in the **sysmaster** database) that contains information about the location of smart large objects, see your *IBM Informix Administrator's Reference*.

CREATE TEMP TABLE

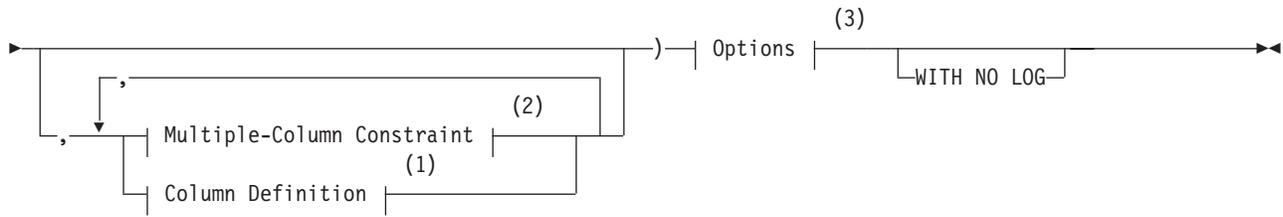
Use the CREATE TEMP TABLE statement to create a temporary table in the current database.

Syntax

```
►► CREATE TEMP TABLE table ( 

|                   |     |
|-------------------|-----|
| Column Definition | (1) |
|-------------------|-----|

 )
```



Notes:

- 1 See "CREATE Temporary TABLE "
- 2 See " Multiple-Column Constraint Format" on page 2-234
- 3 See " Temporary Table Options" on page 2-235

Element	Description	Restrictions	Syntax
<i>table</i>	Name declared here for a temporary table	Must be unique among the names of temporary tables in the same session	"Identifier " on page 5-21

Usage

The CREATE TEMP TABLE statement is a special case of the CREATE Temporary TABLE statement. The CREATE Temporary TABLE statement can also create a SCRATCH table in an Extended Parallel Server database.

For the complete syntax and semantics of the CREATE TEMP TABLE statement, see "CREATE Temporary TABLE ."

CREATE Temporary TABLE

Use the CREATE Temporary TABLE statement to create a temporary table in the current database.

Syntax

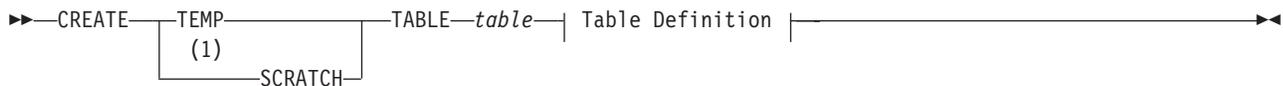
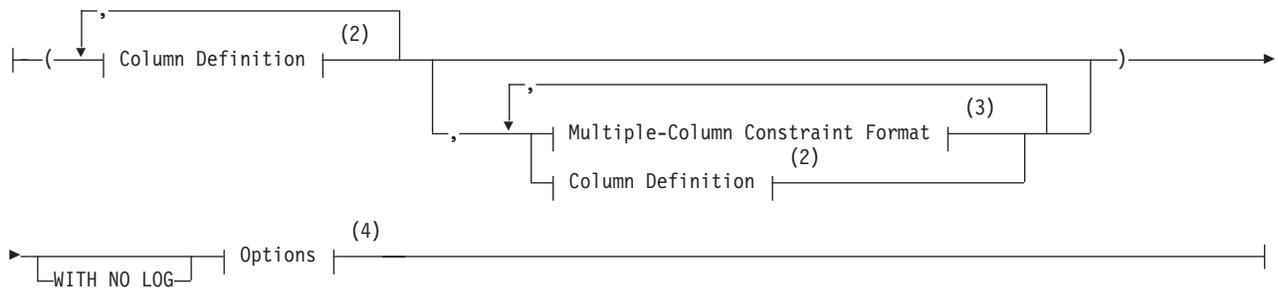


Table Definition:



Notes:

- 1 Extended Parallel Server only
- 2 See “CREATE Temporary TABLE ” on page 2-231
- 3 See “ Multiple-Column Constraint Format” on page 2-234
- 4 See “ Temporary Table Options” on page 2-235

Element	Description	Restrictions	Syntax
<i>table</i>	Name declared here for a table	Must be unique in session	“Identifier ” on page 5-21

Usage

You must have the Connect privilege on the database to create a temporary table. The temporary table is visible only to the user who created it.

In DB–Access, using the CREATE Temporary Table statement outside the CREATE SCHEMA statement generates warnings if you set **DBANSIWARN**.

In ESQ/C, the CREATE Temporary TABLE statement generates warnings if you use the **-ansi** flag or set the **DBANSIWARN** environment variable.

Using the TEMP Option

When you create a TEMP table, you can build indexes and constraints on the table.

Using the SCRATCH Option (XPS)

Use the SCRATCH keyword to reduce the overhead of transaction logging. A scratch table is a nonlogging temporary table that does not support indexes or referential constraints. A scratch table is identical to a TEMP table created with the WITH NO LOG option. Operations on scratch tables are not included in transaction-log operations.

Naming a Temporary Table

A temporary table is associated with a session, not with a database. When you create a temporary table, you cannot create another temporary table with the same name (even for another database) until you drop the first temporary table or end the session. The name must be different from the name of any other table, view, sequence object, or synonym in the current database, but it need not be different from the temporary table names that are declared by other users.

In an ANSI-compliant database, the combination *owner.table* must be unique in the database.

Using the WITH NO LOG Option

Use the WITH NO LOG option to reduce the overhead of transaction logging. If you specify WITH NO LOG, operations on the temporary table are not included in the transaction-log operations. The WITH NO LOG option is required on all temporary tables that you create in temporary dbspaces.

In Dynamic Server, if the ONCONFIG parameter TEMPTAB_NOLOG is set to 1, logging of temporary tables is disabled and all temporary tables are non-logging by default. This setting can improve the performance of operations that use temporary tables, such as HDR operations. The WITH NO LOG option is not

needed when TEMPTAB_NOLOG has disabled logging of temporary tables. For information about how to set the TEMPTAB_NOLOG parameter, see your *IBM Informix Administrator's Reference*.

If you use the WITH NO LOG option in a database that does not use logging, the WITH NO LOG keywords have no effect. If your database does not support transaction logging, every table behaves as if the WITH NO LOG option were specified.

Once you turn off logging on a temporary table, you cannot turn it back on; a temporary table, therefore, is either always logged or else never logged.

The following temporary table is not logged in a database that uses logging:

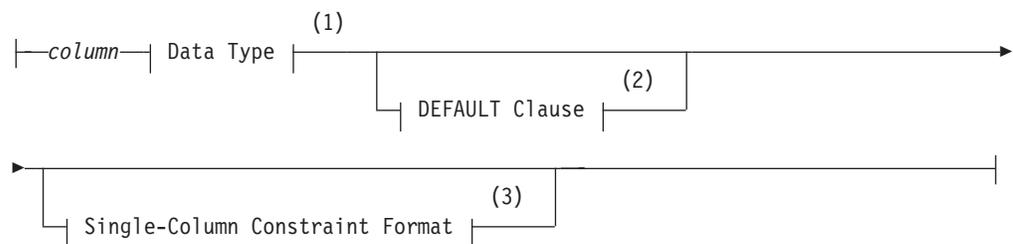
```
CREATE TEMP TABLE tab2 (fname CHAR(15), lname CHAR(15))
    WITH NO LOG;
```

In Extended Parallel Server, you should use a SCRATCH table rather than a TEMP..WITH NO LOG table. The behavior of a temporary table that you create with the WITH NO LOG option is the same as that of a SCRATCH table.

Column Definition

Use the Column Definition segment of CREATE Temporary TABLE to list the name, data type, default value, and constraints of a single column.

Column Definition:



Notes:

- 1 See "Data Type " on page 4-21
- 2 See "DEFAULT Clause" on page 2-195
- 3 See " Single-Column Constraint Format" on page 2-234

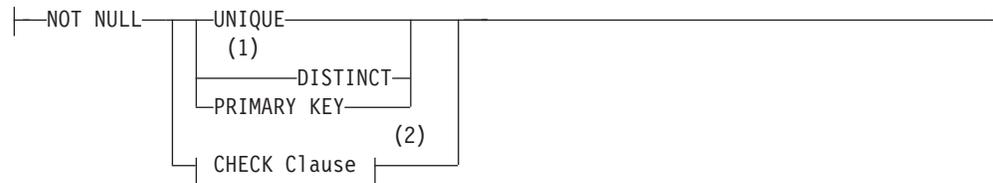
Element	Description	Restrictions	Syntax
<i>column</i>	Name declared here for a column in the table	Must be unique in its table	"Identifier " on page 5-21

This portion of the CREATE Temporary TABLE statement is almost identical to the corresponding section in the CREATE TABLE statement. The difference is that fewer types of constraints are allowed in a temporary table.

Single-Column Constraint Format

Use the single-column constraint format to create one or more data-integrity constraints for a single column in a temporary table.

Single-Column Constraint Format:



Notes:

- 1 Informix extension
- 2 See "CHECK Clause" on page 2-202

This is a subset of the syntax of "Single-Column Constraint Format" on page 2-197 that the CREATE TABLE statement supports.

You can find detailed discussions of specific constraints in these sections.

Constraint

For more information, see

CHECK

"CHECK Clause" on page 2-202

DISTINCT

"Using UNIQUE or DISTINCT Constraints" on page 2-198

NOT NULL

"Using the NOT NULL Constraint" on page 2-198

PRIMARY KEY

"Using the PRIMARY KEY Constraint" on page 2-199

UNIQUE

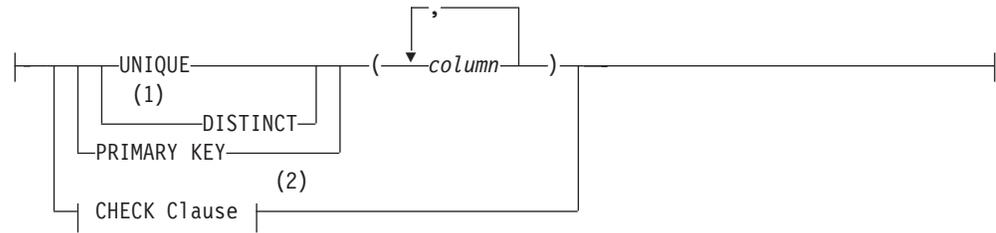
"Using UNIQUE or DISTINCT Constraints" on page 2-198

Constraints that you define on temporary tables are always enabled.

Multiple-Column Constraint Format

Use the multiple-column constraint format to associate one or more columns with a constraint. This alternative to the single-column constraint format allows you to associate multiple columns with a constraint.

Multiple-Column Constraint Format:



Notes:

- 1 Informix extension
- 2 See "CHECK Clause" on page 2-202

Element	Description	Restrictions	Syntax
<i>column</i>	Name of column or columns on which the constraint is placed	Must be unique in a table, but the same name can be in different tables of the same database	"Identifier " on page 5-21

This is a subset of the syntax of "Multiple-Column Constraint Format" on page 2-205 that the CREATE TABLE statement supports.

This alternative to the single-column constraint segment of CREATE Temporary TABLE can associate multiple columns with a constraint. Constraints that you define on temporary tables are always enabled.

The following table indicates where you can find detailed discussions of specific constraints.

Constraint	For more information, see	For an example, see
CHECK	"CHECK Clause" on page 2-202	"Defining Check Constraints Across Columns" on page 2-207
DISTINCT	"Using UNIQUE or DISTINCT Constraints" on page 2-198	"Examples of the Multiple-Column Constraint Format" on page 2-207
PRIMARY KEY	"Using the PRIMARY KEY Constraint" on page 2-199	"Defining Composite Primary and Foreign Keys" on page 2-208
UNIQUE	"Using UNIQUE or DISTINCT Constraints" on page 2-198	"Examples of the Multiple-Column Constraint Format" on page 2-207

See also the section "Differences Between a Unique Constraint and a Unique Index" on page 2-199.

Temporary Table Options

The Options clauses of CREATE Temporary TABLE let you specify storage locations, locking modes, and user-defined access methods. You cannot specify the size of the initial or next extents for a temporary table. Extents for a temporary table always have a size of eight pages.

Options:





Notes:

- 1 Dynamic Server only
- 2 Informix extension
- 3 See "Storage Options"
- 4 See "LOCK MODE Options" on page 2-225
- 5 See "USING Access-Method Clause (IDS)" on page 2-224

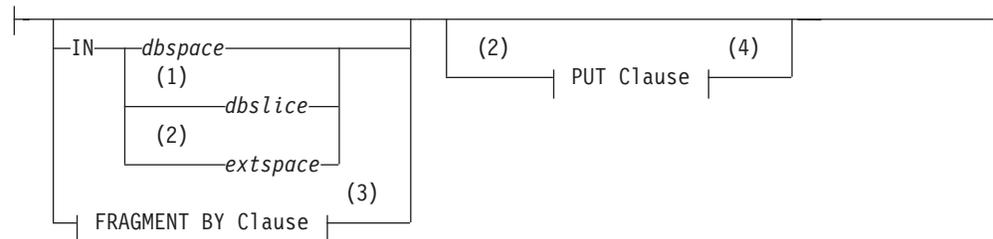
This is a subset of the syntax of "Options Clauses" on page 2-209 that the CREATE TABLE statement supports.

For an explanation of the effect of the WITH CRCOLS keywords for Enterprise Replication, see the section "Using the WITH CRCOLS Option (IDS)" on page 2-209 in the description of the CREATE TABLE statement.

Storage Options

Use the Storage Options segment of the CREATE Temporary Table statement to specify the storage location and distribution scheme for the table. This is an extension to the ANSI/ISO standard for SQL syntax. Unlike the corresponding Storage Options segment of the CREATE TABLE statement, no EXTENT SIZE options are supported for temporary tables.

Storage Options:



Notes:

- 1 Extended Parallel Server only
- 2 Dynamic Server only
- 3 See "FRAGMENT BY Clause" on page 2-212
- 4 See "PUT Clause (IDS)" on page 2-222

Element	Description	Restrictions	Syntax
<i>dbspace</i>	Dbospace in which to store the table.	Must already exist	"Identifier " on page 5-21
<i>dbslice</i>	Name of the dbslice in which to store the table	Must already exist	"Identifier " on page 5-21
<i>extspace</i>	Name that onspaces assigned to a storage area outside the database server	Must already exist	See documentation for access method.

To create a fragmented, unique index on a temporary table, you must specify an explicit expression-based distribution scheme for a temporary table in the CREATE Temporary TABLE statement.

If you are using Extended Parallel Server, you can fragment a temporary table across multiple dbspaces that different coservers manage.

Where Temporary Tables are Stored

The distribution scheme that you specify with the CREATE Temporary TABLE statement (either with the IN clause or the FRAGMENT BY clause) takes precedence over the information that the **DBSPACETEMP** environment variable or the DBSPACETEMP configuration parameter specifies.

If you do not specify an explicit distribution scheme for a temporary table, its storage location depends on the **DBSPACETEMP** (or DBSPACETEMP) setting.

- If **DBSPACETEMP** and DBSPACETEMP are not set, all temporary tables are created without fragmentation in the same dspace where the database was created (or in **rootdbs**, if the database was not created in another dspace).
- If only one dspace for temporary tables is specified by **DBSPACETEMP** (or by DBSPACETEMP, if **DBSPACETEMP** is not set), all temporary tables are created without fragmentation in the specified dspace.
- If **DBSPACETEMP** (or DBSPACETEMP, if **DBSPACETEMP** is not set) specifies two or more dbspaces for temporary tables, then each temporary table is created in one of the specified dbspaces.

In a non-logging database, each temporary table is created in a temporary dspace; in a databases that support transaction logging, the temporary table is created in a standard dspace. The database server keeps track of which of these dbspaces was most recently used, and when it receives the next request to allocate temporary storage, the database server uses the next available dspace (in a round-robin pattern) in order to allocate I/O operations approximately evenly among those dbspaces.

For example, if you create three temporary tables in a database with logging where **DBSPACETEMP** specifies **tempc1**, **tempc2**, and **tempc3** as the default dbspaces for temporary tables, then the first table is created in the dspace called **tempc1**, the second table is created in **tempc2**, and the third one is created in **tempc3**, if these are the only requests for temporary storage.

Temporary tables created with SELECT... INTO TEMP or SELECT... INTO SCRATCH also behave this way, so that **DBSPACETEMP** (or DBSPACETEMP) settings that specify multiple dbspaces result in round-robin fragmentation.

For more information on the **DBSPACETEMP** environment variable, see the *IBM Informix Guide to SQL: Reference*. For more information on the DBSPACETEMP configuration parameter, see your *IBM Informix Administrator's Reference*.

The following example shows how to insert data into a temporary table called **result_tmp** to output to a file the results of a user-defined function (**f_one**) that returns multiple rows:

```
CREATE TEMP TABLE result_tmp( ... );
INSERT INTO result_tmp EXECUTE FUNCTION f_one();
UNLOAD TO 'file' SELECT * FROM temp1;
```

In Extended Parallel Server, to re-create this example, use the EXECUTE PROCEDURE statement instead of the EXECUTE FUNCTION statement.

Differences between Temporary and Permanent Tables

Compared to permanent tables, temporary tables differ in these ways:

- They have fewer types of constraints available.
- They have fewer options that you can specify.
- They are not visible to other users or sessions.
- They do not appear in the system catalog tables.
- They are not preserved.

For more information, see “Duration of Temporary Tables.”

In Extended Parallel Server, you can use the following data definition statements on a temporary table from a secondary coserver: CREATE Temporary TABLE, CREATE INDEX, CREATE SCHEMA, DROP TABLE, and DROP INDEX.

The INFO statement and the **Info Menu** option of DB–Access cannot reference temporary tables.

Duration of Temporary Tables

The duration of a temporary table depends on whether or not it is logged.

A logged temporary table exists until one of the following situations occurs:

- The application disconnects.
- A DROP TABLE statement is issued on the temporary table.
- The database is closed.

When any of these events occur, the temporary table is deleted.

Nonlogging temporary tables include tables that were created using the WITH NO LOG option of CREATE TEMP TABLE and all SCRATCH tables.

A nonlogging temporary table exists until one of the following events occurs:

- The application disconnects.
- A DROP TABLE statement is issued on the temporary table.

Because these tables do not disappear when the database is closed, you can use a nonlogging temporary table to transfer data from one database to another while the application remains connected.

Related Information

Related statements: “ALTER TABLE” on page 2-44, “CREATE TABLE ” on page 2-191, “CREATE DATABASE ” on page 2-97, “DROP TABLE ” on page 2-342, and “SELECT” on page 2-528

If you use Extended Parallel Server, see also “SET Default Table Type ” on page 2-602 and “SET Default Table Space ” on page 2-601.

For additional information about the **DBANSIWARN** and **DBSPACETEMP** environment variables, refer to the *IBM Informix Guide to SQL: Reference*.

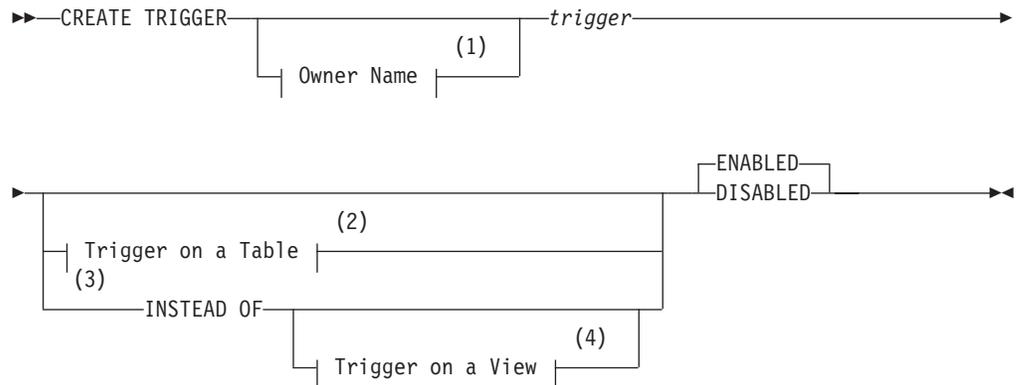
For additional information about the ONCONFIG parameter DBSPACETEMP, see your *IBM Informix Administrator’s Guide*.

CREATE TRIGGER

Use the CREATE TRIGGER statement to define a trigger on a table.

This is an extension to the ANSI/ISO standard for SQL. For Dynamic Server, you can also use CREATE TRIGGER to define an INSTEAD OF trigger on a view.

Syntax



Notes:

- 1 See "Owner Name" on page 5-45
- 2 See "Defining a Trigger Event and Action" on page 2-240
- 3 Dynamic Server only
- 4 See "INSTEAD OF Triggers on Views (IDS)" on page 2-270

Element	Description	Restrictions	Syntax
<i>trigger</i>	Name that you declare here for a new trigger	Must be unique among the names of triggers in the current database	"Identifier" on page 5-21

Usage

A *trigger* is a database object that, unless disabled, automatically executes a specified set of SQL statements, called the *trigger action*, when a specified *trigger event* occurs.

The trigger event that initiates the trigger action can be an INSERT, DELETE, UPDATE, or (for triggers on IDS tables only) a SELECT statement. The event must specify the table or view on which the trigger is defined. (SELECT or UPDATE events for triggers on tables can also specify one or more columns.)

You can use the CREATE TRIGGER statement in two distinct ways:

- You can define a trigger on a table in the current database.
- For Dynamic Server, you can also define an INSTEAD OF trigger on a view in the current database.

Any SQL statement that is an instance of the trigger event is called a *triggering statement*. When the event occurs, triggers defined on tables and triggers defined on views differ in whether the triggering statement is executed:

- For tables, the trigger event and the trigger action both execute.

- For views, only the trigger action executes, instead of the event.

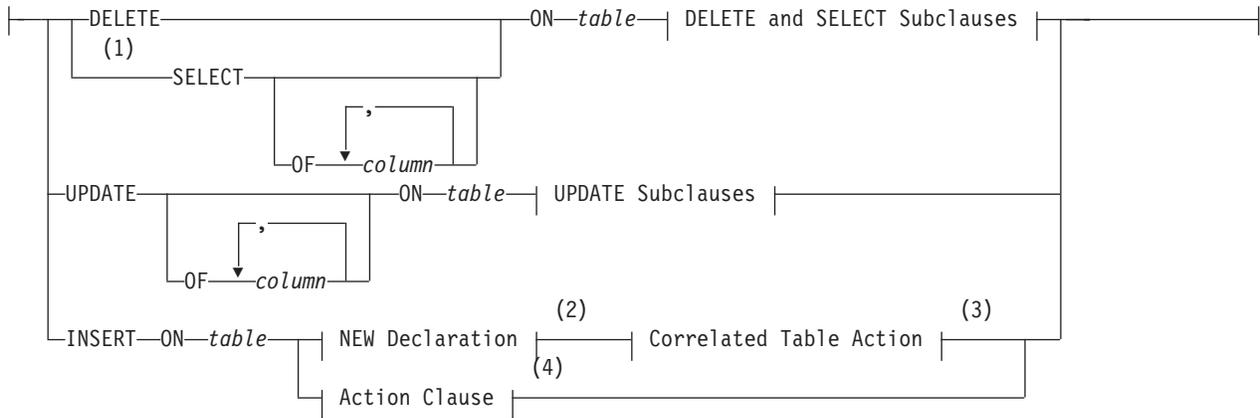
The CREATE TRIGGER statement can support the integrity of data in the database by defining rules by which specified DML operations (the triggering events) cause the database server to take specified actions. The following sections describe the syntax elements. Sections whose names are followed by "(IDS)" describe features that are not supported by Extended Parallel Server.

Clause	Page	Effect
Defining a Trigger Event and Actions	"Defining a Trigger Event and Action"	Associates triggered actions with an event
Trigger Modes (IDS)	" Trigger Modes (IDS)" on page 2-242	Enables or disables the trigger
Insert Events and Delete Events	"INSERT Events and DELETE Events" on page 2-246	Defines Insert events and Delete events
Update Events	"UPDATE Event" on page 2-246	Defines Update events
Select Events (IDS)	" SELECT Event (IDS)" on page 2-247	Defines Select events
Action Clause	" Action Clause" on page 2-251	Defines triggered actions
REFERENCING Clause for Delete	"REFERENCING Clause for Delete" on page 2-254	Declares qualifier for deleted values
REFERENCING Clause for Insert	"REFERENCING Clause for Insert" on page 2-255	Declares qualifier for inserted values
REFERENCING Clause for Update	"REFERENCING Clause for Update" on page 2-255	Declares qualifiers for old and new values
REFERENCING Clause for Select (IDS)	" REFERENCING Clause for Select (IDS)" on page 2-256	Declares qualifier for result set values
Correlated Table Action	" Correlated Table Action" on page 2-257	Defines triggered actions
Triggered Action	"Triggered Action" on page 2-257	Defines triggered actions
INSTEAD OF Trigger on Views (IDS)	"INSTEAD OF Triggers on Views (IDS)" on page 2-270	Defines a trigger on views
Action Clause of INSTEAD OF Triggers (IDS)	"The Action Clause of INSTEAD OF Triggers (IDS)" on page 2-271	Triggered actions on views

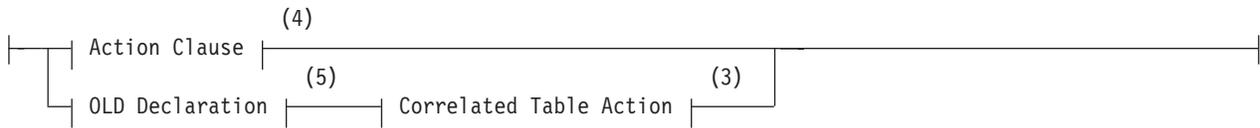
Defining a Trigger Event and Action

This syntax defines the event and action of a trigger on a table or on a view.

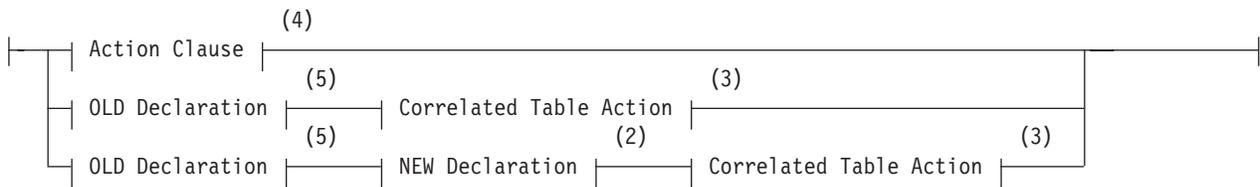
Trigger on a Table:



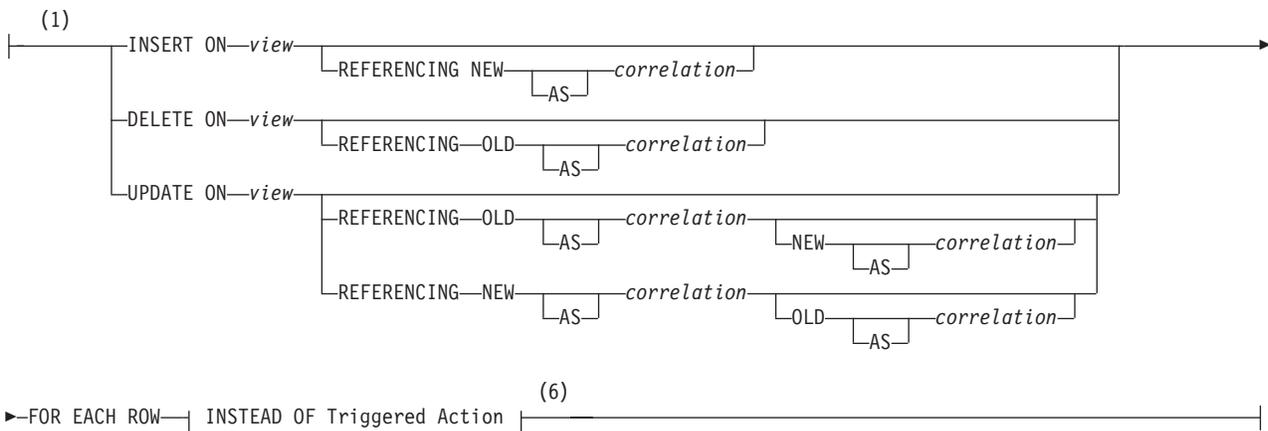
DELETE and SELECT Subclauses:



UPDATE Subclauses:



Trigger on a View:



Notes:

- 1 Dynamic Server only
- 2 See "REFERENCING Clause for Insert" on page 2-255
- 3 See "Correlated Table Action" on page 2-257
- 4 See "Action Clause" on page 2-251

- 5 See “REFERENCING Clause for Update” on page 2-255
- 6 See “INSTEAD OF Triggers on Views (IDS)” on page 2-270

Element	Description	Restrictions	Syntax
column	The name of a column in the triggering <i>table</i>	Must exist	“Identifier ” on page 5-21
correlation	Name that you declare here to qualify an old or new column value (as <i>correlation.column</i>) in a triggered action	Must be unique in this trigger	“Identifier ” on page 5-21
<i>table, view</i>	Name or synonym of the triggering table or view. The <i>table</i> or <i>view</i> can include an <i>owner</i> . qualifier.	Must exist in the current database	“Identifier ” on page 5-21

The left-hand portion of the main diagram (including the *table* or *view*) defines the *trigger event* (sometimes called the *triggering event*). The rest of the diagram declares correlation names and defines the *trigger action* (sometimes called the *triggered action*). (For triggers on tables, see “ Action Clause” on page 2-251 and “ Correlated Table Action” on page 2-257. For INSTEAD OF triggers on views, see “The Action Clause of INSTEAD OF Triggers (IDS)” on page 2-271.)

Restrictions on Triggers

To create a trigger on a table or a view, you must own the table or view, or have DBA privilege. For the relationship between the privileges of the trigger owner and those of other users, see “Privileges to Execute Trigger Actions” on page 2-266.

The table on which you create a trigger must exist in the current database. You cannot create a trigger on any of the following types of tables:

- A diagnostics table, a violations table, or a table in another database
- A raw table, a temporary table, or a system catalog table

Extended Parallel Server does not support triggers on a static or scratch tables, and does not support light appends on operational tables on which a trigger is defined. (Light appends are described in the *IBM Informix Administrator’s Guide*.)

For additional restrictions on INSTEAD OF triggers on views, see “Restrictions on INSTEAD OF Triggers on Views (IDS)” on page 2-272.

In DB–Access, if you want to define a trigger as part of a schema, place the CREATE TRIGGER statement inside a CREATE SCHEMA statement.

If you are embedding the CREATE TRIGGER statement in an Informix ESQL/C program, you cannot use a host variable in the trigger definition.

You can use the DROP TRIGGER statement to remove an existing trigger. If you use DROP TABLE or DROP VIEW to remove triggering tables or views from the database, all triggers on those tables or views are also dropped.

You cannot define a DELETE trigger on a table that has a referential constraint that specifies ON DELETE CASCADE.

Trigger Modes (IDS)

You can set a trigger mode to enable or disable a trigger when you create it.

Trigger Modes



You can create triggers on tables or on views in ENABLED or DISABLED mode.

- When a trigger is created in ENABLED mode, the database server executes the trigger action when the trigger event is encountered. (If you specify no mode when you create a trigger, ENABLED is the default mode.)
- When a trigger is created in DISABLED mode, the trigger event does not cause execution of the trigger action. In effect, the database server ignores the trigger and its action, even though the **systriggers** system catalog table maintains information about the disabled trigger.

You can also use the SET TRIGGERS option of the Database Object Mode statement to set an existing trigger to the ENABLED or DISABLED mode.

After a DISABLED trigger is enabled by the SET TRIGGERS statement, the database server can execute the trigger action when the trigger event is encountered, but the trigger does not perform retroactively. The database server does not attempt to execute the trigger for rows that were selected, inserted, deleted, or updated while the trigger was disabled and before it was enabled.

Warning: Because the behavior of a trigger varies according to its ENABLED or DISABLED mode, be cautious about disabling a trigger. If disabling a trigger will eventually destroy the semantic integrity of the database, do not disable the trigger.

Trigger Inheritance in a Table Hierarchy (IDS)

By default, any trigger that you define on a typed table of Dynamic Server is inherited by all its subtables.

In versions of Dynamic Server earlier than version 11.10, however, if you define a trigger on a subtable of a typed table, that trigger overrides any trigger for the same type of triggering event (Select, Delete, Insert, or Update) that the subtable inherits from its supertable. In this version of Dynamic Server, however, a table can inherit more than one trigger that the same triggering event activates, so both triggers are defined for the same type of event on the subtable.

In all versions of Dynamic Server, a trigger that you set on a subtable is inherited by all its dependent tables, but has no effect on its supertable.

This behavior is important when you require a trigger to be enabled in a supertable, but to be disabled in its subtable. In Dynamic Server 10.00 and in earlier versions, you cannot use the SET TRIGGERS option of the SET Database Object Mode statement to disable an inherited trigger selectively within a hierarchy. In this release, however, disabling a trigger on a table within a table hierarchy has no effect on inherited triggers. For example, the following statement has no effect on triggers on table objects that are above or below *subtable* within a table hierarchy:

```
SET TRIGGERS FOR subtable DISABLED
```

Similarly, the DROP TRIGGER statement cannot destroy an inherited trigger without also destroying the trigger on the supertable. In this situation, you must

instead define a trigger with no Action clause on the subtable. Because triggers are not additive, this empty trigger overrides the inherited trigger and executes for the subtable and for any subtables under the subtable, which are not subject to further overrides.

Triggers and SPL Routines

You cannot define a trigger in an SPL routine that is called inside a DML (data manipulation language) statement, as listed in “Data Manipulation Language Statements” on page 1-8. Thus, the following statement returns an error if the `sp_items` procedure includes the CREATE TRIGGER statement:

```
INSERT INTO items EXECUTE PROCEDURE sp_items;
```

You can use the CREATE FUNCTION or CREATE PROCEDURE statement of SQL with the REFERENCING clause to define *trigger routines* that include the FOR *table* or FOR *view* specification. These UDRs must include the REFERENCING clause that declares *correlation* names for OLD or NEW column values in the specified *table* or *view*. Triggers on the *table* or *view* can invoke the trigger routine from the FOR EACH ROW section of the Triggered Action list. Triggers can also invoke non-trigger routines from the BEFORE and AFTER sections of the Triggered Action list, but these UDRs cannot use *correlation* names to reference the NEW or OLD column values. The REFERENCING clause in a trigger routine supports the same syntax as in the CREATE TRIGGER statement, as described in the section “REFERENCING Clauses” on page 2-254.

Multiple triggers that the same triggering event executes can invoke more than one trigger routine, and these trigger routines can access the same NEW or OLD column values by using SPL variables that have the same names or different names. When a single triggering event executes multiple triggers, the order of execution is not guaranteed, but all of the BEFORE triggered actions execute before any of the FOR EACH ROW triggered actions, and all of the AFTER triggered actions execute after all of the FOR EACH ROW triggered actions.

For UDRs that are not trigger routines, SPL variables are not valid in CREATE TRIGGER statements. An SPL routine cannot perform INSERT, DELETE, or UPDATE operations on any table or view that is not local to the current database. See also “Rules for SPL Routines” on page 2-265 for additional restrictions on SPL routines that are invoked in triggered actions.

Trigger Events

The *trigger event* specifies what DML statements can initiate the trigger. The event can be an INSERT, DELETE, or UPDATE operation on the *table* or *view*, or (for Dynamic Server only) a SELECT operation that queries the *table*. Each CREATE TRIGGER statement must specify exactly one trigger event. Any SQL statement that is an instance of the trigger event is called a *triggering statement*.

For each *table*, you can define triggers that are activated by INSERT, DELETE, UPDATE, or SELECT statements. For each *view*, you can define INSTEAD OF triggers that are activated by INSERT, DELETE, or UPDATE statements. Multiple triggers on the same table or view can be activated by different types of trigger events or by the same type of trigger event.

You cannot specify a DELETE event if the triggering table has a referential constraint that specifies ON DELETE CASCADE.

You are responsible for guaranteeing that the triggering statement returns the same result with and without the trigger action on a table. See also the sections “ Action Clause” on page 2-251 and “Triggered Action” on page 2-257.

A triggering statement from an external database server can activate the trigger.

As the following example shows, an Insert trigger on **newtab**, managed by **dbserver1**, is activated by an INSERT statement from **dbserver2**. The trigger executes as if the INSERT originated on **dbserver1**.

```
-- Trigger on stores_demo@dbserver1:newtab
CREATE TRIGGER ins_tr INSERT ON newtab
  REFERENCING new_AS post_ins
  FOR EACH ROW(EXECUTE PROCEDURE nt_pct (post_ins.mc));
-- Triggering statement from dbserver2
INSERT INTO stores_demo@dbserver1:newtab
  SELECT item_num, order_num, quantity, stock_num, manu_code,
  total_price FROM items;
```

Dynamic Server also supports INSTEAD OF triggers on views, which are initiated when a triggering DML operation references the specified view. The INSTEAD OF trigger replaces the trigger event with the specified trigger action on a view, rather than execute the triggering INSERT, DELETE, or UPDATE operation. A *view* can have any number of INSTEAD OF trigger defined for each type of INSERT, DELETE, or UPDATE triggering event.

Trigger Events with Cursors

For triggers on tables, if the triggering statement uses a cursor, each part of the trigger action (including BEFORE, FOR EACH ROW, and AFTER, if these are specified for the trigger) is activated for each row that the cursor processes.

This behavior differs from what occurs when a triggering statement does not use a cursor and updates multiple rows. In this case, any BEFORE and AFTER triggered actions execute only once, but the FOR EACH ROW action list is executed for each row processed by the triggering statement. For additional information about trigger actions, see “ Action Clause” on page 2-251

Privileges on the Trigger Event

You must have appropriate Insert, Delete, Update, or Select privilege on the triggering table or view to execute a triggering INSERT, DELETE, UPDATE, or SELECT statement as the trigger event. The triggering statement might still fail, however, if you do not also have the privileges necessary to execute one of the SQL statements in the trigger action. When the trigger actions are executed, the database server checks your privileges for each SQL statement in the trigger definition, as if the statement were being executed independently of the trigger. For information on the privileges needed to execute the trigger actions, see “Privileges to Execute Trigger Actions” on page 2-266.

Performance Impact of Triggers

The INSERT, DELETE, UPDATE, and SELECT statements that initiate triggers might appear to execute slowly because they execute additional SQL statements, and the user might not know that other actions are occurring.

The execution time for a trigger event depends on the complexity of the trigger action and whether it initiates other triggers. The time increases as the number of cascading triggers increases. For more information on triggers that initiate other triggers, see “Cascading Triggers” on page 2-267.

INSERT Events and DELETE Events

INSERT and DELETE events on tables are defined by those keywords and by the ON *table* clause, using the following syntax.

INSERT or DELETE Event on a Table:



Element	Description	Restrictions	Syntax
<i>table</i>	Name of the triggering table	Must exist in the database	"Identifier " on page 5-21

An Insert trigger is activated when an INSERT statement includes the specified *table* (or a synonym for *table*) in its INTO clause. Similarly, a Delete trigger is activated when a DELETE statement includes the specified *table* (or a synonym for *table*) in its FROM clause.

For triggers on views, the INSTEAD OF keywords must immediately precede the INSERT, DELETE, or UPDATE keyword that specifies the type of trigger event, and the name or synonym of a *view* (rather than of a table) must follow the ON keyword. The section "INSTEAD OF Triggers on Views (IDS)" on page 2-270 describes the syntax for defining INSTEAD OF trigger events.

Any number of Insert triggers, and any number of Delete trigger, can be defined on the same table.

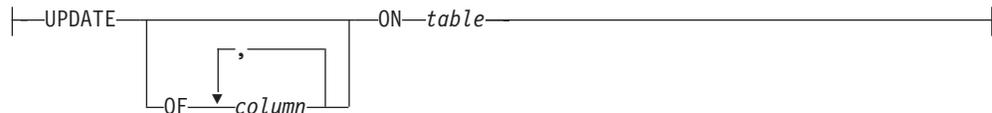
If you define a trigger on a subtable within a table hierarchy, and the subtable supports cascading deletes, then a DELETE operation on the supertable activates the Delete triggers on the subtable.

See also the section "Re-Entrancy of Triggers" on page 2-262 for information about dependencies and restrictions on the actions of Insert triggers and Delete triggers.

UPDATE Event

UPDATE events (and SELECT events) can include an optional *column* list.

UPDATE Event:



Element	Description	Restrictions	Syntax
<i>column</i>	Column that activates the trigger	Must exist in the triggering table	"Identifier " on page 5-21
<i>table</i>	Name of the triggering table	Must exist in the database	"Identifier " on page 5-21

The *column* list is optional. If you omit the OF *column* list, updating any column of *table* activates the trigger.

The `OF column` clause is not valid for an `INSTEAD OF` trigger on a view.

An `UPDATE` on the triggering table can activate the trigger in two cases:

- The `UPDATE` statement references any column in the *column* list.
- The `UPDATE` event definition has no `OF column` list specification.

Whether it updates one column or more than one column from the *column* list, a triggering `UPDATE` statement activates each Update trigger only once.

Defining Multiple Update Triggers

Multiple Update triggers on a table can include the same or different columns. In the following example, **trig3** on the **items** table includes in its column list **stock_num**, which is a triggering column in **trig1**.

```
CREATE TRIGGER trig1 UPDATE OF item_num, stock_num ON items
  REFERENCING OLD AS pre NEW AS post
  FOR EACH ROW(EXECUTE PROCEDURE proc1());
CREATE TRIGGER trig2 UPDATE OF manu_code ON items
  BEFORE(EXECUTE PROCEDURE proc2());
CREATE TRIGGER trig3 UPDATE OF order_num, stock_num ON items
  BEFORE(EXECUTE PROCEDURE proc3());
```

When an `UPDATE` statement updates multiple columns that have different triggers, the firing order is based on the lowest-numbered column that is defined in the triggers that actually fired, regardless of whether that lowest-numbered column was actually the triggering column when the trigger fired. If several Update triggers are set on the same column or on the same set of columns, however, the order of trigger execution is not guaranteed.

The following example shows that table **taba** has four columns (**a**, **b**, **c**, **d**):

```
CREATE TABLE taba (a int, b int, c int, d int);
```

Define **trig1** as an update on columns **a** and **c**, and define **trig2** as an update on columns **b** and **d**, as the following example shows:

```
CREATE TRIGGER trig1 UPDATE OF a, c ON taba
  AFTER (UPDATE tabb SET y = y + 1);

CREATE TRIGGER trig2 UPDATE OF b, d ON taba
  AFTER (UPDATE tabb SET z = z + 1);
```

The following example shows a triggering statement for the Update trigger:

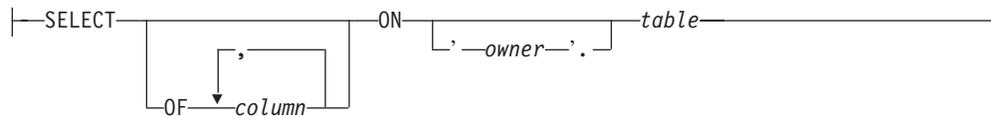
```
UPDATE taba SET (b, c) = (b + 1, c + 1);
```

Then **trig1** for columns **a** and **c** executes first, and **trig2** for columns **b** and **d** executes next. In this case, the smallest column number in the two triggers is column 1 (**a**), and the next is column 2 (**b**).

SELECT Event (IDS)

`DELETE` and `INSERT` events are defined by those keywords (and the `ON table` clause), but `SELECT` and `UPDATE` events also support an optional *column* list.

SELECT Event:



Element	Description	Restrictions	Syntax
<i>column</i>	Column that activates the trigger	Must exist in the triggering <i>table</i>	"Identifier " on page 5-21
<i>owner</i>	Owner of <i>table</i>	Must own <i>table</i>	" Owner Name " on page 5-45
<i>table</i>	Name of the triggering table	Must exist in the database	"Identifier " on page 5-21

If you define more than one Select trigger on the same table, the *column* list is optional, and the *column* lists for each trigger can be unique or can duplicate that of another Select trigger.

A SELECT on the triggering table can activate the trigger in two cases:

- The SELECT statement references any column in the *column* list.
- The SELECT event definition has no OF *column* list specification.

(Sections that follow, however, describe additional circumstances that can affect whether or not a SELECT statement activates a Select trigger.)

Whether it specifies one column or more than one column from the *column* list, a triggering SELECT statement activates the Select trigger only once.

The action of a Select trigger cannot include an UPDATE, INSERT, or DELETE on the triggering table. The action of a Select trigger can include UPDATE, INSERT, and DELETE actions on tables other than the triggering table. The following example defines a Select trigger on one column of a table:

```
CREATE TRIGGER mytrig
  SELECT OF cola ON mytab REFERENCING OLD AS pre
  FOR EACH ROW (INSERT INTO newtab VALUES('for each action'));
```

You cannot specify a SELECT event for an INSTEAD OF trigger on a view.

Circumstances When a Select Trigger Is Activated

A query on the triggering table activates a Select trigger in these cases:

- The SELECT statement is a stand-alone SELECT statement.
- The SELECT statement occurs within a UDR called in a select list.
- The SELECT statement is a subquery in the Projection list.
- The SELECT statement is a subquery in the FROM clause.
- The SELECT statement occurs within a UDR called by EXECUTE PROCEDURE or EXECUTE FUNCTION.
- The SELECT statement selects data from a supertable in a table hierarchy. In this case the SELECT statement activates Select triggers for the supertable and all the subtables in the hierarchy.

For information on SELECT statements that do not activate a Select trigger, see “Circumstances When a Select Trigger Is Not Activated” on page 2-251.

Stand-alone SELECT Statements

A Select trigger is activated if the triggering column appears in the select list of the Projection clause of a stand-alone SELECT statement.

For example, if a Select trigger is defined to execute whenever column **col1** of table **tab1** is selected, then both of the following stand-alone SELECT statements activate the Select trigger:

```
SELECT * FROM tab1;  
SELECT col1 FROM tab1;
```

SELECT Statements Within UDRs in the Select List

A Select trigger is activated by a UDR if the UDR contains a SELECT statement within its statement block, and the UDR also appears in the select list of the Projection clause of a SELECT statement. For example, assume that a UDR named **my_rtn** contains this SELECT statement in its statement block:

```
SELECT col1 FROM tab1;
```

Now suppose that the following SELECT statement invokes the **my_rtn** UDR in its select list:

```
SELECT my_rtn() FROM tab2;
```

This SELECT statement activates the Select trigger defined on column **col1** of table **tab1** when the **my_rtn** UDR is executed.

UDRs That EXECUTE PROCEDURE and EXECUTE FUNCTION Call

A Select trigger is activated by a UDR if the UDR contains a SELECT statement within its statement block and the UDR is called by an EXECUTE PROCEDURE or (for IDS triggers only) the EXECUTE FUNCTION statement. For example, assume that the user-defined procedure named **my_rtn** contains the following SELECT statement in its statement block:

```
SELECT col1 FROM tab1;
```

Now suppose that the following statement invokes the **my_rtn** procedure:

```
EXECUTE PROCEDURE my_rtn();
```

This statement activates the Select trigger defined on column **col1** of table **tab1** when the SELECT statement within the statement block is executed.

Subqueries in the Select List

A Select trigger can be activated by a subquery that appears in the select list of the Projection clause of a SELECT statement.

For example, if an enabled Select trigger was defined on **col1** of **tab1**, the subquery in the following SELECT statement activates that trigger:

```
SELECT (SELECT col1 FROM tab1 WHERE col1=1), colx, col y FROM tabz;
```

Subqueries in the FROM Clause of SELECT

Table expressions in the FROM clause of a SELECT statement can be the triggering event on a table that is referenced by an uncorrelated subquery. In the following example, the subquery that specifies a table expression is the triggering event for any enabled Select triggers that are defined on **col1** of **tab1**:

```
SELECT vcol FROM (SELECT FIRST 5 col1 FROM tab1 ORDER BY col1 ) vtab(vcol);
```

Subqueries in the WHERE Clause of DELETE or UPDATE

Subqueries that are specified using the Condition with Subquery syntax in the WHERE clause of the DELETE statement or the UPDATE statement cannot be the triggering event for a Select trigger. In the following example, the subquery is not the triggering event for any enabled Select triggers that are defined on **col2** of **tab1**:

```
DELETE tab1 WHERE EXISTS  
  (SELECT col2 FROM tab1 WHERE col2 > 1024);
```

The DELETE operation in same example, however, activates any enabled Delete triggers that are defined on **tab1**. No enabled Select trigger on **tab1** can be activated by a subquery within a DELETE statement that modifies a table referenced in the FROM clause of the subquery.

Similarly, the subquery in the WHERE clause of the following UPDATE statement is not the triggering event for any enabled Select triggers that are defined on **col3** of **tab1**:

```
UPDATE tab1 SET col3 = col3 + 10  
  WHERE col3 > ANY  
    (SELECT col3 from tab1 WHERE col3 > 1);
```

The same example activates any enabled Update trigger that is defined on **col3** of **tab1**, but no Select trigger can be updated by the subquery. For additional restrictions on Select triggers, see “Circumstances When a Select Trigger Is Not Activated” on page 2-251.

Select Triggers in Table Hierarchies

A subtable in a Dynamic Server database inherits the Select triggers that are defined on its supertable. When you select from a supertable, the SELECT statement activates the Select triggers on the supertable and the inherited Select triggers on the subtables in the table hierarchy.

For example, assume that table **tab1** is the supertable and table **tab2** is the subtable in a table hierarchy. If the Select trigger **trig1** is defined on table **tab1**, a SELECT statement on table **tab1** activates the Select trigger **trig1** for the rows in table **tab1** and the inherited Select trigger **trig1** for the rows in table **tab2**.

If you add a Select trigger to a subtable, this Select trigger does not override the Select trigger that the subtable inherits from its supertable, but increases the number of Select triggers on the subtable. For example, if the Select trigger **trig1** is defined on column **col1** in supertable **tab1**, the subtable **tab2** inherits this trigger. If you define a Select trigger named **trig2** on column **col1** in subtable **tab2**, and a SELECT statement selects from **col1** in supertable **tab1**, this SELECT statement activates trigger **trig1** for the rows in table **tab1** and both triggers **trig1** and **trig2** for the rows in table **tab2**. (In version 10.0 and in earlier releases of Dynamic Server, however, a Select trigger that you add to the subtable overrides the Select trigger that the subtable inherited from the supertable.)

Circumstances When a Select Trigger Is Not Activated

In Dynamic Server, a SELECT statement on the triggering table does not activate a Select trigger in certain circumstances:

- If a subquery or UDR that contains the triggering SELECT statement appears in any clause of a SELECT statement other than the Projection clause or the FROM clause, the Select trigger is not activated.

For example, if the subquery or UDR appears in the WHERE clause or HAVING clause of a SELECT statement, the SELECT statement within the subquery or UDR does not activate the Select trigger.

- If the trigger action of a Select trigger calls a UDR that includes a triggering SELECT statement, the Select trigger on the SELECT in the UDR is not activated. Cascading Select triggers are not supported.
- If a SELECT statement contains a built-in aggregate or user-defined aggregate in its Projection clause, the Select trigger is not activated. For example, the following SELECT statement does not activate a Select trigger defined on **col1** of **tab1**:

```
SELECT MIN(col1) FROM tab1;
```

- A SELECT statement that includes the UNION or UNION ALL operator does not activate a Select trigger.
- The SELECT clause of INSERT does not activate a Select trigger.
- A subquery in the WHERE clause of the DELETE or UPDATE statement cannot activate a Select trigger on the same table that the DELETE or UPDATE statement is modifying.
- If the Projection clause of a SELECT includes the DISTINCT or UNIQUE keywords, the SELECT statement does not activate a Select trigger.
- Select triggers are not supported on scroll cursors.
- If a SELECT statement refers to a remote triggering table, the Select trigger is not activated on the remote database server.
- Columns in the ORDER BY list of a query activate no Select triggers (nor any other triggers) unless they are also listed in the Projection clause.

An exception to the last restriction is that a Select trigger can be activated by a column in the ORDER BY list of a subquery in the FROM clause, whether or not the same column also appears in the Projection clause. In the following example, a table expression that includes **col1** in the ORDER BY clause (but not in the select list of the Projection clause) is the triggering event for any enabled Select triggers that are defined on **col1** of **tab1**:

```
SELECT vcol FROM (SELECT col2 FROM tab1 ORDER BY col1 ) vtab(vcol);
```

Action Clause

The Action clause defines the SQL statements that are executed when the trigger is activated. For a trigger on a table, there can be up to three sections in the Action clause: BEFORE, AFTER and FOR EACH ROW.

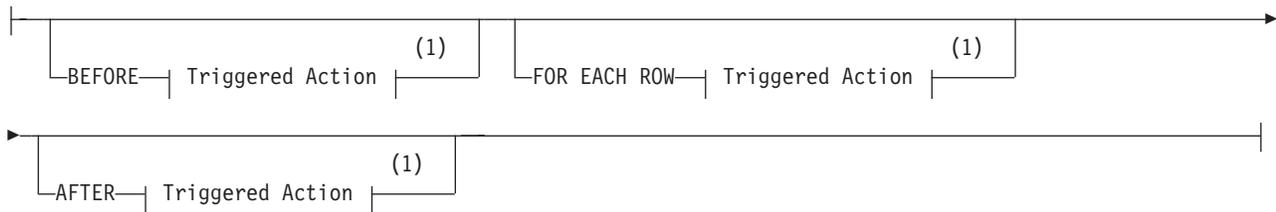
- The BEFORE actions are executed once for each triggering event, before the database server performs the triggering DML operation.
- The AFTER actions are also executed once for each triggering DML event, after the operation on the table is complete, in the context of the triggering statement.
- The FOR EACH ROW actions are executed for each row that is inserted, updated, deleted or selected in the DML operation, after the DML operation is executed on each row, but before the database server writes the values into the log and into the table.

If the same table has multiple triggers that are activated by the same triggering event, the order of trigger execution is not guaranteed, but all of the BEFORE triggered actions execute before any of the FOR EACH ROW triggered actions, and all of the AFTER triggered actions execute after all of the FOR EACH ROW triggered actions.

When you define an INSTEAD OF trigger on a view, the BEFORE and AFTER keywords are not supported, but the FOR EACH ROW section of the Action clause is valid. See the section “INSTEAD OF Triggers on Views (IDS)” on page 2-270 for the syntax of specifying triggered actions on a view.

The Action clause has the following syntax.

Action Clause:



Notes:

- 1 See “Triggered Action” on page 2-257

For the trigger to have any effect on the table, you must define at least one triggered action, using the keywords BEFORE, FOR EACH ROW, or AFTER to indicate when the action occurs relative to execution of the triggering event.

You can specify actions for any or all of these three options on a single trigger, but any BEFORE action list must be specified first, and any AFTER action list must be specified last. For more information on the Action clause when a REFERENCING clause is also specified, see “ Correlated Table Action” on page 2-257.

BEFORE Actions

The list of BEFORE trigger actions execute once before the triggering statement executes. Even if the triggering statement does not process any rows, the database server executes the BEFORE trigger actions.

FOR EACH ROW Actions

After a row of the triggering table is processed, the database server executes all of the statements of the FOR EACH ROW trigger action list; this cycle is repeated for every row that the triggering statement processes. (But if the triggering statement does not insert, delete, update, or select any rows, the database server does not execute the FOR EACH ROW trigger actions.)

In Extended Parallel Server, you cannot define FOR EACH ROW actions on tables that have globally-detached indexes.

In Dynamic Server, the FOR EACH ROW action list of a Select trigger is executed once for each instance of a row. For example, the same row can appear more than once in the result of a query joining two tables. For more information on FOR EACH ROW actions that reference specific values in rows that the triggering statement processes, see “REFERENCING Clauses” on page 2-254.

AFTER Actions

The specified set of AFTER trigger actions executes once after the action of the triggering statement is complete. If the triggering statement does not process any rows, the AFTER trigger actions still execute.

Actions of Multiple Triggers

When an UPDATE statement activates multiple triggers, the trigger actions merge. Assume that **taba** has columns **a**, **b**, **c**, and **d**, as this example shows:

```
CREATE TABLE taba (a INT, b INT, c INT, d INT);
```

Next, assume that you define **trig1** on columns **a** and **c**, and **trig2** on columns **b** and **d**. If both triggers specify BEFORE, FOR EACH ROW, and AFTER actions, then the trigger actions are executed in the following order:

1. BEFORE action list for trigger (**a**, **c**)
2. BEFORE action list for trigger (**b**, **d**)
3. FOR EACH ROW action list for trigger (**a**, **c**)
4. FOR EACH ROW action list for trigger (**b**, **d**)
5. AFTER action list for trigger (**a**, **c**)
6. AFTER action list for trigger (**b**, **d**)

The database server treats all the triggers that are activated by the same triggering statement as a single trigger, and the trigger action is the merged-action list. All the rules that govern a trigger action apply to the merged list as one list, and no distinction is made between the two original triggers.

Guaranteeing Row-Order Independence

In a FOR EACH ROW triggered-action list, the result might depend on the order of the rows being processed. You can ensure that the result is independent of row order by following these suggestions:

- Avoid selecting the triggering table in the FOR EACH ROW section.
If the triggering statement affects multiple rows in the triggering table, the result of the SELECT statement in the FOR EACH ROW section varies as each row is processed. This condition also applies to any cascading triggers. See “Cascading Triggers” on page 2-267.
- In the FOR EACH ROW section, avoid updating a table with values derived from the current row of the triggering table.
If the trigger actions modify any row in the table more than once, the final result for that row depends on the order in which rows from the triggering table are processed.
- Avoid modifying a table in the FOR EACH ROW section that is selected by another statement in the same FOR EACH ROW trigger action, including any cascading trigger actions.

If FOR EACH ROW actions modify a table, the changes might not be complete when a subsequent action of the trigger refers to the table. In this case, the result might differ, depending on the order in which rows are processed.

The database server does not enforce rules to prevent these situations because doing so would restrict the set of tables from which a trigger action can select. Furthermore, the result of most trigger actions is independent of row order. Consequently, you are responsible for ensuring that the results of the trigger actions are independent of row order.

REFERENCING Clauses

The REFERENCING clause for any event declares a *correlation* name (or for Update triggers, two *correlation* names) that can be used to qualify column values in the triggering table. These names enable FOR EACH ROW actions to reference new or old column values in the result of trigger events.

They also enable FOR EACH ROW actions to reference old column values that existed in the triggering table prior to modification by trigger events.

Correlation names are not valid if the triggered action includes both the INSERT statement and the BEFORE WHEN or AFTER WHEN keywords. This restriction does not affect triggered actions that specify the FOR EACH ROW keywords without the BEFORE or AFTER keywords, or that include no INSERT statement.

The REFERENCING clause syntax that is described here for the CREATE TRIGGER statement is also valid in CREATE FUNCTION and CREATE PROCEDURE statements that define a trigger routine, provided that the CREATE FUNCTION or CREATE PROCEDURE statement also includes the FOR *table_object* clause to specify the table or view whose FOR EACH ROW actions can invoke the trigger routine.

REFERENCING Clause for Delete

The REFERENCING clause for a Delete trigger can declare a correlation name for the deleted value in a column.

REFERENCING Clause for Delete:

```

REFERENCING OLD AS correlation

```

Element	Description	Restrictions	Syntax
<i>correlation</i>	Name that you declare here to qualify an old column value (as <i>correlation.column</i>) in a triggered action	Must be unique within this CREATE TRIGGER statement	"Identifier " on page 5-21

The *correlation* is a qualifier for the column value in the triggering table before the triggering statement executed. The *correlation* is in scope in the FOR EACH ROW trigger action list. See "Correlated Table Action" on page 2-257.

To use a correlation name in a trigger action to refer to an old column value, prefix the column name with the correlation name and a period (.) symbol. For example, if the NEW *correlation* name is **post**, refer to the new value for the column **fname** as **post.fname**.

If the trigger event is a DELETE statement, using the new correlation name as a qualifier causes an error, because the column has no value after the row is deleted. For the rules that govern the use of correlation names, see "Using Correlation Names in Triggered Actions" on page 2-260.

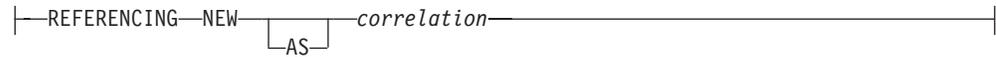
You can use the REFERENCING clause for Delete only if you define a FOR EACH ROW trigger action.

In Extended Parallel Server, the OLD *correlation* value cannot be a BYTE or TEXT value. That is, it cannot refer to a BYTE or TEXT column.

REFERENCING Clause for Insert

The REFERENCING clause for an Insert trigger can declare a correlation name for the inserted value in a column.

REFERENCING Clause for Insert:



Element	Description	Restrictions	Syntax
<i>correlation</i>	Name that you declare here to qualify a new column value (as <i>correlation.column</i>) in a triggered action	Must be unique within this CREATE TRIGGER statement	"Identifier " on page 5-21

The *correlation* is a name for the new column value after the triggering statement has executed. Its scope of reference is only the FOR EACH ROW trigger action list; see "Correlated Table Action" on page 2-257. To use the correlation name, precede the column name with the *correlation* name, followed by a period (.) symbol. Thus, if the NEW *correlation* name is **post**, refer to the new value for the column **fname** as **post.fname**.

If the trigger event is an INSERT statement, using the old *correlation* name as a qualifier causes an error, because no value exists before the row is inserted. For the rules that govern how to use correlation names, see "Using Correlation Names in Triggered Actions" on page 2-260. You can use the INSERT REFERENCING clause only if you define a FOR EACH ROW trigger action.

The following example illustrates use of the INSERT REFERENCING clause. This example inserts a row into **backup_table1** for every row that is inserted into **table1**. The values that are inserted into **col1** and **col2** of **backup_table1** are an exact copy of the values that were just inserted into **table1**.

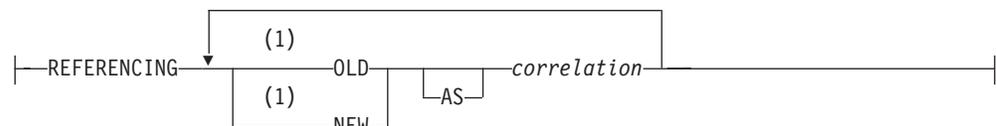
```
CREATE TABLE table1 (col1 INT, col2 INT);
CREATE TABLE backup_table1 (col1 INT, col2 INT);
CREATE TRIGGER before_trig
  INSERT ON table1 REFERENCING NEW AS new
  FOR EACH ROW
  (
    INSERT INTO backup_table1 (col1, col2)
    VALUES (new.col1, new.col2)
  );
```

As the preceding example shows, the INSERT REFERENCING clause allows you to refer to data values produced by the trigger action.

REFERENCING Clause for Update

The REFERENCING clause for an Update trigger can declare correlation names for the original value and for the updated value in a column.

REFERENCING Clause for Update:



Notes:

- 1 Use path no more than once

Element	Description	Restrictions	Syntax
<i>correlation</i>	Name that you declare here to qualify an old or new column value (as <i>correlation.column</i>) in a triggered action	Must be unique within this CREATE TRIGGER statement	"Identifier " on page 5-21

The OLD *correlation* is the name of the value of the column in the triggering table before execution of the triggering statement; the NEW *correlation* identifies the corresponding value after the triggering statement executes.

The scope of reference of the *correlation* names that you declare here is only within the FOR EACH ROW trigger action list. See "Correlated Table Action" on page 2-257.

To refer to an old or new column value, prefix the column name with the *correlation* name and a period (.) symbol. For example, if the new *correlation* name is **post**, you can refer to the new value in column **fname** as **post.fname**.

If the trigger event is an UPDATE statement, you can define both old and new *correlation* names to refer to column values before and after the triggering UPDATE statement. For rules that govern the use of *correlation* names, see "Using Correlation Names in Triggered Actions" on page 2-260.

You can use the UPDATE REFERENCING clause only if you define a FOR EACH ROW trigger action.

In Extended Parallel Server, the OLD *correlation* value cannot be a BYTE or TEXT value. That is, it cannot refer to a BYTE or TEXT column.

REFERENCING Clause for Select (IDS)

The REFERENCING clause for a Select trigger can declare a correlation name for the value in a column.

REFERENCING Clause for Select:

```
|—REFERENCING—OLD—AS—correlation—|
```

Element	Description	Restrictions	Syntax
<i>correlation</i>	Name that you declare here to qualify an old or new column value (as <i>correlation.column</i>) in a triggered action	Must be unique within this CREATE TRIGGER statement	"Identifier " on page 5-21

This has the same syntax as the "REFERENCING Clause for Delete" on page 2-254. The scope of reference of the *correlation* name that you declare here is only within the FOR EACH ROW trigger action list. See "Correlated Table Action" on page 2-257.

You use the *correlation* name to refer to an old column value by preceding the column name with the *correlation* name and a period (.) symbol. For example, if the old *correlation* name is **pre**, you can refer to the old value for the column **fname** as **pre.fname**.

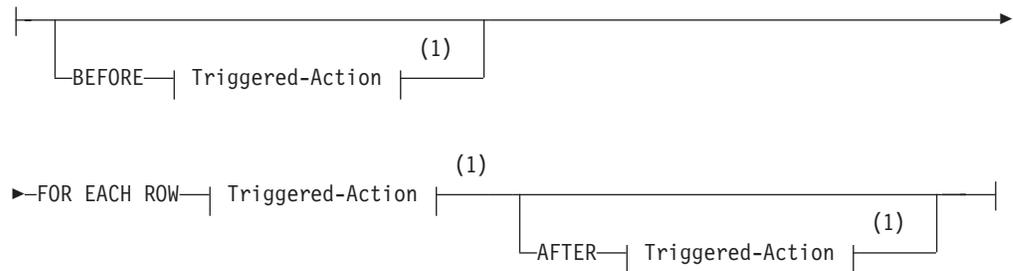
If the trigger event is a SELECT statement, using the new *correlation* name as a qualifier causes an error because the column does not have a new value after the column is selected. For the rules that govern the use of correlation names, see “Using Correlation Names in Triggered Actions” on page 2-260.

You can use the SELECT REFERENCING clause only if you define a FOR EACH ROW trigger action.

Correlated Table Action

Use the Correlated Trigger Action clause to define the SQL statements that are executed as the trigger action when a triggering event activates a trigger on a table.

Correlated Table Action:



Notes:

- 1 See “Triggered Action”

If the CREATE TRIGGER statement contains an INSERT REFERENCING clause, a DELETE REFERENCING clause, an UPDATE REFERENCING clause, or (for Dynamic Server only) a SELECT REFERENCING clause, you must include a FOR EACH ROW triggered-action list in the Action clause. You can also include BEFORE and AFTER triggered-action lists, but they are optional.

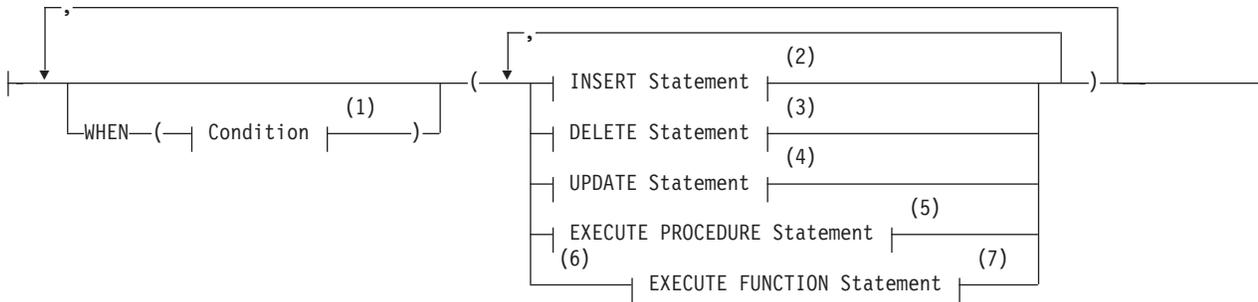
For information on the BEFORE, FOR EACH ROW, and AFTER triggered-action lists, see “ Action Clause” on page 2-251

In Extended Parallel Server, you cannot specify FOR EACH ROW actions on tables that have globally-detached indexes.

Triggered Action

The Triggered Action specifies a list of SQL statements to execute when a trigger is activated. The BEFORE, FOR EACH ROW, and AFTER sections of the Action Clause can each specify different list of triggered actions for the same trigger.

Triggered Action:



Notes:

- 1 See "Condition " on page 4-5
- 2 See "INSERT " on page 2-430
- 3 See "DELETE " on page 2-303
- 4 See "UPDATE" on page 2-689
- 5 See "EXECUTE PROCEDURE " on page 2-364
- 6 Dynamic Server only
- 7 See "EXECUTE FUNCTION " on page 2-356

For a trigger on a table, the trigger action consists of an optional WHEN condition and the action statements. You can specify a triggered-action list for each WHEN clause, or you can specify a single list (of one or more trigger actions) if you include no WHEN clause.

Database objects that are referenced explicitly in the trigger action or in the definition of the trigger event, such as tables, columns, and UDRs, must exist when the CREATE TRIGGER statement defines the new trigger.

Attention: When you specify a date expression in the WHEN condition or in an action statement, make sure to specify four digits instead of two digits for the year. For more about abbreviated years, see the description of **DBCENTURY** in the *IBM Informix Guide to SQL: Reference*, which also describes how the behavior of some database objects can be affected by environment variable settings. Like fragmentation expressions, check constraints, and UDRs, triggers are stored in the system catalog with the creation-time settings of environment variables that can affect the evaluation of expressions like the WHEN condition. The database server ignores any subsequent changes to those settings when evaluating expressions in those database objects.

WHEN Condition

The WHEN condition makes the triggered action dependent on the outcome of a test. When you include a WHEN condition in a triggered action, the statements in the triggered action list execute only if the condition evaluates to true. If the WHEN condition evaluates to false or unknown, then the statements in the triggered action list are not executed.

If the triggered action is in a FOR EACH ROW section, its condition is evaluated for each row. For example, the triggered action in the following trigger executes only if the condition in the WHEN clause is true:

```

CREATE TRIGGER up_price
  UPDATE OF unit_price ON stock
  REFERENCING OLD AS pre NEW AS post
  WHEN (pre.unit_price < NEW.unit_price)
  UPDATE unit_price SET unit_price = NEW.unit_price;
  
```

```
FOR EACH ROW WHEN(post.unit_price > pre.unit_price * 2)
(ININSERT INTO warn_tab VALUES(pre.stock_num, pre.order_num,
pre.unit_price, post.unit_price, CURRENT));
```

An SPL routine that executes inside the WHEN condition carries the same restrictions as a UDR that is called in a data manipulation statement. That is, the SPL routine cannot contain certain SQL statements. For information on which statements are restricted, see “Restrictions on SPL Routines in Data-Manipulation Statements” on page 5-78.

Action Statements

The triggered-action statements can be INSERT, DELETE, UPDATE, EXECUTE FUNCTION, or EXECUTE PROCEDURE statements. If the action list contains multiple statements, and the WHEN condition is satisfied (or is absent), then these statements execute in the order in which they appear in the list.

UDRs as Triggered Actions: Calls to user-defined functions and procedures, including trigger routines, can be triggered actions. The triggered action list of the FOR EACH ROW clause can include calls to UDRs that call **mi_trigger*()** functions. Triggered actions are the only context in which a trigger routine of Dynamic Server can be invoked. For restrictions on the calling context and the syntax of trigger routines, see “The REFERENCING and FOR Clauses” on page 2-116.

In Dynamic Server, you can use the EXECUTE FUNCTION statement to call any user-defined function or trigger function. Use the EXECUTE PROCEDURE statement to call any user-defined procedure or trigger procedure.

In contexts where Boolean expressions are valid, the Boolean operators SELECTING, INSERTING, DELETING, and UPDATING are valid in trigger routines and in other UDRs that are invoked in triggered action statements. These operators return TRUE ('t') if the triggering event matches the DML operation that matches the name of the operator; otherwise they return FALSE ('f'). A single trigger routine can be designed to perform different triggered actions for different types of triggering events, using these Boolean operators to execute program blocks that are appropriate to the type of trigger.

In Extended Parallel Server, you can use the EXECUTE PROCEDURE statement to execute any SPL routine.

For restrictions on using SPL routines as triggered actions, see “Rules for SPL Routines” on page 2-265 and “Triggers and SPL Routines” on page 2-244.

Achieving a Consistent Result: To guarantee that the triggering statement returns the same result with and without the triggered actions, make sure that the triggered actions in the BEFORE and FOR EACH ROW sections do not modify any table referenced in the following clauses:

- WHERE clause
- SET clause in the UPDATE statement
- SELECT clause
- EXECUTE PROCEDURE clause or EXECUTE FUNCTION clause in a multiple-row INSERT statement.

Using Reserved Words: If you use the INSERT, DELETE, UPDATE, or EXECUTE reserved words as a *correlation* identifier in any of the following clauses inside a triggered action list, you must qualify them by the *owner* name, the *table* name, or both:

- FROM clause of a SELECT statement
- INTO clause of the EXECUTE PROCEDURE or EXECUTE FUNCTION statement
- GROUP BY clause
- SET clause of the UPDATE statement.

You get a syntax error if these keywords are *not* qualified when you use these clauses inside a triggered action.

If you use the keyword as a column name, it must be qualified by the table name; for example, **table.update**. If both the table name and the column name are keywords, they must be qualified by the owner name (for example, *owner.insert.update*). If the owner name, table name, and column name are all keywords, the owner name must be in quotation marks; for example, **'delete'.insert.update**. (These are general rules regarding reserved words as identifiers, rather than special cases for triggers. Your code will be easier to read and to maintain if you avoid using the keywords of SQL as identifiers.)

The only exception is when these keywords are the first table or column name in the list, and you do not need to qualify them. For example, **delete** in the following statement does not need to be qualified because it is the first column listed in the INTO clause:

```
CREATE TRIGGER t1 UPDATE OF b ON tab1
  FOR EACH ROW (EXECUTE PROCEDURE p2() INTO delete, d);
```

The following statements show examples in which you must qualify the column name or the table name:

- FROM clause of a SELECT statement

```
CREATE TRIGGER t1 INSERT ON tab1
  BEFORE (INSERT INTO tab2 SELECT * FROM tab3, 'owner1'.update);
```

- INTO clause of an EXECUTE PROCEDURE statement

```
CREATE TRIGGER t3 UPDATE OF b ON tab1
  FOR EACH ROW (EXECUTE PROCEDURE p2() INTO
  d, tab1.delete);
```

In Dynamic Server, an INSTEAD OF trigger on a view cannot include the EXECUTE PROCEDURE INTO statement among its triggered actions.

- GROUP BY clause of a SELECT statement

```
CREATE TRIGGER t4 DELETE ON tab1
  BEFORE (INSERT INTO tab3 SELECT deptno, SUM(exp)
  FROM budget GROUP BY deptno, budget.update);
```

- SET clause of an UPDATE statement

```
CREATE TRIGGER t2 UPDATE OF a ON tab1
  BEFORE (UPDATE tab2 SET a = 10, tab2.insert = 5);
```

Using Correlation Names in Triggered Actions

These rules apply when you use correlation names in triggered actions:

- You can use the correlation names as qualifiers for the old and new column values in SQL statements of the FOR EACH ROW triggered-action list and in the WHEN condition.

- The WHEN conditions and FOR EACH ROW clauses of multiple triggers on the same table can use different correlated variables in the REFERENCING clauses of triggers and of trigger routines to reference values in the same column.
- The old and new correlation names refer to all rows affected by the triggering statement.
- You cannot use the correlation name to qualify a column name in the GROUP BY, the SET, or the COUNT DISTINCT clause.
- The scope of reference of the correlation names is the entire trigger definition. This scope is statically determined, meaning that it is limited to the trigger definition; it does not encompass cascading triggers or columns that are qualified by a table name in a UDR that is a triggered action, except for trigger routines that are invoked in the FOR EACH ROW clause.

For additional information on using correlation names in trigger routines, see “Rules for SPL Routines” on page 2-265.

When to Use Correlation Names

In SQL statements of the FOR EACH ROW list, you must qualify all references to columns in the triggering table with either the old or new correlation name, unless the statement is valid independent of the triggered action.

In other words, if a column name inside a FOR EACH ROW triggered action list is not qualified by a correlation name, even if it is qualified by the triggering table name, it is interpreted as if the statement were independent of the triggered action. No special effort is made to search the definition of the triggering table for the non-qualified column name.

For example, assume that the following DELETE statement is a triggered action inside the FOR EACH ROW section of a trigger:

```
DELETE FROM tab1 WHERE col_c = col_c2;
```

For the statement to be valid, both **col_c** and **col_c2** must be columns from **tab1**. If **col_c2** is intended to be a correlation reference to a column in the triggering table, it must be qualified by either the old or the new correlation name. If **col_c2** is not a column in **tab1** and is not qualified by either the old or new correlation name, you get an error.

In a statement that is valid independent of the triggered action, a column name with no *correlation* qualifier refers to the current value in the database.

In the triggered action for trigger **t1** in the next example, **mgr** in the WHERE clause of the correlated subquery is an unqualified column in the triggering table. In this case, **mgr** refers to the current column value in **empsal** because the INSERT statement is valid independent of the triggered action.

```
CREATE DATABASE db1;
CREATE TABLE empsal (empno INT, salary INT, mgr INT);
CREATE TABLE mgr (eno INT, bonus INT);
CREATE TABLE biggap (empno INT, salary INT, mgr INT);

CREATE TRIGGER t1 UPDATE OF salary ON empsal
AFTER (INSERT INTO biggap SELECT * FROM empsal WHERE salary <
(SELECT bonus FROM mgr WHERE eno = mgr));
```

In a triggered action, an unqualified column name from the triggering table refers to the current column value, but only when the triggered statement is valid independent of the triggered action.

Qualified Versus Unqualified Value

The following table summarizes what value is retrieved when the **column** name is qualified by the *old* or by the *new* correlation name after various trigger events.

Trigger Event	<i>old.column</i>	<i>new.column</i>
INSERT	No value (error)	Inserted value
UPDATE (column updated)	Original value	Current value (U)
UPDATE (column not updated)	Original value	Original value (N)
DELETE	Original value	No value (error)
SELECT	Original value	No value (error)

When a correlation name has no value, an error is issued only when an SQL or SPL statement referencing the undefined correlation is executed, rather than when the correlation name is declared. Refer to the following key when you read the previous table.

Term Meaning

Original value

Value before the triggering event

Current value

Value after the triggering event

(N) Cannot be changed by triggered action

(U) Can be updated by triggered actions; updated value might be different from the original value because of preceding triggered actions.

Outside a FOR EACH ROW triggered-action list, you cannot qualify a column from the triggering table with either the old correlation name or the new correlation name; it always refers to the current value in the database.

In Dynamic Server, statements in the trigger action list use whatever collating order was in effect when the trigger was created, even if a different collation is in effect when the trigger action is executed. See "SET COLLATION" on page 2-584 for details of how to specify a collating order different from what **DB_LOCALE** specifies.

Re-Entrancy of Triggers

In some cases a trigger can be re-entrant. In these cases the triggered action can reference the triggering table. In other words, both the trigger event and the triggered action can operate on the same table. The following list summarizes the situations in which triggers can be re-entrant and the situations in which triggers cannot be re-entrant:

- The trigger action of an Update trigger cannot be an INSERT or DELETE of the table that the trigger event updated.
- Similarly, the trigger action of an Update trigger cannot be an UPDATE of a column that the trigger event updated. (But the trigger action of an Update trigger can update a column that was not updated by the trigger event.)

For example, assume that the following UPDATE statement, which updates columns **a** and **b** of **tab1**, is the triggering statement:

```
UPDATE tab1 SET (a, b) = (a + 1, b + 1);
```

Now consider the trigger actions in the following example. The first UPDATE statement is a valid trigger action, but the second one is not, because it updates column **b** again.

```
UPDATE tab1 SET c = c + 1;    -- OK
UPDATE tab1 SET b = b + 1;    -- INVALID
```

- If the trigger has an UPDATE event, the trigger action can be an EXECUTE PROCEDURE or EXECUTE FUNCTION statement with an INTO clause that references a column that was updated by the trigger event or any other column in the triggering table.

When an EXECUTE PROCEDURE or EXECUTE FUNCTION statement is the trigger action, the INTO clause for an UPDATE trigger is valid only in FOR EACH ROW trigger actions, and column names that appear in the INTO clause must be from the triggering table.

The following statement illustrates the appropriate use of the INTO clause:

```
CREATE TRIGGER upd_totpr UPDATE OF quantity ON items
  REFERENCING OLD AS pre_upd NEW AS post_upd
  FOR EACH ROW (EXECUTE PROCEDURE
    calc_totpr(pre_upd.quantity, post_upd.quantity,
    pre_upd.total_price) INTO total_price);
```

The column that follows the INTO keyword must be in the triggering table, but need not have been updated by the trigger event.

When the INTO clause appears in the EXECUTE PROCEDURE or EXECUTE FUNCTION statement, the database server updates the specified columns with values returned from the UDR, immediately upon returning from the UDR.

- If the trigger has an INSERT event, the trigger action cannot be an INSERT or DELETE statement that references the triggering table.
- If the trigger has an INSERT event, the trigger action can be an UPDATE statement that references a column in the triggering table, but this column cannot be a column for which a value was supplied by the trigger event.

If the trigger has an INSERT event, and the trigger action updates the triggering table, the columns in both statements must be mutually exclusive. For example, assume that the triggering statement inserts values for columns **cola** and **colb** of table **tab1**:

```
INSERT INTO tab1 (cola, colb) VALUES (1,10);
```

Now consider the following trigger actions. The first UPDATE is valid, but the second one is not, because it updates column **colb** even though the trigger event already supplied a value for column **colb**:

```
UPDATE tab1 SET colc=100; --OK
UPDATE tab1 SET colb=100; --INVALID
```

- If the trigger has an INSERT event, the trigger action can be an EXECUTE PROCEDURE or EXECUTE FUNCTION statement with an INTO clause that references a column that was supplied by the trigger event or a column that was not supplied by the trigger event.

When an EXECUTE PROCEDURE or EXECUTE FUNCTION statement is the trigger action, you can specify the INTO clause for an INSERT trigger only when the trigger action occurs in the FOR EACH ROW list. In this case, the INTO clause can contain only column names from the triggering table.

The following statement illustrates the valid use of the INTO clause:

```
CREATE TRIGGER ins_totpr INSERT ON items
  REFERENCING NEW AS new_ins
  FOR EACH ROW (EXECUTE PROCEDURE calc_totpr
    (0, new_ins.quantity, 0) INTO total_price);
```

The column that follows the INTO keyword can be a column in the triggering table that was supplied by the trigger event, or a column in the triggering table that was not supplied by the trigger event.

When the INTO clause appears in the EXECUTE PROCEDURE or (for Dynamic Server only) the EXECUTE FUNCTION statement, the database server immediately updates the specified columns with values returned from the UDR.

- If the trigger action is a SELECT statement, the SELECT statement can reference the triggering table. The SELECT statement can be a trigger action in the following instances:
 - The SELECT statement appears in a subquery in the WHEN clause or in a trigger-action statement.
 - The trigger action is a UDR, and the SELECT statement appears inside the UDR.

Re-Entrancy and Cascading Triggers

The cases when a trigger cannot be re-entrant apply recursively to all cascading triggers, which are considered part of the initial trigger. In particular, this rule means that a cascading trigger cannot update any columns in the triggering table that were updated by the original triggering statement, including any nontriggering columns affected by that statement. For example, assume this UPDATE statement is the triggering statement:

```
UPDATE tab1 SET (a, b) = (a + 1, b + 1);
```

In the cascading triggers of the next example, **trig2** fails at runtime because it references column **b**, which the triggering UPDATE statement updates:

```
CREATE TRIGGER trig1 UPDATE OF a ON tab1-- Valid
  AFTER (UPDATE tab2 SET e = e + 1);

CREATE TRIGGER trig2 UPDATE OF e ON tab2-- Invalid
  AFTER (UPDATE tab1 SET b = b + 1);
```

Now consider the following SQL statements. When the final UPDATE statement is executed, column **a** is updated and the trigger **trig1** is activated.

The trigger action again updates column **a** with an EXECUTE PROCEDURE INTO statement.

```
CREATE TABLE temp1 (a INT, b INT, e INT);
INSERT INTO temp1 VALUES (10, 20, 30);

CREATE PROCEDURE proc(val iINT) RETURNING INT,INT;
  RETURN val+10, val+20;
END PROCEDURE;

CREATE TRIGGER trig1 UPDATE OF a ON temp1
  FOR EACH ROW (EXECUTE PROCEDURE proc(50) INTO a, e);

CREATE TRIGGER trig2 UPDATE OF e ON temp1
  FOR EACH ROW (EXECUTE PROCEDURE proc(100) INTO a, e);

UPDATE temp1 SET (a,b) = (40,50);
```

In Extended Parallel Server, to re-create this example, use the CREATE PROCEDURE statement instead of the CREATE FUNCTION statement.

Several questions arise from this example of cascading triggers. First, should the update of column **a** activate trigger **trig1** again? The answer is no. Because the trigger was activated, it is not activated a second time. If the trigger action is an

EXECUTE PROCEDURE INTO or EXECUTE FUNCTION INTO statement, the only triggers that are activated are those that are defined on columns that are mutually exclusive from the columns updated until then (in the cascade of triggers) in that table. Other triggers are ignored.

Another question that arises from the example is whether trigger **trig2** should be activated. The answer is yes. The trigger **trig2** is defined on column **e**. Until now, column **e** in table **temp1** has not been modified. Trigger **trig2** is activated.

A final question that arises from the example is whether triggers **trig1** and **trig2** should be activated after the trigger action in **trig2** is performed. The answer is no. Neither trigger is activated. By this time columns **a** and **e** have been updated once, and triggers **trig1** and **trig2** have been executed once. The database server ignores and does not activate these triggers. For more about cascading triggers, see “Cascading Triggers” on page 2-267.

In Dynamic Server, as noted earlier, an INSTEAD OF trigger on a view cannot include the EXECUTE PROCEDURE INTO statement among its trigger actions. In addition, an error results if two views each have INSERT INSTEAD OF triggers with actions defined to perform INSERT operations on the other view.

Rules for SPL Routines

In addition to the rules listed in “Re-Entrancy of Triggers” on page 2-262, the following guidelines apply to an SPL routine that is specified as a trigger action:

- The SPL routine cannot be a cursor function (one that returns more than one row) in a context where only one row is expected.
- You cannot use the old or new correlation name inside the SPL routine unless the CREATE FUNCTION or CREATE PROCEDURE statement includes the REFERENCING clause that defines the UDR as a trigger routine. If you need to use the corresponding values in a routine that is not a trigger routine, you must pass them as parameters. In this case, the routine should be independent of triggers, and the old or new correlation name does not have any meaning outside the trigger.
- A trigger routine must include the REFERENCING clause that can declare a correlation name for OLD or NEW column values that SPL statements in the trigger routine can reference.
- A trigger routine must include the FOR *table_object* clause that specifies the name of the table or view in the local database whose triggers can invoke this routine. The triggered action cannot call a trigger routine that does not specify the triggering table or view.
- Only trigger routines invoked in the FOR EACH ROW section of the Triggered Action list can operate directly on old or new correlation names that are defined in the REFERENCING clause of the trigger or of the trigger routine.
- Trigger routines can be invoked only in the FOR EACH ROW section of the Triggered Action list in the trigger definition.
- Correlated variables for OLD or NEW values can appear in the IF statement of SPL and in CASE expressions.
- Only correlated variables for NEW values can be on the left-hand side of a LET expression that references correlated variables. In this case, the FOR clause of the SPL routine must specify a table, rather than a view, and the trigger whose action invokes the SPL routine cannot be an INSTEAD OF trigger.
- Both OLD and NEW values can be on the right-hand side of a LET expression.

- Only trigger routines that are invoked in the FOR EACH ROW clause can use the Boolean operators SELECTING, INSERTING, DELETING, and UPDATING. These operators return TRUE ('t') if the triggering event matches the DML operation referenced by the name of the operator, and they return FALSE ('f') otherwise.
- The IF statement of SPL and CASE expressions of SQL can specify these operators as the condition in a trigger routine.
- Trigger routines must be written in the SPL language. They cannot be written in an external language, such as the C or Java language, but the trigger routine can include calls to external language routines, such as the `mi_trigger` application programming interface for trigger introspection.
- Trigger routines cannot reference savepoint objects. Any changes to the data values or to the schema of the database by a triggered action must be committed or rolled back in their entirety. Dynamic Server does not support the partial rollback of a triggered action.

For more information about the `mi_trigger` API, refer to the *IBM Informix Datablade API Programmer's Guide* and to the *IBM Informix Datablade API Function Reference*.

When you use an SPL routine as a trigger action, the database objects that the routine references are not checked until the routine is executed.

See also the SPL restrictions in “Triggers and SPL Routines” on page 2-244.

Privileges to Execute Trigger Actions

If you do not own the trigger, but the access privileges held by the trigger owner include WITH GRANT OPTION, then for each triggered SQL statement you inherit the privileges of the trigger owner (with grant option), in addition to any privileges that have been granted to you individually, or through an active or default role that you hold, or that you hold as a member of the PUBLIC group. If the triggered action calls a UDR, you need Execute privilege on the UDR, or the trigger owner must have Execute privilege with grant option.

While executing the UDR, however, you do not inherit the privileges of the trigger owner; instead, you receive the privileges granted with the UDR, depending on whether the routine is a DBA-privileged or an owner-privileged UDR:

1. Privileges for a DBA-privileged UDR

When a UDR is registered with the DBA keyword, and you are granted the Execute privilege on the UDR, the database server automatically grants you temporary DBA privileges that are available only when you are executing the UDR.

2. Privileges for an owner-privileged UDR

If the UDR was created without the DBA keyword, but the owner of the UDR was granted the necessary privileges on the underlying database objects with the WITH GRANT OPTION keywords, then you inherit these privileges when you are granted the Execute privilege on the UDR.

For a UDR that is not DBA privileged, all non-qualified database objects that the UDR references are implicitly qualified by the name of the UDR owner.

If the UDR owner has no WITH GRANT OPTION privilege, you have your original privileges on the underlying database objects when the UDR executes. For more information on privileges on SPL routines, refer to the *IBM Informix Guide to SQL: Tutorial*.

In Dynamic Server, a view that has no INSTEAD OF trigger has only Select (with grant option) privilege. If an INSTEAD OF trigger is created on it, however, then the view has Insert (with grant option) privilege during creation of the trigger. The view owner can now grant only Select and Insert privileges to others. This is independent of the trigger action. It is not necessary to obtain Execute (with grant option) privilege on the procedure or function. By default, Execute privilege (without grant option) is granted on each UDR in the action list.

You can use roles with triggers. Role-related statements (CREATE ROLE, DROP ROLE, GRANT, REVOKE, and SET ROLE) and SET SESSION AUTHORIZATION statements are valid in a UDR that the triggered action invokes. Privileges that a user acquired by enabling a role or by a SET SESSION AUTHORIZATION statement are not relinquished when a trigger is executed.

On a complex view (one with columns from more than one table), only the owner or DBA can create an INSTEAD OF trigger. The owner receives Select privileges when the trigger is created. Only after obtaining the required Execute privileges can the owner of the view grant privileges to other users. When the trigger on the complex view is dropped, all of these privileges are revoked.

Creating a Trigger Action That Anyone Can Use

For a trigger to be executable by anyone who has the privileges to execute the triggering statement, you can ask the DBA to create a DBA-privileged UDR and grant you the Execute privilege with the WITH GRANT OPTION right.

You then use the DBA-privileged UDR as the trigger action. Anyone can execute the trigger action because the DBA-privileged UDR carries the WITH GRANT OPTION right. When you activate the UDR, the database server applies privilege-checking rules for a DBA.

Cascading Triggers

In this section and in sections that follow, references to *nonlogging databases* do not apply to Extended Parallel Server. (In Extended Parallel Server, all databases support transaction logging.)

The database server allows triggers other than Select triggers to cascade, meaning that the trigger actions of one trigger can activate another trigger. (For further information on the restriction against cascading Select triggers, see “Circumstances When a Select Trigger Is Activated” on page 2-248.)

The maximum number of triggers in a cascading series is 61: the initial trigger plus a maximum of 60 cascading triggers. When the number of cascading triggers in a series exceeds the maximum, the database server returns error number -748, with the following message:

```
Exceeded limit on maximum number of cascaded triggers.
```

The next example illustrates a series of cascading triggers that enforce referential integrity on the **manufact**, **stock**, and **items** tables in the **stores_demo** database. When a manufacturer is deleted from the **manufact** table, the first trigger, **del_manu**, deletes all the items of that manufacturer from the **stock** table. Each DELETE in the **stock** table activates a second trigger, **del_items**, that deletes all **items** of that manufacturer from the **items** table. Finally, each DELETE in the **items** table triggers SPL routine **log_order**, creating a record of any orders in the **orders** table that can no longer be filled.

```

CREATE TRIGGER del_manu
  DELETE ON manufact REFERENCING OLD AS pre_del
  FOR EACH ROW(DELETE FROM stock WHERE manu_code = pre_del.manu_code);
CREATE TRIGGER del_stock
  DELETE ON stock REFERENCING OLD AS pre_del
  FOR EACH ROW(DELETE FROM items WHERE manu_code = pre_del.manu_code);
CREATE TRIGGER del_items
  DELETE ON items REFERENCING OLD AS pre_del
  FOR EACH ROW(EXECUTE PROCEDURE log_order(pre_del.order_num));

```

When you are not using logging, referential integrity constraints on both the **manufact** and **stock** tables prohibit the triggers in this example from executing. When you use logging, however, the triggers execute successfully because constraint checking is deferred until all the trigger actions are complete, including the actions of cascading triggers. For more information about how constraints are handled when triggers execute, see “Constraint Checking.”

The database server prevents loops of cascading triggers by not allowing you to modify the triggering table in any cascading trigger action, except with an UPDATE statement that does not modify any column that the triggering UPDATE statement updated, or with an INSERT statement. An INSERT trigger can define UPDATE trigger actions on the same table.

Constraint Checking

When you use logging, the database server defers constraint checking on the triggering statement until after the statements in the triggered-action list execute. This is equivalent to executing a SET CONSTRAINTS ALL DEFERRED statement before executing the triggering statement. After the trigger action is completed, the database server effectively executes a SET CONSTRAINTS *constraint* IMMEDIATE statement to check the constraints that were deferred. This action allows you to write triggers so that the trigger action can resolve any constraint violations that the triggering statement creates. For more information, see “SET Database Object Mode” on page 2-591.

Consider the following example, in which the table **child** has constraint **r1**, which references the table **parent**. You define trigger **trig1** and activate it with an INSERT statement. In the trigger action, **trig1** checks to see if **parent** has a row with the value of the current **cola** in **child**; if not, it inserts it.

```

CREATE TABLE parent (cola INT PRIMARY KEY);
CREATE TABLE child (cola INT REFERENCES parent CONSTRAINT r1);
CREATE TRIGGER trig1 INSERT ON child
  REFERENCING NEW AS new
  FOR EACH ROW
  WHEN((SELECT COUNT (*) FROM parent
        WHERE cola = new.cola) = 0)
  -- parent row does not exist
  (INSERT INTO parent VALUES (new.cola));

```

When you insert a row into a table that is the child table in a referential constraint, the row might not exist in the parent table. The database server does not immediately return this error on a triggering statement. Instead, it allows the trigger action to resolve the constraint violation by inserting the corresponding row into the parent table. As the previous example shows, you can check within the trigger action to see whether the parent row exists, and if so, you can provide logic to bypass the INSERT action.

For a database without logging, the database server does not defer constraint checking on the triggering statement. In this case, the database server immediately returns an error if the triggering statement violates a constraint.

You cannot use the SET Transaction Mode statement in a trigger action. The database server checks this restriction when you activate a trigger, because the statement could occur inside a UDR.

In Extended Parallel Server, rows that cause constraint violations might appear in the violations table, even if a subsequent trigger action corrects the violation.

Preventing Triggers from Overriding Each Other

When you activate multiple triggers with an UPDATE statement, a trigger can possibly override the changes that an earlier trigger made. If you do not want the trigger actions to interact, you can split the UPDATE statement into multiple UPDATE statements, each of which updates an individual column.

As another alternative, you can create a single update trigger for all columns that require a trigger action. Then, inside the trigger action, you can test for the column being updated and apply the actions in the desired order. This approach, however, is different from having the database server apply the actions of individual triggers, and it has the following disadvantages:

- If the triggering UPDATE statement sets a column to the current value, you cannot detect the UPDATE, so the trigger action is skipped. You might wish to execute the trigger action, even though the value of the column has not changed.
- If the trigger has a BEFORE action, it applies to all columns, because you cannot yet detect whether a column has changed.

External Tables

The trigger action can affect tables of other database servers. The following example shows an Update trigger on **dbserver1**, which triggers an UPDATE to **items** on **dbserver2**:

```
CREATE TRIGGER upd_nt UPDATE ON newtab
  REFERENCING NEW AS post
  FOR EACH ROW(UPDATE stores_demo@dbserver2:items
    SET quantity = post.qty WHERE stock_num = post.stock
    AND manu_code = post.mc);
```

If, however, a statement from an external database server initiates a trigger whose action affects tables in an external database, the trigger actions fail.

For example, the following combination of trigger action and triggering statement results in an error when the triggering statement executes:

```
-- Trigger action from dbserver1 to dbserver3:

CREATE TRIGGER upd_nt UPDATE ON newtab
  REFERENCING NEW AS post
  FOR EACH ROW(UPDATE stores_demo@dbserver3:items
    SET quantity = post.qty WHERE stock_num = post.stock
    AND manu_code = post.mc);
-- Triggering statement from dbserver2:

UPDATE stores_demo@dbserver1:newtab
  SET qty = qty * 2 WHERE s_num = 5
  AND mc = 'ANZ';
```

Logging and Recovery

You can create triggers for databases, with and without logging. If the trigger fails in a database that has transaction logging, the triggering statement and trigger actions are rolled back, as if the actions were an extension of the triggering statement, but the rest of the transaction is not rolled back.

In a database that does not have transaction logging, however, you cannot roll back when the triggering statement fails. In this case, you are responsible for maintaining data integrity in the database. The UPDATE, INSERT, or DELETE action of the triggering statement occurs before the trigger actions in the FOR EACH ROW section. If the trigger action fails for a database without logging, the application must restore the row that was changed by the triggering statement to its previous value.

If a trigger action calls a UDR, but the UDR terminates in an exception-handling section, any actions that modify data inside that section are rolled back with the triggering statement. In the following partial example, when the exception handler traps an error, it inserts a row into the table **logtab**:

```
ON EXCEPTION IN (-201)
  INSERT INTO logtab values (errno, errstr);
  RAISE EXCEPTION -201
END EXCEPTION;
```

When the RAISE EXCEPTION statement returns the error, however, the database server rolls back this INSERT because it is part of the trigger actions. If the UDR is executed outside a trigger action, the INSERT is not rolled back.

The UDR that implements a trigger action cannot contain any BEGIN WORK, COMMIT WORK, or ROLLBACK WORK statements. If the database has transaction logging, you must either begin an explicit transaction before the triggering statement, or the statement itself must be an implicit transaction. In any case, no other transaction-related statement is valid inside the UDR.

You can use triggers to enforce referential actions that the database server does not currently support. In a database without logging, you are responsible for maintaining data integrity when the triggering statement fails.

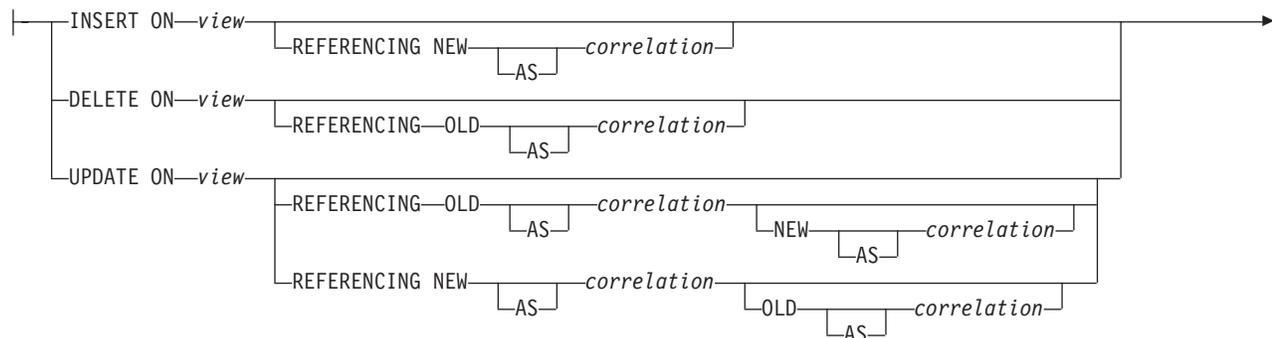
INSTEAD OF Triggers on Views (IDS)

Use INSTEAD OF triggers to perform a specified trigger action on a view, rather than execute the triggering INSERT, DELETE, or UPDATE event.

Syntax

```
►► CREATE TRIGGER trigger INSTEAD OF Trigger on a View [ENABLED | DISABLED]
```

Trigger on a View:



►FOR EACH ROW—| INSTEAD OF Triggered Action | (1)

Notes:

1 See “The Action Clause of INSTEAD OF Triggers (IDS)”

Element	Description	Restrictions	Syntax
<i>correlation</i>	Name that you declare here to qualify an old or new column value (as <i>correlation.column</i>) in a triggered action	Must be unique in this statement	“Identifier ” on page 5-21
<i>trigger</i>	Name declared here for the trigger	Must be unique among the names of triggers in the database	“Identifier ” on page 5-21
<i>view</i>	Name or synonym of the triggering view. Can include <i>owner.</i> qualifier.	The view or synonym must exist in the current database	“Identifier ” on page 5-21

You can use the trigger action to update the tables underlying the view, in some cases updating an otherwise “non-updatable” view. You can also use INSTEAD OF triggers to substitute other actions when INSERT, DELETE, or UPDATE statements reference specific columns within the database.

In the optional REFERENCING clause of an INSTEAD OF UPDATE trigger, the *new* correlation name can appear before or after the *old* correlation name.

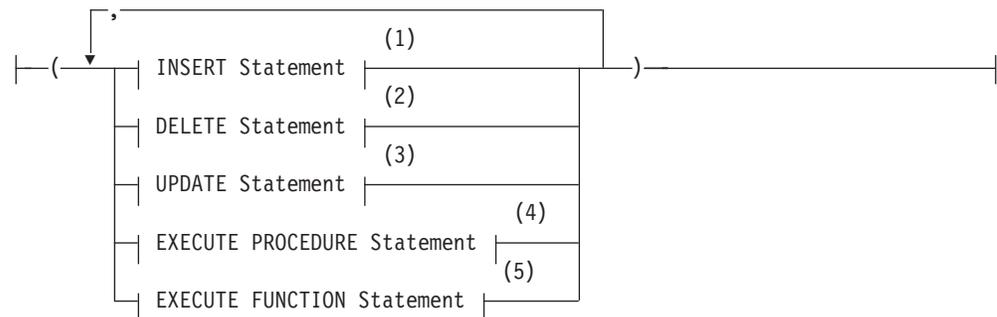
With Dynamic Server, the same REFERENCING OLD and REFERENCING NEW syntax is supported in the CREATE FUNCTION and CREATE PROCEDURE statements for defining correlation names in trigger routines. A trigger routine can be invoked in the Action clause for INSTEAD OF triggers on the view that is specified in the FOR clause of the CREATE FUNCTION or CREATE PROCEDURE statement that defines the trigger routine.

The specified *view* is sometimes called the *triggering view*. The left-hand portion of this diagram (including the *view* specification) defines the *trigger event*. The rest of the diagram defines correlation names and the *trigger action*.

The Action Clause of INSTEAD OF Triggers (IDS)

When the trigger event for the specified *view* is encountered, the SQL statements of the trigger action are executed, instead of the triggering statement. Triggers defined on a view support the following syntax in the action clause.

INSTEAD OF Triggered Action:



Notes:

- 1 See "INSERT " on page 2-430
- 2 See "DELETE " on page 2-303
- 3 See "UPDATE" on page 2-689
- 4 See "EXECUTE PROCEDURE " on page 2-364
- 5 See "EXECUTE FUNCTION " on page 2-356

This is not identical to the syntax of the trigger action for a trigger on a table, as described in the section "Triggered Action" on page 2-257. Because no WHEN (*condition*) is supported, the same trigger action is executed whenever the INSTEAD OF trigger event is encountered, and only one action list can be specified, rather than a separate list for each *condition*.

Restrictions on INSTEAD OF Triggers on Views (IDS)

You must be either the owner of the *view* or have the DBA status to create an INSTEAD OF trigger on a view. The owner of a simple view (based on only one table) has Insert, Update, and Delete privileges. For information about the relationship between the privileges of the trigger owner and the privileges of other users, see "Privileges to Execute Trigger Actions" on page 2-266.

If multiple tables underlie a view, only the owner of the view can create a trigger, but that owner can grant DML privileges on the view to other users.

An INSTEAD OF trigger defined on a view cannot violate the "Restrictions on Triggers" on page 2-242 and must observe the following additional rules:

- You can define an INSTEAD OF trigger only on a view, not on a table.
- The view must be local to the current database.
- The view cannot be an updatable view WITH CHECK OPTION.
- No SELECT event or WHEN clause is valid in an INSTEAD OF trigger.
- No BEFORE nor AFTER action is valid in an INSTEAD OF trigger.
- No OF *column* clause is valid in an INSTEAD OF UPDATE trigger.
- Every INSTEAD OF trigger must specify FOR EACH ROW.
- UDRs called in a trigger action cannot reference savepoint objects.

A view can have any number of INSTEAD OF triggers defined for each type of event (INSERT, DELETE, or UPDATE).

Updating Views

INSERT, DELETE, or UPDATE statements can directly modify a view only if all of the following are true of the SELECT statement that defines the view:

- All of the columns in the view are from a single table.
- No columns in the projection list are aggregate values.
- No UNIQUE or DISTINCT keyword is in the SELECT projection list.
- No GROUP BY clause nor UNION operator is in the view definition.
- The query selects no calculated values and no literal values.

By using INSTEAD OF triggers, however, you can circumvent these restrictions on the view, if the trigger action modifies the base table.

Example of an INSTEAD OF Trigger on a View

Suppose that **dept** and **emp** are tables that list departments and employees:

```
CREATE TABLE dept (  
    deptno INTEGER PRIMARY KEY,  
    deptname CHAR(20),  
    manager_num INT  
);  
CREATE TABLE emp (  
    empno INTEGER PRIMARY KEY,  
    empname CHAR(20),  
    deptno INTEGER REFERENCES dept(deptno),  
    startdate DATE  
);  
ALTER TABLE dept ADD CONSTRAINT(FOREIGN KEY (manager_num)  
    REFERENCES emp(empno));
```

The next statement defines **manager_info**, a view of columns in the **dept** and **emp** tables that includes all the managers of each department:

```
CREATE VIEW manager_info AS  
    SELECT d.deptno, d.deptname, e.empno, e.empname  
    FROM emp e, dept d WHERE e.empno = d.manager_num;
```

The following CREATE TRIGGER statement creates **manager_info_insert**, an INSTEAD OF trigger that is designed to insert rows into the **dept** and **emp** tables through the **manager_info** view:

```
CREATE TRIGGER manager_info_insert  
    INSTEAD OF INSERT ON manager_info      --defines trigger event  
    REFERENCING NEW AS n                  --new manager data  
    FOR EACH ROW                          --defines trigger action  
    (EXECUTE PROCEDURE instab(n.deptno, n.empno));  
  
CREATE PROCEDURE instab (dno INT, eno INT)  
    INSERT INTO dept(deptno, manager_num) VALUES(dno, eno);  
    INSERT INTO emp (empno, deptno) VALUES (eno, dno);  
END PROCEDURE;
```

After the tables, view, trigger, and SPL routine have been created, the database server treats the following INSERT statement as a triggering event:

```
INSERT INTO manager_info(deptno, empno) VALUES (08, 4232);
```

This triggering INSERT statement is not executed, but this event causes the trigger action to be executed instead, invoking the **instab()** SPL routine. The INSERT statements in the SPL routine insert new values into both the **emp** and **dept** base tables of the **manager_info** view.

Related Information

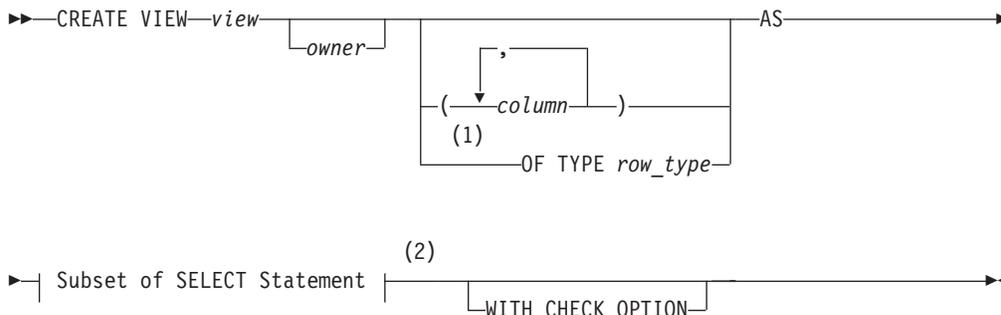
Related statements: “CREATE FUNCTION ” on page 2-113, “CREATE PROCEDURE ” on page 2-153, “CREATE VIEW ” on page 2-274, “DROP TRIGGER ” on page 2-344, “EXECUTE PROCEDURE ” on page 2-364, and “SET Database Object Mode” on page 2-591

For a task-oriented discussion of triggers, and for examples of INSTEAD OF DELETE (and UPDATE) triggers on views, see the *IBM Informix Guide to SQL: Tutorial*. For performance implications of triggers, see your *IBM Informix Performance Guide*.

CREATE VIEW

Use the CREATE VIEW statement to create a new view that is based on one or more existing tables and views that reside in the database, or in another database of the local database server or of a different database server.

Syntax



Notes:

- 1 Dynamic Server only
- 2 See "Subset of SELECT Statements Valid in View Definitions" on page 2-276

Element	Description	Restrictions	Syntax
<i>column</i>	Name that you declare here for a column in view. Default is a column name from Projection list of SELECT.	See "Naming View Columns" on page 2-276.	"Identifier " on page 5-21
<i>owner</i>	Owner of the view. If omitted, default is the user ID that issues the statement.	To specify another user ID requires DBA access privilege.	" Owner Name " on page 5-45
<i>row_type</i>	Named-row type for typed view	Must already exist in the database	"Data Type " on page 4-21
<i>view</i>	Name that you declare here for the view	Must be unique among view, table, sequence, and synonym names in the database.	"Identifier " on page 5-21

Usage

A *view* is a virtual table, defined by a SELECT statement. Except for the statements in the following list, you can specify the name or synonym of a view in any SQL statement where the name of a table is syntactically valid:

- ALTER FRAGMENT
- CREATE INDEX
- CREATE TABLE
- CREATE TRIGGER
- RENAME TABLE
- START VIOLATIONS TABLE
- STOP VIOLATIONS TABLE
- TRUNCATE
- UPDATE STATISTICS

You must specify the name of a view when you use the CREATE TRIGGER statement to define an INSTEAD OF trigger on a view, but the syntax and functionality are different from those of a trigger defined on a table.

“Updating Through Views” on page 2-278 prohibits non-updatable views in INSERT, DELETE, or UPDATE statements (where other views are valid).

To create a view, you must have the Select privilege on all columns from which the view is derived. You can query a view as if it were a table, and in some cases, you can update it as if it were a table; but a view is not a table.

The view consists of the set of rows and columns that the SELECT statement in the view definition returns each time you refer to the view in a query.

In some cases, the database server merges the SELECT statement of the user with the SELECT statement defining the view and executes the combined statements. In other cases, a query against a view might execute more slowly than expected, if the complexity of the view definition causes the database server to create a temporary table (referred to as a materialized view). For more information on materialized views, see the *IBM Informix Performance Guide*.

The view reflects changes to the underlying tables with one exception. If a SELECT * specification defines the view, the view has only the columns that existed in the underlying tables when the view was defined by CREATE VIEW. Any new columns that are subsequently added to the underlying tables with the ALTER TABLE statement do not appear in the view.

The view inherits the data types of the columns in the tables from which the view is derived. The database server determines data types of virtual columns from the nature of the expression.

The SELECT statement is stored in the **sysviews** system catalog table. When you subsequently refer to a view in another statement, the database server performs the defining SELECT statement while it executes the new statement.

In DB–Access, if you create a view outside the CREATE SCHEMA statement, you receive warnings if you use the **-ansi** flag or if you set the **DBANSIWARN** environment variable.

The following statement creates a view that is based on the **person** table. When you create a view like this, which has no OF TYPE clause, the view is referred to as an *untyped view*.

```
CREATE VIEW v1 AS SELECT * FROM person;
```

Typed Views

You can create *typed views* if you have Usage privileges on the named-ROW type or if you are its owner or the DBA. If you omit the OF TYPE clause, rows in the view are considered untyped and default to an unnamed-ROW type.

Typed views, like typed tables, are based on a named-ROW type. Each column in the view corresponds to a field in the named-ROW type. The following statement creates a typed view that is based on the table **person**.

```
CREATE VIEW v2 OF TYPE person_t AS SELECT * FROM person;
```

To create a typed view, you must include an OF TYPE clause. When you create a typed view, the named-ROW type that you specify immediately after the OF TYPE keywords must already exist.

Subset of SELECT Statements Valid in View Definitions

You cannot create a view on a temporary table. The FROM clause of the SELECT statement cannot include the name of a temporary table.

If Select privileges are revoked from a user for a table that is referenced in the SELECT statement defining a view that the same user owns, then that view is dropped, unless it also includes columns from tables in another database.

You cannot create a view on typed tables (including any table that is part of a table hierarchy) that reside in a remote database.

Do not use display labels in the select list of the Projection clause. Display labels in the Projection clause are interpreted as column names.

The SELECT statement in CREATE VIEW cannot include the SKIP, FIRST, or LIMIT keywords, the INTO TEMP clause, or the ORDER BY clause. For complete information about SELECT statement syntax and usage, see “SELECT” on page 2-528.

Union Views

A view that contains a UNION or UNION ALL operator in its SELECT statement is known as a *union view*. Certain restrictions apply to union views:

- If a CREATE VIEW statement defines a union view, you cannot specify the WITH CHECK OPTION keywords in the CREATE VIEW statement.
- All restrictions that apply to UNION or UNION ALL operations in stand-alone SELECT statements also apply to UNION and UNION ALL operations in the SELECT statement of a union view.

For a list of these restrictions, see “Restrictions on a Combined SELECT” on page 2-578. For an example of a CREATE VIEW statement that defines a union view, see “Naming View Columns.”

Naming View Columns

The number of columns that you specify in the *column* list must match the number of columns returned by the SELECT statement that defines the view. If you do not specify a list of columns, the view inherits the column names of the underlying tables. In the following example, the view **herostock** has the same column names as the columns in Projection clause of the SELECT statement:

```
CREATE VIEW herostock AS
  SELECT stock_num, description, unit_price, unit, unit_descr
  FROM stock WHERE manu_code = 'HR0';
```

You must specify at least one column name in the following circumstances:

- If you provide names for some of the columns in a view, then you must provide names for all the columns. That is, the column list must contain an entry for every column that appears in the view.
- If the SELECT statement returns an expression, the corresponding column in the view is called a *virtual* column. You must provide a name for a virtual column.

In the following example, the user must specify the column parameter because the select list of the Projection clause of the SELECT statement contains an aggregate expression:

```
CREATE VIEW newview (firstcol, secondcol) AS
    SELECT sum(colA), colB FROM oldtab;
```

- You must also specify column names in cases where any of the selected columns have duplicate column names without the table qualifiers. For example, if both **orders.order_num** and **items.order_num** appear in the SELECT statement, the CREATE VIEW statement, must provide two separate column names to label them:

```
CREATE VIEW someorders (custnum,ocustnum,newprice) AS
    SELECT orders.order_num,items.order_num,
           items.total_price*1.5
    FROM orders, items
    WHERE orders.order_num = items.order_num
    AND items.total_price > 100.00;
```

Here **custnum** and **ocustnum** replace the two identical column names.

- The CREATE VIEW statement must also provide column names in the column list when the SELECT statement includes a UNION or UNION ALL operator and the names of the corresponding columns in the SELECT statements are not identical.

For example, code in the following CREATE VIEW statement must specify the column list because the second column in the first SELECT statement has a different name from the second column in the second SELECT statement:

```
CREATE VIEW myview (colA, colB) AS
    SELECT colX, colY from firsttab
    UNION
    SELECT colX, colZ from secondtab;
```

Using a View in the SELECT Statement

You can define a view whose columns are based on other views, but you must abide by the restrictions on creating views that are discussed in the *IBM Informix Database Design and Implementation Guide*.

WITH CHECK OPTION Keywords

The WITH CHECK OPTION keywords instruct the database server to ensure that all modifications that are made through the view to the underlying tables satisfy the definition of the view.

The following example creates a view that is named **palo_alto**, which uses all the information in the **customer** table for customers in the city of Palo Alto. The database server checks any modifications made to the **customer** table through **palo_alto** because the WITH CHECK OPTION keywords are specified.

```
CREATE VIEW palo_alto AS
    SELECT * FROM customer WHERE city = 'Palo Alto'
    WITH CHECK OPTION
```

You can insert into a view a row that does not satisfy the conditions of the view (that is, a row that is not visible through the view). You can also update a row of a view so that it no longer satisfies the conditions of the view. For example, if the view was created without the WITH CHECK OPTION keywords, you could insert a row through the view where the city is Los Altos, or you could update a row through the view by changing the city from Palo Alto to Los Altos.

To prevent such inserts and updates, you can add the WITH CHECK OPTION keywords when you create the view. These keywords ask the database server to test every inserted or updated row to ensure that it meets the conditions that are set by the WHERE clause of the view. The database server rejects the operation with an error if the row does not meet the conditions.

Even if the view was created with the WITH CHECK OPTION keywords, however, you can perform inserts and updates through the view to change columns that are not part of the view definition. A column is not part of the view definition if it does not appear in the WHERE clause of the SELECT statement that defines the view.

Updating Through Views

If a view is built on a single table, the view is *updatable* if the SELECT statement that defines the view does not contain any of the following elements:

- Columns in the projection list that are aggregate values
- Columns in the projection list that use the UNIQUE or DISTINCT keyword
- A GROUP BY clause
- A UNION operator
- A query that selects calculated or literal values.

You can DELETE from a view that selects calculated values from a single table, but INSERT and UPDATE operations are not valid on such views.

In an updatable view, you can update the values in the underlying table by inserting values into the view. If a view is built on a table that has a derived value for a column, however, that column is not updatable through the view. Other columns in the view, however, can be updated.

See also “Updating Views” on page 2-272 for information about using INSTEAD OF triggers to update views that are based on more than one table or that include columns containing aggregates or other calculated values.

Important: You cannot update or insert rows in a remote table through views that were created using the WITH CHECK OPTION keywords.

Related Information

Related statements: “CREATE TABLE ” on page 2-191, “CREATE TRIGGER” on page 2-239, “DROP VIEW ” on page 2-346, “GRANT ” on page 2-395, “REVOKE” on page 2-494, “SELECT” on page 2-528, and “SET SESSION AUTHORIZATION ” on page 2-650

For a discussion of views, see the *IBM Informix Database Design and Implementation Guide*.

For a discussion of how to use privileges and views to restrict access to the database, see the *IBM Informix Guide to SQL: Tutorial*.

CREATE XADATASOURCE

Use the CREATE XADATASOURCE statement to create a new XA-compliant data source and create an entry for it in the **sysxadatasources** system catalog table. Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

►► CREATE XADATASOURCE *xa_source* USING *xa_type* ◀◀

Element	Description	Restrictions	Syntax
<i>xa_source</i>	Name that you declare here for the new XA data source	Must be unique among XA data source names in sysxadatasources	"Identifier " on page 5-21
<i>xa_type</i>	Name of an existing XA data source type	Must already exist in the database in the sysxasourcetypes system catalog table	"Identifier " on page 5-21

Usage

An XA-compliant data source is an external data source that complies with the X/Open DTP XA Standard for managing interactions between a transaction manager and a resource manager. To register XA-compliant data sources in the database requires two SQL statements:

- First create one or more XA-compliant data source types by using the CREATE XADATASOURCE TYPE statement.
- Then create one or more instances of XA-compliant data sources with the CREATE XADATASOURCE statement.

You can integrate transactions at the XA data source with the Dynamic Server transaction, using a 2-phase commit protocol to ensure that transactions are uniformly committed or rolled back among multiple database servers, and manage multiple external XA data sources within the same global transaction.

Any user can create an XA data source, which follows standard owner-naming rules, according to the ANSI-compliance status of the database. Only a database that uses transaction logging can support the X/Open DTP XA Standard.

Both XA data source types and instances of XA data sources are specific to one database. To support distributed transactions, they must be created in each database that interacts with the external XA data source.

The following statement creates a new XA data source instance called **informix.NewYork** of type **informix.MQSeries**.

```
CREATE XADATASOURCE informix.NewYork USING informix.MQSeries;
```

Related Information

Related statements: "CREATE XADATASOURCE TYPE " on page 2-280, "DROP XADATASOURCE" on page 2-347, "DROP XADATASOURCE TYPE" on page 2-347

For the schema of the **sysxadatasources** system catalog table, see the documentation notes to the *IBM Informix Guide to SQL: Reference*. For more

information on XA data sources, see the *IBM Informix Datablade API Programmer's Guide*.

CREATE XADATASOURCE TYPE

Use the CREATE XADATASOURCE TYPE statement to create a new XA-compliant data source type and create an entry for it in the **sysxasourcetypes** system catalog table. Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```

▶▶ CREATE XADATASOURCE TYPE xa_type ( Purpose Options (1) )

```

Notes:

- 1 See "Purpose Options" on page 5-50

Element	Description	Restrictions	Syntax
<i>xa_type</i>	Name that you declare here for a new XA data source type	Must be unique among XA data source type names in the sysxasourcetypes system catalog table	"Identifier " on page 5-21

Usage

The CREATE XADATASOURCE TYPE statement adds an XA-compliant data source type to the database.

Any user can create an XA data source type, whose owner-naming rules depend on the ANSI-compliance status of the database. Only a database that uses transaction logging can support the X/Open DTP XA Standard.

To create a data source type, you must declare its name and specify *purpose functions* and *purpose values* as attributes of the XA source type. Most of the purpose options that follow the source type name associate columns in the **sysxasourcetypes** system catalog table with the name of a UDR.

Both XA data source types and instances of XA data sources are specific to one database. To support distributed transactions, they must be created in each database that interacts with the external XA data source.

The following statement creates a new XA data source type called **MQSeries®**, owned by user **informix**.

```

CREATE XADATASOURCE TYPE 'informix'.MQSeries(
    xa_flags      = 1,
    xa_version    = 0,
    xa_open       = informix.mqseries_open,
    xa_close      = informix.mqseries_close,
    xa_start      = informix.mqseries_start,
    xa_end        = informix.mqseries_end,
    xa_rollback   = informix.mqseries_rollback,
    xa_prepare    = informix.mqseries_prepare,
    xa_commit     = informix.mqseries_commit,

```

```

xa_recover = informix.mqseries_recover,
xa_forget  = informix.mqseries_forget,
xa_complete = informix.mqseries_complete);

```

You need to provide one value or UDR name for each of the options listed above, but the sequence in which you list them is not critical. (The order of purpose options in this example corresponds to the order of column names in the **sysxasourcetypes** system catalog table.)

After this statement executes successfully, you can create instances of type **informix.MQSeries**. The following statement creates a new instance called **informix.MenloPark** of XA-compliant data source type **informix.MQSeries**:

```
CREATE XADATASOURCE informix.MenloPark USING informix.MQSeries;
```

Related Information

Related statements: “CREATE XADATASOURCE ” on page 2-279, “DROP XADATASOURCE” on page 2-347, “DROP XADATASOURCE TYPE” on page 2-347

For information about how to enter purpose-option specifications, see “Purpose Options” on page 5-50.

For information on how to use XA data sources, see the *IBM Informix Datablade API Programmer’s Guide*.

For the schema of the **sysxasourcetypes** table, see the documentation notes to the *IBM Informix Guide to SQL: Reference*.

DATABASE

Use the DATABASE statement to open an accessible database as the current database.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax

```

▶▶ DATABASE database [EXCLUSIVE] ▶▶

```

Element	Description	Restrictions	Syntax
<i>database</i>	Name of the database	The database must exist	“Database Name” on page 5-15

Usage

You can use the DATABASE statement to select any database on your database server. To select a database on another database server, specify the name of the database server with the database name.

If you include the name of the current (or another) database server with the database name, the database server name cannot be uppercase. (See “Database Name” on page 5-15 for the syntax of specifying the database server name.)

Issuing a DATABASE statement when a database is already open closes the current database before opening the new one. Closing the current database releases any cursor resources that the database server holds, invalidating any cursors that you have declared up to that point. If the *user* specification was changed through a SET SESSION AUTHORIZATION statement, the original user name is restored when the new database is opened.

If a previous CONNECT statement has established an explicit connection to a database, and that connection is still your current connection, you cannot use the DATABASE statement (nor any SQL statement that creates an implicit connection) until after you use DISCONNECT to close the explicit connection.

The current user (or PUBLIC) must have the Connect privilege on the database that is specified in the DATABASE statement. The current user cannot have the same user name as an existing role in the database.

DATABASE is not a valid statement in multistatement PREPARE operations.

SQLCA.SQLWARN Settings Immediately after DATABASE Executes (ESQL/C)

Immediately after DATABASE executes, you can identify characteristics of the specified database by examining warning flags in the `sqlca` structure.

- If the first field of `sqlca.sqlwarn` is blank, then no warnings were issued.
- The second `sqlca.sqlwarn` field is set to the letter W if the *database* that was opened supports transaction logging.
- The third field is set to W if *database* is an ANSI-compliant database.
- The fourth field is set to W if *database* is a Dynamic Server database.
- The fifth field is set to W if *database* converts all floating-point data to DECIMAL format. (System lacks FLOAT and SMALLFLOAT support.)
- The seventh field is set to W if *database* is the secondary server (that is, running in read-only mode) in a data-replication pair.
- The eighth field is set to W if *database* has DB_LOCALE set to a locale different from the DB_LOCALE setting on the client system.

EXCLUSIVE Keyword

The EXCLUSIVE keyword opens the database in exclusive mode and prevents access by anyone but the current user. To allow others to access the database, you must first execute the CLOSE DATABASE statement and then reopen the database without the EXCLUSIVE keyword. The following statement opens the `stores_demo` database on the `training` database server in exclusive mode:

```
DATABASE stores_demo@training EXCLUSIVE
```

If another user has already opened the specified database, exclusive access is denied, an error is returned, and no database is opened.

Related Information

Related statements: "CLOSE DATABASE " on page 2-80, "CONNECT " on page 2-83, "CREATE DATABASE " on page 2-97, "DISCONNECT " on page 2-320, and "SET CONNECTION" on page 2-587

For discussions of how to use different data models to design and implement a database, see the *IBM Informix Database Design and Implementation Guide*.

For descriptions of the **sqlca** structure, see the *IBM Informix Guide to SQL: Tutorial* or the *IBM Informix ESQL/C Programmer's Manual*.

DEALLOCATE COLLECTION

Use the DEALLOCATE COLLECTION statement to release memory for a collection variable that was previously allocated with the ALLOCATE COLLECTION statement. Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL. Use this statement with ESQL/C.

Syntax

►►—DEALLOCATE COLLECTION—:*variable*—◀◀

Element	Description	Restrictions	Syntax
<i>variable</i>	Name that identifies a typed or untyped collection variable for which to deallocate memory	Must be the name of an Informix ESQL/C collection variable that has already been allocated	Name must conform to language-specific rules for names of variables

Usage

The DEALLOCATE COLLECTION statement frees all the memory that is associated with the Informix ESQL/C collection variable that *variable* identifies. You must explicitly release memory resources for a collection variable with DEALLOCATE COLLECTION. Otherwise, deallocation occurs automatically at the end of the program.

The DEALLOCATE COLLECTION statement releases resources for both typed and untyped collection variables.

Tip: The DEALLOCATE COLLECTION statement deallocates memory for an Informix ESQL/C collection variable only. To deallocate memory for an Informix ESQL/C row variable, use the DEALLOCATE ROW statement.

If you deallocate a nonexistent collection variable or a variable that is not an Informix ESQL/C collection variable, an error results. Once you deallocate a collection variable, you can use the ALLOCATE COLLECTION to reallocate resources and you can then reuse a collection variable.

This example shows how to deallocate resources with the DEALLOCATE COLLECTION statement for the untyped collection variable, **a_set**:

```
EXEC SQL BEGIN DECLARE SECTION;
      client collection a_set;
EXEC SQL END DECLARE SECTION;
. . .
EXEC SQL allocate collection :a_set;
. . .
EXEC SQL deallocate collection :a_set;
```

Related Information

Related example: refer to the collection variable example in "PUT " on page 2-479.

Related statements: "ALLOCATE COLLECTION " on page 2-1 and "DEALLOCATE ROW " on page 2-285

For a discussion of collection data types, see the *IBM Informix ESQL/C Programmer's Manual*.

DEALLOCATE DESCRIPTOR

Use the DEALLOCATE DESCRIPTOR statement to free a previously allocated, system-descriptor area. Use this statement with Informix ESQL/C.

Syntax

```
▶▶ DEALLOCATE DESCRIPTOR { 'descriptor' | descriptor_var } ▶▶
```

Element	Description	Restrictions	Syntax
<i>descriptor</i>	Name of a system-descriptor area	Use single quotation marks. System-descriptor area must already be allocated	"Quoted String" on page 4-189
<i>descriptor_var</i>	Host variable that contains the name of a system-descriptor area	System-descriptor area must already be allocated, and the variable must already have been declared	Name must conform to language-specific rules

Usage

The DEALLOCATE DESCRIPTOR statement frees all the memory that is associated with the system-descriptor area that *descriptor* or *descriptor_var* identifies. It also frees all the item descriptors (including memory for data items in the item descriptors).

You can reuse a descriptor or descriptor variable after it is deallocated. Otherwise, deallocation occurs automatically at the end of the program.

If you deallocate a nonexistent descriptor or descriptor variable, an error results.

You cannot use the DEALLOCATE DESCRIPTOR statement to deallocate an **sqllda** structure. You can use it only to free the memory that is allocated for a system-descriptor area.

The following examples show valid DEALLOCATE DESCRIPTOR statements. The first line uses an embedded-variable name, and the second line uses a quoted string to identify the allocated system-descriptor area.

```
EXEC SQL deallocate descriptor :descname;
```

```
EXEC SQL deallocate descriptor 'desc1';
```

Related Information

Related statements: "ALLOCATE DESCRIPTOR " on page 2-2, "DECLARE " on page 2-286, "DESCRIBE" on page 2-311, "EXECUTE " on page 2-348, "FETCH " on page 2-367, "GET DESCRIPTOR " on page 2-381, "OPEN " on page 2-459, "PREPARE " on page 2-468, "PUT " on page 2-479, and "SET DESCRIPTOR " on page 2-605

For more information on system-descriptor areas, refer to the *IBM Informix ESQL/C Programmer's Manual*.

DEALLOCATE ROW

Use the DEALLOCATE ROW statement to release memory for a ROW variable.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL. Use this statement with ESQL/C.

Syntax

►►—DEALLOCATE ROW—:*variable*—►►

Element	Description	Restrictions	Syntax
<i>variable</i>	Typed or untyped row variable	Must be declared and allocated	Language specific

Usage

DEALLOCATE ROW frees all the memory that is associated with the Informix ESQL/C typed or untyped **row** variable that *variable* identifies. If you do not explicitly release memory resources with DEALLOCATE ROW, deallocation occurs automatically at the end of the program. To deallocate memory for an Informix ESQL/C collection variable, use the DEALLOCATE COLLECTION statement.

After you deallocate a ROW variable, you can use the ALLOCATE ROW statement to reallocate resources, and you can then reuse a ROW variable. The following example shows how to deallocate resources for the ROW variable, **a_row**, using the DEALLOCATE ROW statement:

```
EXEC SQL BEGIN DECLARE SECTION; row (a int, b int) a_row;
EXEC SQL END DECLARE SECTION;
. . .
EXEC SQL allocate row :a_row;
. . .
EXEC SQL deallocate row :a_row;
```

Related Information

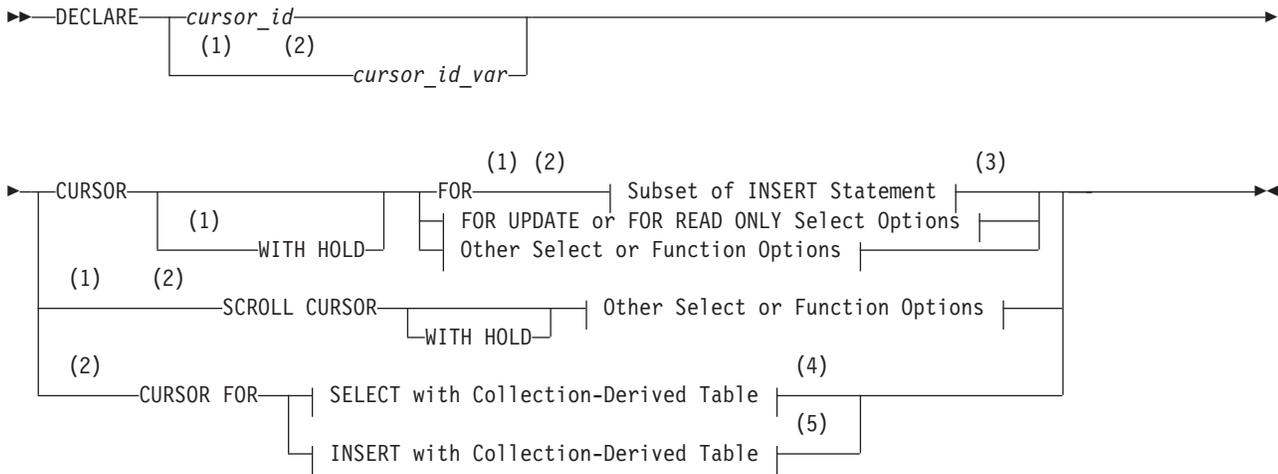
Related statements: "ALLOCATE ROW " on page 2-4 and "DEALLOCATE COLLECTION " on page 2-283

For a discussion of ROW data types, see the *IBM Informix Guide to SQL: Tutorial*. For complex data types, see the *IBM Informix ESQL/C Programmer's Manual*.

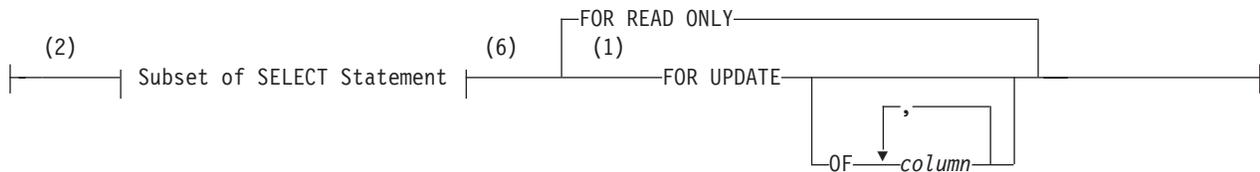
DECLARE

Use the DECLARE statement of dynamic SQL to declare a cursor and to associate it with an SQL statement that returns a set of rows to an Informix ESQL/C or SPL routine.

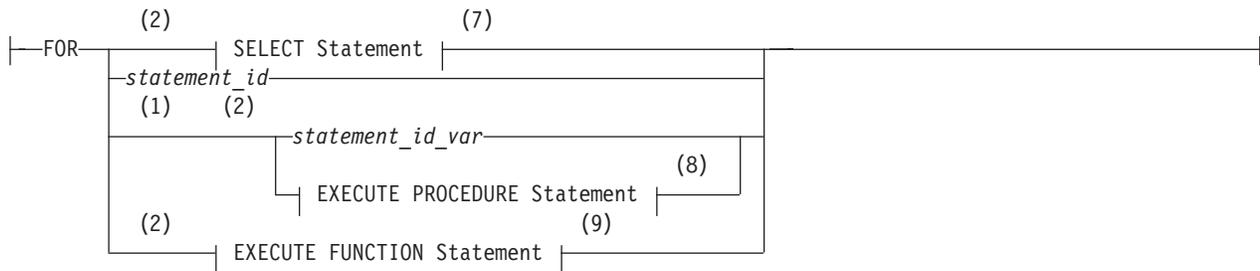
Syntax



FOR UPDATE or FOR READ ONLY Select Options:



Other Select or Function Options:



Notes:

- 1 Informix extension
- 2 ESQL/C only
- 3 See "Subset of INSERT Statement with a Sequential Cursor" on page 2-292
- 4 See "Select with a Collection-Derived Table" on page 2-298
- 5 See "Insert with a Collection-Derived Table" on page 2-299

- 6 See "Subset of SELECT Statement Associated with Cursors" on page 2-296
- 7 See "SELECT" on page 2-528
- 8 See "EXECUTE PROCEDURE " on page 2-364
- 9 See "EXECUTE FUNCTION " on page 2-356

Element	Description	Restrictions	Syntax
<i>column</i>	Column to update with cursor	Must exist, but need not be listed in Select list of the Projection clause	"Identifier " on page 5-21
<i>cursor_id</i>	Name declared here for cursor	Must be unique in the routine among names of cursors and prepared objects (and in SPL, among variables)	"Identifier " on page 5-21
<i>cursor_id_var</i>	Variable holding <i>cursor_id</i>	Must have a character data type	Language-specific
<i>statement_id</i>	Name of prepared statement	Must have been declared by a PREPARE statement	"Identifier " on page 5-21
<i>statement_id_var</i>	Variable holding <i>statement_id</i>	Must have a character data type	Language-specific

Usage

Except as noted, sections that follow describe how to use the DECLARE statement in Informix ESQL/C routines. For information about the more restricted syntax and semantics of the DECLARE statement in SPL routines, see "Declaring a Dynamic Cursor in an SPL Routine" on page 2-301.

A *cursor* is an identifier that you associate with an SQL statement that returns a group of rows. The DECLARE statement associates the cursor with one of the following database objects:

- With an SQL statement, such as SELECT, EXECUTE FUNCTION (or EXECUTE PROCEDURE), or INSERT.
Each of these SQL statements creates a different type of cursor. For more information, see "Overview of Cursor Types" on page 2-288.
- With the statement identifier (*statement id* or *statement id variable*) of a prepared statement
You can prepare one of the previous SQL statements and associate the prepared statement with a cursor. For more information, see "Associating a Cursor with a Prepared Statement" on page 2-297.
- Using Dynamic Server with a collection variable in an Informix ESQL/C program
The name of the collection variable appears in the FROM clause of a SELECT or the INTO clause of an INSERT. For more information, see "Associating a Cursor with a Prepared Statement" on page 2-297.

DECLARE assigns an identifier to the cursor, specifies its uses, and directs the Informix ESQL/C preprocessor to allocate storage for it. DECLARE must precede any other statement that refers to the cursor during program execution.

The maximum length of a DECLARE statement is 64 kilobytes. The cursors and prepared objects that can exist concurrently in a single program are limited by available memory. To avoid exceeding the limit, use the FREE statement to release the resources of prepared statements or cursors that are no longer needed.

An ESQL/C program can consist of one or more source-code files. By default, the scope of reference of a cursor is global to a program, so a cursor that was declared in one source file can be referenced from a statement in another file. If you want to limit the scope of each cursor name to the file where it was declared, you must preprocess each file with the **-local** command-line option.

Multiple cursors can be declared for the same prepared statement identifier. For example, the following Informix ESQL/C example does not return an error:

```
EXEC SQL prepare id1 from 'select * from customer';
EXEC SQL declare x cursor for id1;
EXEC SQL declare y scroll cursor for id1;
EXEC SQL declare z cursor with hold for id1;
```

If you include the **-ansi** compilation flag (or if **DBANSIWARN** is set), warnings are generated for statements that use dynamic cursor names, dynamic statement identifiers, or statements that use collection-derived tables. Some error checking is performed at runtime, such as these typical checks:

- Invalid use of sequential cursors as scroll cursors
- Use of undeclared cursors
- Invalid cursor names or statement names (empty)

Checks for multiple declarations of a cursor of the same name are performed at compile time only if the cursor or statement is specified as an identifier. The following example uses a host variable to store the cursor name:

```
EXEC SQL declare x cursor for select * from customer;
. . .
strcpy("x", s);
EXEC SQL declare :s cursor for select * from customer;
```

A cursor uses the collating order that was in effect when the cursor was declared, even if this is different from the collation of the session at runtime.

Overview of Cursor Types

Cursors are typically required for data manipulation language (DML) operations on more than one row of data (or on an Informix ESQL/C collection variable). You can declare the following types of cursors with the DECLARE statement:

- A *Select cursor* is a cursor associated with a SELECT statement.
- A *Function cursor* is a cursor associated with an EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement.
- An *Insert cursor* is a cursor associated with an INSERT statement.

Sections that follow describe each of these cursor types. Cursors can also have *sequential*, *scroll*, and *hold* characteristics (but an Insert cursor cannot be a scroll cursor). These characteristics determine the structure of the cursor; see “Cursor Characteristics” on page 2-293. In addition, a Select or Function cursor can specify *read-only* or *update* mode. For more information, see “Select Cursor or Function Cursor” on page 2-289.

Tip: Function cursors behave the same as Select cursors that are enabled as update cursors.

Besides associating a cursor directly with the text of an SQL statement, the FOR keyword of the DECLARE statement can be followed by the identifier of a prepared SQL statement, associating the cursor with the result set of an INSERT, SELECT, EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement that was

prepared dynamically. This feature enables you to associate different SQL statements with the same cursor at different times. In this case, the type of cursor depends on the prepared SQL statement that the *statement_id* or *statement_id* variable specifies when the cursor is opened. (For more information, see “Associating a Cursor with a Prepared Statement” on page 2-297.)

Select Cursor or Function Cursor

When an SQL statement returns more than one group of values to an Informix ESQL/C program, you must declare a cursor to save the multiple groups, or rows, of data and to access these rows one at a time. You must associate one of the following SQL statements with a cursor:

- If you associate a SELECT statement with a cursor, the cursor is called a *Select cursor*.

A Select cursor is a data structure that represents a specific location within the active set of rows that the SELECT statement retrieved.

- If you associate an EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement with a cursor, the cursor is called a *Function cursor*.

The Function cursor represents the columns or values that a user-defined function returns. Function cursors behave the same as Select cursors that are enabled as update cursors.

In Extended Parallel Server, to create a Function cursor, you must use the EXECUTE PROCEDURE statement. Extended Parallel Server does not support the EXECUTE FUNCTION statement.

In Dynamic Server, for compatibility with legacy applications, if an SPL function was created with the CREATE PROCEDURE statement, you can create a Function cursor with the EXECUTE PROCEDURE statement. With external functions, you must use the EXECUTE FUNCTION statement.

When you associate a SELECT or EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement with a cursor, the statement can include an INTO clause. However, if you prepare the SELECT or EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement, you must omit the INTO clause in the PREPARE statement and use the INTO clause of the FETCH statement to retrieve the values from the Collection cursor.

A Select or Function cursor can scan returned rows of data and to move data row by row into a set of receiving variables, as the following steps describe:

1. DECLARE

Use DECLARE to define a cursor and associate it with a statement.

2. OPEN

Use OPEN to open the cursor. The database server processes the query until it locates or constructs the first row of the active set.

3. FETCH

Use FETCH to retrieve successive rows of data from the cursor.

4. CLOSE

Use CLOSE to close the cursor when its active set is no longer needed.

5. FREE

Use FREE to release the resources that are allocated for the cursor.

Using the FOR READ ONLY Option

Use the FOR READ ONLY keywords to define a cursor as a read-only cursor. A cursor declared to be read-only cannot be used to update (or delete) any row that it fetches.

The need for the FOR READ ONLY keywords depends on whether your database is ANSI compliant or not ANSI compliant.

In a database that is not ANSI compliant, the cursor that the DECLARE statement defines is a read-only cursor by default, so you do not need to specify the FOR READ ONLY keywords if you want the cursor to be a read-only cursor. The only advantage of specifying the FOR READ ONLY keywords explicitly is for better program documentation.

In an ANSI-compliant database, the cursor associated with a SELECT statement through the DECLARE statement is an update cursor by default, provided that the SELECT statement conforms to all of the restrictions for update cursors listed in “Subset of SELECT Statement Associated with Cursors” on page 2-296. If you want a Select cursor to be read-only, you must use the FOR READ ONLY keywords when you declare the cursor.

The database server can use less stringent locking for a read-only cursor than for an update cursor.

The following example creates a read-only cursor:

```
EXEC SQL declare z_curs cursor for
      select * from customer_ansi
      for read only;
```

Using the FOR UPDATE Option

Use the FOR UPDATE option to declare an update cursor. You can use the update cursor to modify (update or delete) the current row.

In an ANSI-compliant database, you can use a Select cursor to update or delete data if the cursor was not declared with the FOR READ ONLY keywords and it follows the restrictions on update cursors that are described in “Subset of SELECT Statement Associated with Cursors” on page 2-296. You do not need to use the FOR UPDATE keywords when you declare the cursor.

The following example declares an update cursor:

```
EXEC SQL declare new_curs cursor for
      select * from customer_notansi
      for update;
```

In an update cursor, you can update or delete rows in the active set. After you create an update cursor, you can update or delete the currently selected row by using an UPDATE or DELETE statement with the WHERE CURRENT OF clause. The words CURRENT OF refer to the row that was most recently fetched; they take the place of the usual test expressions in the WHERE clause.

An update cursor lets you perform updates that are not possible with the UPDATE statement because the decision to update and the values of the new data items can be based on the original contents of the row. Your program can evaluate or manipulate the selected data before it decides whether to update. The UPDATE statement cannot interrogate the table that is being updated.

You can specify particular columns that can be updated. The columns need not appear in the Select list of the Projection clause.

Using FOR UPDATE with a List of Columns: When you declare an update cursor, you can limit the update to specific columns by including the OF keyword and a list of columns. You can modify only those named columns in subsequent UPDATE statements. The columns need not be in the select list of the SELECT clause.

The next example declares an update cursor and specifies that this cursor can update only the **fname** and **lname** columns in the **customer_notansi** table:

```
EXEC SQL declare name_curs cursor for
  select * from customer_notansi
  for update of fname, lname;
```

By default, unless declared as FOR READ ONLY, a Select cursor in a database that is ANSI compliant is an update cursor, so the FOR UPDATE keywords are optional. If you want an update cursor to be able to modify only some of the columns in a table, however, you must specify these columns in the FOR UPDATE OF *column* list.

The principal advantage to specifying columns is documentation and preventing programming errors. (The database server refuses to update any other columns.) An additional advantage is improved performance, when the SELECT statement meets the following criteria:

- The SELECT statement can be processed using an index.
- The columns that are listed are not part of the index that is used to process the SELECT statement.

If the columns that you intend to update are part of the index that is used to process the SELECT statement, the database server keeps a list of each updated row, to ensure that no row is updated twice. If the OF keyword specifies which columns can be updated, the database server determines whether or not to keep the list of updated rows. If the database server determines that the work of keeping the list is unnecessary, performance improves. If you do not use the OF *column* list, the database server always maintains a list of updated rows, although the list might be unnecessary.

The following example contains Informix ESQL/C code that uses an update cursor with a DELETE statement to delete the current row.

Whenever the row is deleted, the cursor remains between rows. After you delete data, you must use a FETCH statement to advance the cursor to the next row before you can refer to the cursor in a DELETE or UPDATE statement.

```
EXEC SQL declare q_curs cursor for
  select * from customer where lname matches :last_name for update;

EXEC SQL open q_curs;
for (;;)
{
  EXEC SQL fetch q_curs into :cust_rec;
  if (strncmp(SQLSTATE, "00", 2) != 0)
    break;

  /* Display customer values and prompt for answer */
  printf("\n%s %s", cust_rec.fname, cust_rec.lname);
  printf("\nDelete this customer? ");
  scanf("%s", answer);
}
```

```

    if (answer[0] == 'y')
        EXEC SQL delete from customer where current of q_curs;
    if (strncmp(SQLSTATE, "00", 2) != 0)
        break;
}
printf("\n");
EXEC SQL close q_curs;

```

Locking with an Update Cursor: The FOR UPDATE keywords notify the database server that updating is possible and cause it to use more stringent locking than with a Select cursor. You declare an update cursor to let the database server know that the program might update (or delete) any row that it fetches as part of the SELECT statement. The update cursor employs *promotable* locks for rows that the program fetches. Other programs can read the locked row, but no other program can place a promotable lock (also called a *write lock*). Before the program modifies the row, the row lock is promoted to an exclusive lock.

It is possible to declare an update cursor with the WITH HOLD keywords, but the only reason to do so is to break a long series of updates into smaller transactions. You must fetch and update a particular row in the same transaction.

If an operation involves fetching and updating a large number of rows, the lock table that the database server maintains can overflow. The usual way to prevent this overflow is to lock the entire table that is being updated. If this action is impossible, an alternative is to update through a hold cursor and to execute COMMIT WORK at frequent intervals. You must plan such an application carefully, however, because COMMIT WORK releases all locks, even those that are placed through a hold cursor.

Subset of INSERT Statement with a Sequential Cursor

As indicated in the diagram for "DECLARE " on page 2-286, to create an Insert cursor, you associate a sequential cursor with a restricted form of the INSERT statement. The INSERT statement must include a VALUES clause; it cannot contain an embedded SELECT statement.

The following example contains Informix ESQL/C code that declares an Insert cursor:

```

EXEC SQL declare ins_cur cursor for
    insert into stock values
        (:stock_no, :manu_code, :descr, :u_price, :unit, :u_desc);

```

The Insert cursor simply inserts rows of data; it cannot be used to fetch data. When an Insert cursor is opened, a buffer is created in memory. The buffer receives rows of data as the program executes PUT statements. The rows are written to disk only when the buffer is full. You can flush the buffer (that is, to write its contents into the database) when it is less than full, using the CLOSE, FLUSH, or COMMIT WORK statements. This topic is discussed further under the CLOSE, FLUSH, and PUT statements.

You must close an Insert cursor to insert any buffered rows into the database before the program ends. You can lose data if you do not close the cursor properly. For a complete description of INSERT syntax and usage, see "INSERT " on page 2-430.

Insert Cursor

When you associate an INSERT statement with a cursor, the cursor is called an *Insert cursor*. An Insert cursor is a data structure that represents the rows that the INSERT statement is to add to the database. The Insert cursor simply inserts rows of data; it cannot be used to fetch data. To create an Insert cursor, you associate a cursor with a restricted form of the INSERT statement. The INSERT statement must include a VALUES clause; it cannot contain an embedded SELECT statement.

Create an Insert cursor if you want to add multiple rows to the database in an INSERT operation. An Insert cursor allows bulk insert data to be buffered in memory and written to disk when the buffer is full, as these steps describe:

1. Use DECLARE to define an Insert cursor for the INSERT statement.
2. Open the cursor with the OPEN statement. The database server creates the insert buffer in memory and positions the cursor at the first row of the insert buffer.
3. Copy successive rows of data into the insert buffer with the PUT statement.
4. The database server writes the rows to disk only when the buffer is full. You can use the CLOSE, FLUSH, or COMMIT WORK statement to flush the buffer when it is less than full. This topic is discussed further under the PUT and CLOSE statements.
5. Close the cursor with the CLOSE statement when the insert cursor is no longer needed. You must close an Insert cursor to insert any buffered rows into the database before the program ends. You can lose data if you do not close the cursor properly.
6. Free the cursor with the FREE statement. The FREE statement releases the resources that are allocated for an Insert cursor.

Using an Insert cursor is more efficient than embedding the INSERT statement directly. This process reduces communication between the program and the database server and also increases the speed of the insertions.

Insert cursors also have the sequential cursor characteristic. To create an Insert cursor, you associate a sequential cursor with a restricted form of the INSERT statement. (For more information, see “Insert Cursor.”) The following example contains IBM Informix ESQL/C code that declares a sequential Insert cursor:

```
EXEC SQL declare ins_cur cursor for
      insert into stock values
      (:stock_no, :manu_code, :descr, :u_price, :unit, :u_desc);
```

Cursor Characteristics

You can declare a cursor as a *sequential* cursor (the default), a *scroll* cursor (by using the SCROLL keyword), or a *hold* cursor (by using the WITH HOLD keywords). The SCROLL and WITH HOLD keywords are not mutually exclusive. Sections that follow explain these structural characteristics.

A Select or Function cursor can be either a sequential or a scroll cursor. An Insert cursor can only be a sequential cursor. In ESQL/C routines, Select, Function, and Insert cursors can optionally be hold cursors. (In SPL routines, all cursors are sequential cursors, but only Select cursors can be hold cursors.)

Creating a Sequential Cursor by Default

If you use only the `CURSOR` keyword, you create a sequential cursor, which can fetch only the next row in sequence from the active set. The sequential cursor can read through the active set only once each time it is opened.

If you are using a sequential cursor for a Select cursor, on each execution of the `FETCH` statement, the database server returns the contents of the current row and locates the next row in the active set.

The following example creates a read-only sequential cursor in a database that is not ANSI compliant and an update sequential cursor in an ANSI-compliant database:

```
EXEC SQL declare s_cur cursor for
      select fname, lname into :st_fname, :st_lname
      from orders where customer_num = 114;
```

Insert cursors also have the sequential cursor characteristic. To create a Insert cursor, you associate a sequential cursor with a restricted form of the `INSERT` statement. (For more information, see “Insert Cursor” on page 2-293.) The following example declares an Insert cursor:

```
EXEC SQL declare ins_cur cursor for
      insert into stock values
      (:stock_no, :manu_code, :descr, :u_price, :unit, :u_desc);
```

Using the `SCROLL` Keyword to Create a Scroll Cursor

Use the `SCROLL` keyword to create a scroll cursor, which can fetch rows of the active set in any sequence.

The database server retains the active set of the cursor as a temporary table until the cursor is closed. You can fetch the first, last, or any intermediate rows of the active set as well as fetch rows repeatedly without having to close and reopen the cursor. (See `FETCH`.)

On a multiuser system, the rows in the tables from which the active-set rows were derived might change after the cursor is opened and a copy is made in the temporary table. If you use a scroll cursor within a transaction, you can prevent copied rows from changing either by setting the isolation level to Repeatable Read or by locking the entire table in share mode during the transaction. (See `SET ISOLATION` and `LOCK TABLE`.)

The following example creates a scroll cursor for a `SELECT` statement:

```
DECLARE sc_cur SCROLL CURSOR FOR SELECT * FROM orders;
```

You can create scroll cursors as Select and Function cursors but *not* as Insert cursors. Scroll cursors cannot be declared as `FOR UPDATE`.

Using the `WITH HOLD` Keywords to Create a Hold Cursor

Use the `WITH HOLD` keywords to create a hold cursor. A hold cursor allows uninterrupted access to a set of rows across multiple transactions. Ordinarily, all cursors close at the end of a transaction. A hold cursor does not close; it remains open after a transaction ends.

A hold cursor can be either a sequential cursor or (in `ESQL/C`) a scroll cursor.

The WITH HOLD keywords are valid in SPL routines only for Select cursors. For the syntax of the DECLARE statement in SPL routines, see “Declaring a Dynamic Cursor in an SPL Routine” on page 2-301.

In ESQL/C, you can use the WITH HOLD keywords to declare Select and Function cursors (with the sequential attribute or the scroll attribute) and also to declare Insert cursors. These keywords follow the CURSOR keyword in the DECLARE statement. The following example creates a sequential hold cursor for a SELECT statement:

```
DECLARE hld_cur CURSOR WITH HOLD FOR
  SELECT customer_num, lname, city FROM customer;
```

You can use a select hold cursor as the following Informix ESQL/C code example shows. This code fragment uses a hold cursor as a *master* cursor to scan one set of records and a sequential cursor as a *detail* cursor to point to records that are located in a different table. The records that the master cursor scans are the basis for updating the records to which the detail cursor points. The COMMIT WORK statement at the end of each iteration of the first WHILE loop leaves the hold cursor **c_master** open but closes the sequential cursor **c_detail** and releases all locks. This technique minimizes the resources that the database server must devote to locks and unfinished transactions, and it gives other users immediate access to updated rows.

```
EXEC SQL BEGIN DECLARE SECTION;
  int p_custnum, int save_status; long p_orddate;
EXEC SQL END DECLARE SECTION;

EXEC SQL prepare st_1 from
  'select order_date from orders where customer_num = ? for update';
EXEC SQL declare c_detail cursor for st_1;
EXEC SQL declare c_master cursor with hold for
  select customer_num from customer where city = 'Pittsburgh';

EXEC SQL open c_master;
if(SQLCODE==0) /* the open worked */
  EXEC SQL fetch c_master into :p_custnum; /* discover first customer */
while(SQLCODE==0) /* while no errors and not end of pittsburgh customers */
  {
  EXEC SQL begin work; /* start transaction for customer p_custnum */
  EXEC SQL open c_detail using :p_custnum;
  if(SQLCODE==0) /* detail open succeeded */
    EXEC SQL fetch c_detail into :p_orddate; /* get first order */
  while(SQLCODE==0) /* while no errors and not end of orders */
    {
    EXEC SQL update orders set order_date = '08/15/94'
      where current of c_detail;
    if(status==0) /* update was ok */
      EXEC SQL fetch c_detail into :p_orddate; /* next order */
    }
  if(SQLCODE==SQLNOTFOUND) /* correctly updated all found orders */
    EXEC SQL commit work; /* make updates permanent, set status */
  else /* some failure in an update */
    {
    save_status = SQLCODE; /* save error for loop control */
    EXEC SQL rollback work;
    SQLCODE = save_status; /* force loop to end */
    }
  if(SQLCODE==0) /* all updates, and the commit, worked ok */
    EXEC SQL fetch c_master into :p_custnum; /* next customer? */
  }
EXEC SQL close c_master;
```

Use either the CLOSE statement to close the hold cursor explicitly or the CLOSE DATABASE or DISCONNECT statements to close it implicitly. The CLOSE DATABASE statement closes all cursors.

Releases earlier than Version 9.40 of Dynamic Server do not support the PDQPRIORITY feature with cursors that were declared WITH HOLD.

Using an Insert Cursor with Hold: If you associate a hold cursor with an INSERT statement, you can use transactions to break a long series of PUT statements into smaller sets of PUT statements. Instead of waiting for the PUT statements to fill the buffer and cause an automatic write to the database, you can execute a COMMIT WORK statement to flush the row buffer. With a hold cursor, COMMIT WORK commits the inserted rows but leaves the cursor open for further inserts. This method can be desirable when you are inserting a large number of rows, because pending uncommitted work consumes database server resources.

Subset of SELECT Statement Associated with Cursors

As indicated in the syntax diagram for “DECLARE ” on page 2-286, not all SELECT statements can be associated with a read-only or update cursor.

If the DECLARE statement includes one of these options, you must observe certain restrictions on the SELECT statement that is included in the DECLARE statement (either directly or as a prepared statement).

If the DECLARE statement includes the FOR READ ONLY option, the SELECT statement cannot have a FOR READ ONLY or FOR UPDATE option. (For a description of SELECT syntax and usage, see “SELECT” on page 2-528.)

If the DECLARE statement includes the FOR UPDATE option, the SELECT statement must conform to the following restrictions:

- The statement can select data from only one table.
- The statement cannot include any aggregate functions.
- The statement cannot include any of the following clauses or keywords: DISTINCT, FOR READ ONLY, FOR UPDATE, GROUP BY, INTO TEMP, ORDER BY, UNION, or UNIQUE.
- In Extended Parallel Server, the statement cannot include the INTO EXTERNAL and INTO SCRATCH clauses.

Examples of Cursors in Non-ANSI Compliant Databases

In a database that is not ANSI compliant, a cursor associated with a SELECT statement is a read-only cursor by default. The following example declares a read-only cursor in a non-ANSI compliant database:

```
EXEC SQL declare cust_curs cursor for
      select * from customer_notansi;
```

If you want to make it clear in the program code that this cursor is a read-only cursor, specify the FOR READ ONLY option as the following example shows:

```
EXEC SQL declare cust_curs cursor for
      select * from customer_notansi for read only;
```

If you want this cursor to be an update cursor, specify the FOR UPDATE option in your DECLARE statement. This example declares an update cursor:

```
EXEC SQL declare new_curs cursor for
      select * from customer_notansi for update;
```

If you want an update cursor to be able to modify only some columns in a table, you must specify those columns in the FOR UPDATE clause. The following example declares an update cursor that can update only the **fname** and **lname** columns in the **customer_notansi** table:

```
EXEC SQL declare name_curs cursor for
    select * from customer_notansi for update of fname, lname;
```

Examples of Cursors in ANSI-Compliant Databases

In an ANSI-compliant database, a cursor associated with a SELECT statement is an update cursor by default.

The following example declares an update cursor in an ANSI-compliant database:

```
EXEC SQL declare x_curs cursor for select * from customer_ansi;
```

To make it clear in the program documentation that this cursor is an update cursor, you can specify the FOR UPDATE option as in this example:

```
EXEC SQL declare x_curs cursor for
    select * from customer_ansi for update;
```

If you want an update cursor to be able to modify only some of the columns in a table, you must specify these columns in the FOR UPDATE option. The following example declares an update cursor and specifies that this cursor can update only the **fname** and **lname** columns in the **customer_ansi** table:

```
EXEC SQL declare y_curs cursor for
    select * from customer_ansi for update of fname, lname;
```

If you want a cursor to be a read-only cursor, you must override the default behavior of the DECLARE statement by specifying the FOR READ ONLY option in your DECLARE statement. The following example declares a read-only cursor:

```
EXEC SQL declare z_curs cursor for
    select * from customer_ansi for read only;
```

Associating a Cursor with a Prepared Statement

The PREPARE statement lets you assemble the text of an SQL statement at runtime and pass the statement text to the database server for execution. If you anticipate that a dynamically prepared SELECT, EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement that returns values could produce more than one row of data, the prepared statement must be associated with a cursor. (See PREPARE.)

The result of a PREPARE statement is a statement identifier (*statement id* or *id variable*), which is a data structure that represents the prepared statement text. To declare a cursor for the statement text, associate a cursor with the statement identifier.

You can associate a sequential cursor with any prepared SELECT or EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement. You cannot associate a scroll cursor with a prepared INSERT statement or with a SELECT statement that was prepared to include a FOR UPDATE clause.

After a cursor is opened, used, and closed, a different statement can be prepared under the same statement identifier. In this way, it is possible to use a single cursor with different statements at different times. The cursor must be redeclared before you use it again.

The following example contains Informix ESQL/C code that prepares a SELECT statement and declares a sequential cursor for the prepared statement text. The statement identifier `st_1` is first prepared from a SELECT statement that returns values; then the cursor `c_detail` is declared for `st_1`.

```
EXEC SQL prepare st_1 from
    'select order_date
      from orders where customer_num = ?';
EXEC SQL declare c_detail cursor for st_1;
```

If you want to use a prepared SELECT statement to modify data, add a FOR UPDATE clause to the statement text that you want to prepare, as the following Informix ESQL/C example shows:

```
EXEC SQL prepare sel_1 from
    'select * from customer for update';
EXEC SQL declare sel_curs cursor for sel_1;
```

DDL operations that change the schema of a table can invalidate a cursor whose associated prepared statement or associated routine references the modified table, unless the prepared objects are recompiled, or unless the routine is reoptimized. For more information, see the section “DDL Operations on Tables Referenced by Cursors” on page 2-466.

The DECLARE statement allows you to declare a cursor for an Informix ESQL/C collection variable. Such a cursor is called a *Collection cursor*. You use a collection variable to access the elements of a collection (SET, MULTISSET, LIST) column. Use a cursor when you want to access one or more elements in a collection variable.

The Collection-Derived Table segment identifies the collection variable for which to declare the cursor. For more information, see “Collection-Derived Table” on page 5-4.

Select with a Collection-Derived Table

The diagram for “DECLARE ” on page 2-286 refers to this section.

To declare a Select cursor for a collection variable, include the Collection-Derived Table segment with the SELECT statement that you associate with the Collection cursor. A Select cursor allows you to select one or more elements *from* the collection variable. (For a description of SELECT syntax and usage, see “SELECT” on page 2-528.)

When you declare a Select cursor for a collection variable, the DECLARE statement has the following restrictions:

- It cannot include the FOR READ ONLY keywords as cursor mode.
The Select cursor is an update cursor.
- It cannot include the SCROLL or WITH HOLD keywords.
The Select cursor must be a sequential cursor.

In addition, the SELECT statement that you associate with the collection cursor has the following restrictions:

- It cannot include the following clauses or options: WHERE, GROUP BY, ORDER BY, HAVING, INTO TEMP, and WITH REOPTIMIZATION.
- It cannot contain expressions in the projection list.
- If the collection contains elements of opaque, distinct, built-in, or other collection data types, the projection list must be an asterisk (*).

- Column names in the projection list must be simple column names, without qualifiers.

These columns cannot use the following syntax:

```
database@server:table.column --INVALID SYNTAX
```

- It *must* specify the name of the collection variable in the FROM clause.
You cannot specify an input parameter (the question-mark (?) symbol) for the collection variable. Likewise, you cannot use the virtual table format of the Collection-Derived Table segment.

Using a Select Cursor with a Collection Variable:

A Collection cursor that includes a SELECT statement with the Collection- Derived Table clause provides access to the elements in a collection variable.

To select more than one element

1. Create a client collection variable in your Informix ESQL/C program.
2. Declare the Collection cursor for the SELECT statement with the DECLARE statement. To modify elements of the collection variable, declare the Select cursor as an update cursor with the FOR UPDATE keywords. You can then use the WHERE CURRENT OF clause of the DELETE and UPDATE statements to delete or update elements of the collection.
3. Open this cursor with the OPEN statement.
4. Fetch the elements from the Collection cursor with the FETCH statement and the INTO clause.
5. If necessary, perform any updates or deletes on the fetched data and save the modified collection variable in the collection column. Once the collection variable contains the correct elements, use the UPDATE or INSERT statement to save the contents of the collection variable in the actual collection column (SET, MULTISSET, or LIST).
6. Close the Collection cursor with the CLOSE statement.

This DECLARE statement declares a Select cursor for a collection variable:

```
EXEC SQL BEGIN DECLARE SECTION;
      client collection set(integer not null) a_set;
EXEC SQL END DECLARE SECTION;
...
EXEC SQL declare set_curs cursor for select * from table(:a_set);
```

For an extended code example that uses a Collection cursor for a SELECT statement, see “ Fetching from a Collection Cursor (IDS)” on page 2-374.

Insert with a Collection-Derived Table

To declare an Insert cursor for a collection variable, include the Collection- Derived Table segment in the INSERT statement associated with the Collection cursor. An Insert cursor can insert one or more elements in the collection. For a description of INSERT syntax and usage, see “INSERT ” on page 2-430.

The Insert cursor must be a sequential cursor. That is, the DECLARE statement cannot specify the SCROLL keyword.

When you declare an Insert cursor for a collection variable, the Collection- Derived Table clause of the INSERT statement *must* contain the name of the collection variable. You cannot specify an input parameter (the question-mark (?) symbol)

for the collection variable. However, you can use an input parameter in the VALUES clause of the INSERT statement. This parameter indicates that the collection element is to be provided later by the FROM clause of the PUT statement.

A Collection cursor that includes an INSERT statement with the Collection-Derived Table clause allows you to insert more than one element into a collection variable.

To insert more than one element

1. Create a client collection variable in your Informix ESQL/C program.
2. Declare the Collection cursor for the INSERT statement with the DECLARE statement.
3. Open the cursor with the OPEN statement.
4. Put the elements into the Collection cursor with the PUT statement and the FROM clause.
5. Once the collection variable contains all the elements, use the UPDATE statement or the INSERT statement on a table name to save the contents of the collection variable in a collection column (SET, MULTISSET, or LIST).
6. Close the Collection cursor with the CLOSE statement.

This example declares an Insert cursor for the `a_set` collection variable:

```
EXEC SQL BEGIN DECLARE SECTION;
    client collection multiset(smallint not null) a_mset;
    int an_element;
EXEC SQL END DECLARE SECTION;
...
EXEC SQL declare mset_curs cursor for
    insert into table(:a_mset) values (?);
EXEC SQL open mset_curs;
while (1)
{
    ...
    EXEC SQL put mset_curs from :an_element;
    ...
}
```

To insert the elements into the collection variable, use the PUT statement with the FROM clause. For a code example that uses a Collection cursor for an INSERT statement, see “Inserting into a Collection Cursor (IDS)” on page 2-483.

Using Cursors with Transactions

To roll back a modification, you must perform the modification within a transaction. A transaction in a database that is not ANSI compliant begins only when the BEGIN WORK statement is executed.

In an ANSI-compliant database, transactions are always in effect.

The database server enforces these guidelines for select and update cursors to ensure that modifications can be committed or rolled back properly:

- Open an insert or update cursor within a transaction.
- Include PUT and FLUSH statements within one transaction.
- Modify data (update, insert, or delete) within one transaction.

The database server lets you open and close a hold cursor for an update outside a transaction; however, you should fetch all the rows that pertain to a given modification and then perform the modification all within a single transaction. You cannot open and close a hold cursor or an update cursor outside a transaction.

The following example uses an update cursor within a transaction:

```
EXEC SQL declare q_curs cursor for
    select customer_num, fname, lname from customer
    where lname matches :last_name for update;
EXEC SQL open q_curs;
EXEC SQL begin work;
EXEC SQL fetch q_curs into :cust_rec; /* fetch after begin */
EXEC SQL update customer set lname = 'Smith'
    where current of q_curs;
/* no error */
EXEC SQL commit work;
```

When you update a row within a transaction, the row remains locked until the cursor is closed or the transaction is committed or rolled back. If you update a row when no transaction is in effect, the row lock is released when the modified row is written to disk. If you update or delete a row outside a transaction, you cannot roll back the operation.

In a database that uses transactions, you cannot open an Insert cursor outside a transaction unless it was also declared with the WITH HOLD keywords.

Declaring a Dynamic Cursor in an SPL Routine

Use the DECLARE statement in an SPL routine to declare the name of a dynamic cursor and to associate that cursor with the statement identifier of a prepared object that the PREPARE statement has declared in the same SPL routine.

Syntax

The syntax of the DECLARE statement in SPL routines is a subset of the syntax that DECLARE supports in Informix ESQ/C routines.

```

▶▶—DECLARE— cursor_id ————— (1) ————— WITH HOLD ————— FOR— statement_id —————▶▶

```

Notes:

- 1 Informix extension

Element	Description	Restrictions	Syntax
<i>cursor_id</i>	Name declared here for a new dynamic cursor	Must be unique among names of cursors, variables, and prepared objects in the SPL routine	"Identifier " on page 5-21
<i>statement_id</i>	Identifier of a single prepared SQL statement	Must have been declared in a PREPARE statement of the same SPL routine	"Identifier " on page 5-21

Usage

In UDRs written in the SPL language, the *statement_id* associated with the cursor must have been declared earlier in the same UDR by a PREPARE statement from the text of a single SQL statement of one of these statement types:

- EXECUTE FUNCTION
- EXECUTE PROCEDURE
- SELECT.

This prepared statement text that *statement_id* specifies can include question mark (?) symbols as placeholders for values that the user supplies at runtime, but these placeholders in the PREPARE statement can represent only data values, not SQL identifiers.

Dynamic cursors that the DECLARE statement can define in SPL routines resemble ESQL/C Select cursors or Function cursors in their functionality, but with these restrictions:

- Cursors that DECLARE defines in an SPL routine can be Select cursors or Function cursors, but they cannot be Insert cursors nor Collection cursors.
- The identifier of the cursor or of the prepared statement cannot be specified as an SPL variable, because in SPL, the identifiers of variables, cursors, and prepared objects all share the same namespace.
- By default, dynamic cursors of SPL are sequential. They cannot be scroll cursors.
- The semantics of dynamic cursors that you create with the WITH HOLD keywords are the same as for hold cursors that the FOREACH statement declares.
- The WITH HOLD keywords are valid in SPL routines only for Select cursors. If *statement_id* references the prepared text of an EXECUTE FUNCTION or EXECUTE PROCEDURE statement, the DECLARE statement fails with error -26056.
- The FOR UPDATE and FOR READ ONLY keywords that ESQL/C supports in DECLARE statements are not supported in SPL routines. Use the FOREACH statement of SPL to declare direct cursors that can emulate the functionality of ESQL/C update cursors. (But queries associated with direct cursors are defined when the UDR is compiled, rather than at runtime.)
- The DECLARE statement in SPL routines does not support SELECT operations on collection-derived tables.
- Syntax errors in DECLARE statements of SPL routines are reported at runtime, unlike syntax errors of ESQL/C, which are reported when the routine is compiled.

The names of the dynamic cursors that the DECLARE statement associates with a prepared statement in SPL routines can be referenced by the OPEN, CLOSE, FETCH, and FREE statements of dynamic SQL in the same SPL routine.

In the following program fragment, a cursor called **equi_noctis** is declared, opened, closed, and freed.

```
CREATE FUNCTION lente
  DEFINE first, last VARCHAR(30);
  . . .
  DATABASE stores_demo;
  LET first = "select * from state";
  LET lsst = "where code < ?";
  PREPARE stmt_1 FROM first || last;
  DECLARE cursor_1 FOR stmt_1;
```

```

OPEN cursor_1
. . .
CLOSE cursor_1;
FREE cursor_1;
FREE stmt_1;
. . .
END FUNCTION;

```

Related Information

Related statements: "CLOSE " on page 2-77, "DELETE ," "EXECUTE PROCEDURE " on page 2-364, "FETCH " on page 2-367, "FLUSH " on page 2-377, "FREE " on page 2-379, "INSERT " on page 2-430, "OPEN " on page 2-459, "PREPARE " on page 2-468, "PUT " on page 2-479, "SELECT" on page 2-528, " SET AUTOFREE " on page 2-582 and "UPDATE" on page 2-689.

For a discussion of direct cursors in SPL routines that for some tasks are an alternative to dynamic cursors, see "FOREACH " on page 3-22.

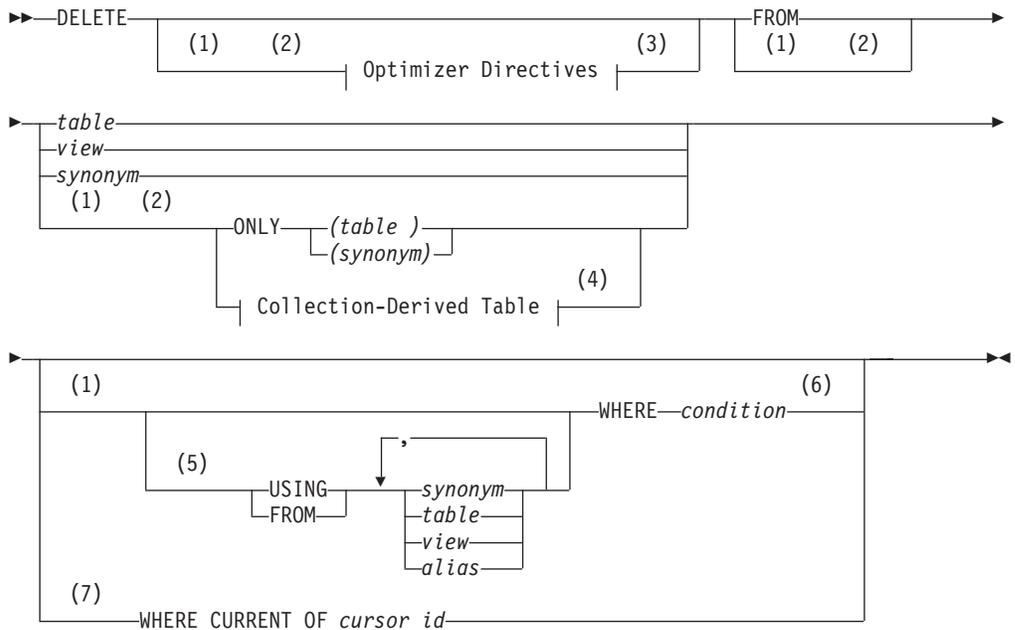
For discussions of cursors and data modification, see the *IBM Informix Guide to SQL: Tutorial*.

For more advanced issues related to cursors or to their use with collection variables, see the *IBM Informix ESQL/C Programmer's Manual*.

DELETE

Use the DELETE statement to delete one or more rows from a table, or to delete one or more elements in an SPL or Informix ESQL/C collection variable.

Syntax



Notes:

- 1 Informix extension

- 2 Dynamic Server only
- 3 See "Optimizer Directives" on page 5-35
- 4 See "Collection-Derived Table" on page 5-4
- 5 Extended Parallel Server only
- 6 See "Using the WHERE Keyword to Specify a Condition" on page 2-306
- 7 ESQL/C and Stored Procedure Language only

Element	Description	Restrictions	Syntax
<i>alias</i>	Temporary name that you declare here for a table	You cannot declare an <i>alias</i> for an indexed table	"Identifier" on page 5-21
<i>condition</i>	Logical criteria that deleted rows must satisfy	Cannot be a UDR nor a correlated subquery	"Condition" on page 4-5
<i>cursor_id</i>	Previously declared cursor	Must have been declared FOR UPDATE	"Identifier" on page 5-21
<i>synonym, table, view</i>	Table, view, or synonym with rows to be deleted	The <i>table</i> or <i>view</i> (or <i>synonym</i> and the table or view to which it points) must exist	"Database Object Name" on page 5-16

Usage

To execute the DELETE statement, you must hold the DBA access privilege on the database, or the Delete access privilege on the table.

In a database with explicit transaction logging, any DELETE statement that you execute outside a transaction is treated as a single transaction.

If you specify a *view* name, the view must be updatable. For an explanation of an updatable view, see "Updating Through Views" on page 2-278.

The database server locks each row affected by a DELETE statement within a transaction for the duration of the transaction. The type of lock that the database server uses is determined by the lock mode of the table, as set by a CREATE TABLE or ALTER TABLE statement, as follows:

- If the lock mode is ROW, the database server acquires one lock for each row affected by the delete.
- In Extended Parallel Server, if the lock mode is PAGE, the database server acquires one lock for each page affected by the delete.

If the number of rows affected is very large and the lock mode is ROW, you might exceed the limits your operating system places on the maximum number of simultaneous locks. If this occurs, you can either reduce the scope of the DELETE statement or lock the table in exclusive mode before you execute the statement.

If you use DELETE without a WHERE clause (to specify either a condition or the active set of the cursor), all rows in the table are deleted. It is typically more efficient, however, to use the TRUNCATE statement, rather than the DELETE statement, to remove all rows from a table.

In DB-Access, if you omit the WHERE clause while working at the SQL menu, DB-Access prompts you to verify that you want to delete all rows from a table. You do not receive a prompt if you run execute DELETE within a command file.

In an ANSI-compliant database, data manipulation language (DML) statements are always in a transaction. You cannot execute a DELETE statement outside a transaction.

On Dynamic Server, the FROM keyword that immediately follows DELETE can be omitted if the DELIMIT environment variable has been set.

Using the ONLY Keyword (IDS)

If you use DELETE to remove rows of a supertable, rows from both the supertable and its subtables can be deleted. To delete rows from the supertable only, specify the ONLY keyword before the table name.

```
DELETE FROM ONLY(super_tab)
  WHERE name = "johnson";
```

Warning: If you use the DELETE statement on a supertable and omit the ONLY keyword and WHERE clause, all rows of the supertable and its subtables are deleted.

You cannot specify the ONLY keyword if you plan to use the WHERE CURRENT OF clause to delete the current row of the active set of a cursor.

Considerations When Tables Have Cascading Deletes

When you use the ON DELETE CASCADE option of the REFERENCES clause of either the CREATE TABLE or ALTER TABLE statement, you specify that you want deletes to cascade from one table to another. For example, in the **stores_demo** database, the **stock** table contains the column **stock_num** as a primary key. The **catalog** and **items** tables each contain the column **stock_num** as foreign keys with the ON DELETE CASCADE option specified. When a delete is performed from the **stock** table, rows are also deleted in the **catalog** and **items** tables, which are referenced through the foreign keys.

To have DELETE actions cascade to a table that has a referential constraint on a parent table, you need the Delete privilege only on the parent table that you reference in the DELETE statement.

If a DELETE operation with no WHERE clause is performed on a table that one or more child tables reference with cascading deletes, Dynamic Server deletes all rows from that table and from any affected child tables. (This resembles the effect of the TRUNCATE statement, but Dynamic Server does not support TRUNCATE operations on any table that has a child table referencing it.)

For an example of how to create a referential constraint that uses cascading deletes, see "Using the ON DELETE CASCADE Option" on page 2-201.

Restrictions on DELETE When Tables Have Cascading Deletes

You cannot use a child table in a correlated subquery to delete a row from a parent table. If two child tables reference the same parent table, and one child specifies cascading deletes but the other child does not, then if you attempt to delete a row that applies to both child tables from the parent table, the delete fails, and no rows are deleted from the parent or child tables.

Locking and Logging Implications of Cascading Deletes

During deletes, the database server places locks on all qualifying rows of the referenced and referencing tables.

Dynamic Server requires transaction logging for cascading deletes. If logging is turned off in a database that is not ANSI-compliant, even temporarily, deletes do not cascade, because you cannot roll back any actions. For example, if a parent row is deleted, but the system fails before the child rows are deleted, the database will have dangling child records, in violation of referential integrity. After logging is turned back on, however, subsequent deletes cascade.

Using the WHERE Keyword to Specify a Condition

Use the WHERE *condition* clause to specify which rows you want to delete from the table. The *condition* after the WHERE keyword is equivalent to the *condition* in the SELECT or UPDATE statement. For example, the next statement deletes all the rows of the **items** table where the order number is less than 1034:

```
DELETE FROM items WHERE order_num < 1034;
```

In DB–Access, if you include a WHERE clause that selects all rows in the table, DB–Access gives no prompt and deletes all rows.

In Dynamic Server, if you are deleting from a supertable in a table hierarchy, a subquery in the WHERE clause cannot reference a subtable.

When deleting from a subtable, a subquery in the WHERE clause can reference the supertable only in SELECT ... FROM ONLY (*supertable*)... syntax.

Subqueries in the WHERE Clause of DELETE

The FROM clause of a subquery in the WHERE clause of the DELETE statement can specify as a data source the same table or view that the FROM clause of the DELETE statement specifies. DELETE operations with subqueries that reference the same table object are supported only if all of the following conditions are true:

- The subquery either returns a single row, or else has no correlated column references.
- The subquery is in the DELETE statement WHERE clause, using Condition with Subquery syntax.
- Any SPL routines within the subquery cannot reference the table that is being modified.

Unless all of these conditions are satisfied, DELETE statements that include subqueries that reference the same table or view that the DELETE statement modifies return error -360.

The following example deletes from the **orders** table a subset of rows whose **paid_date** column value satisfies the condition in the WHERE clause. The WHERE clause specifies which rows to delete by applying the IN operator to the rows returned by a subquery that selects only the rows of the **orders** table where the **paid_date** value is earlier than the current date:

```
DELETE FROM orders WHERE paid_date IN  
(SELECT paid_date FROM orders WHERE paid_date < CURRENT );
```

This subquery includes only uncorrelated column references, because its only referenced column is in a table specified in its FROM clause. The requirements listed above are in effect, because the data source of the subquery is the same orders table that the FROM clause of the outer UPDATE statement specifies. The previous example illustrates Dynamic Server support for uncorrelated subqueries in the WHERE clause of the DELETE statement. rather than how to write short SQL statements. The next example achieves the same result with simpler syntax:

```
DELETE orders WHERE paid_date < CURRENT;
```

The following example deletes from the **stock** table the row (or rows) with the largest **unit_price** value. The WHERE clause identifies which **unit_price** value is the largest by applying the equality operator to the result of a subquery that calls the built-in **MAX** aggregate function for the **unit_price** column values:

```
DELETE FROM stock WHERE unit_price =  
    (SELECT MAX(unit_price) FROM stock );
```

If an enabled Select trigger is defined on a table that is the data source of a subquery in the WHERE clause of a DELETE statement that modifies the same table, executing that subquery within the DELETE statement does not activate the trigger.

A subquery in the DELETE statement can include the UNION or UNION ALL operators.

If the table that the outer DELETE statement modifies a typed table within a table hierarchy, Dynamic Server supports all of the following operations that use valid subqueries in the WHERE clause of DELETE:

- DELETE from parent table with subquery (SELECT from parent table)
- DELETE from parent table with subquery (SELECT from child table)
- DELETE from child table with subquery (SELECT from parent table)
- DELETE from child table with subquery (SELECT from child table).

See the Condition with Subquery topic for more information about the syntax of subqueries to return multiple rows as predicates in the WHERE clause of the DELETE statement.

Using the WHERE CURRENT OF Keywords (ESQL/C, SPL)

The WHERE CURRENT OF clause deletes the current row of the active set of a cursor. When you include this clause, the DELETE statement removes the row of the active set at the current position of the cursor. After the deletion, no current row exists; you cannot use the cursor to delete or update a row until you reposition the cursor with a FETCH statement (in ESQL/C routines) or with a FOREACH statement (in SPL routines).

You access the current row of the active set of a cursor with an update cursor. Before you can use the WHERE CURRENT OF clause, you must first create an update cursor by using the FOREACH statement (in SPL) or the DECLARE statement with the FOR UPDATE clause (in Informix ESQL/C). The *cursor_id* that follows the OF keyword cannot be declared, however, by the DECLARE statement in an SPL routine

Unless they are declared with the FOR READ ONLY keywords, all Select cursors are potentially update cursors in an ANSI-compliant database. You can use the WHERE CURRENT OF clause with any Select cursor that was not declared with the FOR READ ONLY keywords.

You cannot use WHERE CURRENT OF if you are selecting from only one table in a table hierarchy. That is, this clause is not valid with the ONLY keyword.

The WHERE CURRENT OF clause can be used to delete an element from a collection by deleting the current row of the collection-derived table that a collection variable holds. For more information, see “Collection-Derived Table” on page 5-4.

Using the USING or FROM Keyword to Introduce a Join Condition (XPS)

To delete information from a table based on information contained in one or more other tables, use the USING keyword or a second FROM keyword to introduce the list of tables that you want to join in the WHERE clause. When you use this syntax, the WHERE clause can include any complex join.

If you do not list a join in the WHERE clause, the database server ignores the tables listed after the introductory keyword (either USING or FROM). That is, the database server performs the query as if you specified no list of tables.

You can use a second FROM keyword to introduce the list of tables, but your code will be easier to read if you use the USING keyword instead.

When you introduce a list of tables that you want to join in the WHERE clause, the following restrictions for the DELETE statement exist:

- You must list the target table (the one from which rows are to be deleted) and any table that will be part of a join after the USING (or second from) keyword.
- The WHERE clause cannot contain outer joins.
- The target table cannot be a static or external table.
- The statement cannot contain cursors.
- If the target table is a view, the view must be updatable.

That implies that the SELECT statement that defines the view cannot contain any of the following syntax elements:

- Aggregate expressions
- UNIQUE or DISTINCT keywords
- UNION operator
- GROUP BY keywords

The next example deletes the rows from the **lineitem** table whose corresponding rows in the **order** table show a **qty** of less than one.

```
DELETE FROM lineitem USING order o, lineitem l
WHERE o.qty < 1 AND o.order_num = l.order_num;
```

A delete join makes it easier to incorporate new data into a database. For example, you can:

1. Store new values in a temporary table.
2. Use a delete join (DELETE...USING statement) to remove any records from the temporary table that already exist in the table into which you want to insert the new records.
3. Insert the remaining records into the table.

In addition, you can use this syntax instead of deleting from the results of a SELECT statement that includes a join.

In ESQL/C, when you use a delete join, the entire operation occurs as a single transaction. For example, if a delete join query is supposed to delete 100 rows and an error occurs after the 50th row, the first 50 rows that are already deleted will reappear in the table.

Deleting Rows That Contain Opaque Data Types (IDS)

Some opaque data types require special processing when they are deleted. For example, if an opaque data type contains spatial or multirepresentational data, it might provide a choice of how to store the data: inside the internal structure or, for large objects, in a smart large object.

To accomplish this process, call a user-defined support function called **destroy()**. When you use DELETE to remove a row that contains one of these opaque types, the database server automatically invokes **destroy()** for the opaque type. This function decides how to remove the data, regardless of where it is stored. For more information on the **destroy()** support function, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

Deleting Rows That Contain Collection Data Types (IDS)

When a row contains a column that is a collection data type (LIST, MULTISSET, or SET), you can search for a particular element in the collection, and delete the row or rows in which the element is found.

For example, the following statement deletes any rows from the **new_tab** table in which the **set_col** column contains the element jimmy smith:

```
DELETE FROM new_tab WHERE 'jimmy smith' IN set_col;
```

You can also use a collection variable to delete values in a collection column by deleting one or more individual elements in a collection. For more information, see "Collection-Derived Table" on page 5-4 and the examples in "Database Name" on page 5-15 and "Example of Deleting from a Collection" on page 5-12.

Data Types in Distributed DELETE Operations (IDS)

In Dynamic Server, a DELETE statement (or any other SQL data-manipulation language statement) that accesses a database of another Dynamic Server instance can reference only the following data types:

- Built-in data types that are not opaque
- BOOLEAN
- LVARCHAR
- DISTINCT of the built-in data types that are not opaque
- DISTINCT of BOOLEAN
- DISTINCT of LVARCHAR
- DISTINCT of any of the DISTINCT data types that appear in this list.

Cross-server distributed DELETE operations can support these DISTINCT types only if the DISTINCT types are cast explicitly to built-in types, and all of the DISTINCT types, their data type hierarchies, and their casts are defined exactly the same way in each participating database.

Cross-server DML operations cannot reference a column or expression of a complex, large-object, nor user-defined opaque data type (UDT), nor of an unsupported DISTINCT type or built-in opaque type. For additional information

about the data types that Dynamic Server supports in cross-server DML operations, see “Data Types in Cross-Server Transactions” on page 2-537.

Distributed operations that access other databases of the local Dynamic Server instance, however, can access the same data types that are listed above for cross-server operations, and also the following additional data types:

- Most of the *built-in opaque data types*, as listed in “Data Types in Cross-Database Transactions” on page 2-536
- DISTINCT of the built-in types that are referenced in the line above
- DISTINCT of any of the data types that are listed in either of the two lines above
- Opaque user-defined data types (UDTs) that can be cast explicitly to built-in data types.

Cross-database distributed DELETE operations can support these DISTINCT and opaque UDTs only if all the opaque and DISTINCT UDTs are cast explicitly to built-in types, and all of the opaque UDTs, DISTINCT types, data type hierarchies, and casts are defined exactly the same way in each of the participating databases.

Deletes cannot access a database of another database server unless both servers define TCP/IP connections in the DBSERVERNAME or DBSERVERALIASES configuration parameters. This applies to any communication between Dynamic Server instances, even if both database servers reside on the same computer.

SQLSTATE Values in an ANSI-Compliant Database

If no rows satisfy the WHERE clause of a DELETE operation on a table in an ANSI-compliant database, the database server issues a warning. You can detect this warning condition in either of the following ways:

- The GET DIAGNOSTICS statement sets the **RETURNED_SQLSTATE** field to the value 02000. In an SQL API application, the **SQLSTATE** variable contains this same value.
- In an SQL API application, the **sqlca.sqlcode** and **SQLCODE** variables contain the value 100.

The database server also sets **SQLSTATE** and **SQLCODE** to these values if the DELETE . . . WHERE statement is part of a multistatement prepared object, and the database server returns no rows.

SQLSTATE Values in a Database That Is Not ANSI-Compliant

In a database that is not ANSI compliant, the database server does not return a warning when it finds no rows satisfying the WHERE clause of a DELETE statement. In this case, the **SQLSTATE** code is 00000 and the **SQLCODE** code is zero (0). If the DELETE . . . WHERE is part of a multistatement prepared object, however, and no rows are returned, the database server does issue a warning. It sets **SQLSTATE** to 02000 and sets the **SQLCODE** value to 100.

For information on the ANSI/ISO-compliance status of values returned to the **SQLSTATE** variable, see the section “SQLSTATE Support for the ANSI/ISO Standard for SQL” on page 2-387.

Related Information

Related Statements: “DECLARE ” on page 2-286, “FETCH ” on page 2-367, “GET DIAGNOSTICS ” on page 2-386, “INSERT ” on page 2-430, “OPEN ” on page 2-459, “SELECT” on page 2-528, and “UPDATE” on page 2-689

For discussions of the DELETE statement, SPL routines, statement modification, cursors, and the **SQLCODE** code, see the *IBM Informix Guide to SQL: Tutorial*.

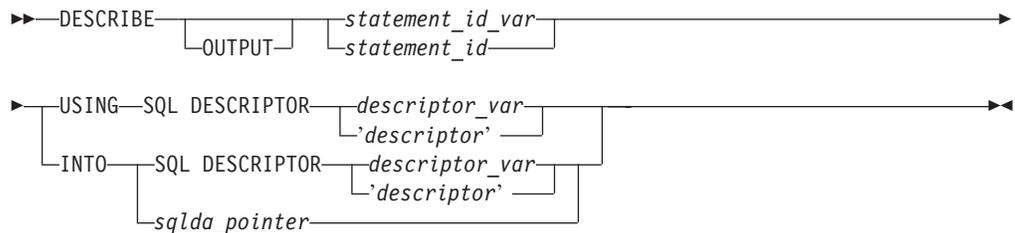
For information on how to access row and collections with Informix ESQL/C host variables, see the chapter on complex data types in the *IBM Informix ESQL/C Programmer’s Manual*.

For a discussion of the GLS aspects of the DELETE statement, see the *IBM Informix GLS User’s Guide*.

DESCRIBE

Use the DESCRIBE statement to obtain information about output parameters and other features of a prepared statement before you execute it. Use this statement with Informix ESQL/C. (See also “DESCRIBE INPUT ” on page 2-315.)

Syntax



Element	Description	Restrictions	Syntax
<i>descriptor</i>	Name of a system-descriptor area	System-descriptor area must already be allocated	“Quoted String” on page 4-189
<i>descriptor_var</i>	Host variable specifying a system-descriptor area	Must contain the name of an allocated system-descriptor area	Language-specific rules for names
<i>sqlda_pointer</i>	Pointer to an sqlda structure	Cannot begin with dollar (\$) sign or colon (:). An sqlda structure is required if dynamic SQL is used.	See the sqlda structure in the <i>IBM Informix ESQL/C Programmer’s Manual</i>
<i>statement_id</i>	Statement identifier for a prepared SQL statement	Must be defined in a previous PREPARE statement	“PREPARE ” on page 2-468; “Identifier ” on page 5-21
<i>statement_id_var</i>	Host variable that contains the value of <i>statement_id</i>	Must be declared in a previous PREPARE statement	Language-specific rules for names

Usage

DESCRIBE can provide information at runtime about a prepared statement:

- The type of SQL statement that was prepared
- Whether an UPDATE or DELETE statement contains a WHERE clause

- In Dynamic Server, for a SELECT, EXECUTE FUNCTION (or EXECUTE PROCEDURE), INSERT, or UPDATE statement, the DESCRIBE statement also returns the number, data types, and size of the values, and the name of the column or expression that the query returns.
- In Extended Parallel Server, for a SELECT, EXECUTE PROCEDURE, or INSERT statement, DESCRIBE also returns the number, types, and size of the values, and the name of the column or expression that the query returns.

With this information, you can write code to allocate memory to hold retrieved values and display or process them after they are fetched.

The OUTPUT Keyword

The OUTPUT keyword specifies that only information about output parameters of the prepared statement are stored in the `sqllda` descriptor area. If you omit this keyword, DESCRIBE can return input parameters, but only for INSERT statements (and for UPDATE, if the `IFX_UPDDESC` environment variable is set in the environment where the database server is initialized).

Describing the Statement Type

The DESCRIBE statement takes a statement identifier from a PREPARE statement as input. When the DESCRIBE statement executes, the database server sets the value of the SQLCODE field of the `sqlca` to indicate the statement type (that is, the keyword with which the statement begins). If the prepared statement text contains more than one SQL statement, the DESCRIBE statement returns the type of the first statement in the text.

SQLCODE is set to zero to indicate a SELECT statement *without* an INTO TEMP clause. This situation is the most common. For any other SQL statement, SQLCODE is set to a positive integer. You can test the number against the constant names that are defined. In Informix ESQL/C, the constant names are defined in the `sqlstypes.h` header file.

The DESCRIBE statement (and the DESCRIBE INPUT statement) use the SQLCODE field differently from any other statement, possibly returning a nonzero value when it executes successfully. You can revise standard error-checking routines to accommodate this behavior, if desired.

Checking for the Existence of a WHERE Clause

If the DESCRIBE statement detects that a prepared statement contains an UPDATE or DELETE statement without a WHERE clause, the DESCRIBE statement sets the `sqlca.sqlwarn.sqlwarn4` variable to W.

When you do not specify a WHERE clause in either a DELETE or UPDATE statement, the database server performs the delete or update operation on the entire table. Check the `sqlca.sqlwarn.sqlwarn4` variable to avoid unintended global changes to your table.

Describing a Statement with Runtime Parameters

If the prepared statement contains parameters for which the number of parameters or parameter data types is to be supplied at runtime, you can describe these input values. If the prepared statement text includes one of the following statements, the DESCRIBE statement returns a description of each column or expression that is included in the list:

- EXECUTE FUNCTION (or EXECUTE PROCEDURE)
- INSERT
- SELECT (without an INTO TEMP clause)
- UPDATE

In Dynamic Server, the **IFX_UPDDESC** environment variable, as described in the *IBM Informix Guide to SQL: Reference*, must be set before you can use DESCRIBE to obtain information about an UPDATE statement.

The description includes the following information:

- The data type of the column, as defined in the table
- The length of the column, in bytes
- The name of the column or expression

For a prepared INSERT or UPDATE statement, DESCRIBE returns only the dynamic parameters (those expressed with a question mark (?) symbol). Using the OUTPUT keyword, however, prevents these from being returned.

You can specify a destination for the returned informations as a new or existing system-descriptor area, or as a pointer to an **sqlda** structure.

A system-descriptor area conforms to the X/Open standards.

Using the SQL DESCRIPTOR Keywords

Use the USING SQL DESCRIPTOR clause to store the description of a prepared statement list in a previously allocated system-descriptor area.

Use the INTO SQL DESCRIPTOR clause to create a new system-descriptor structure and store the description of a statement list in that structure.

To describe one of the previously mentioned statements into a system-descriptor area, DESCRIBE updates the system-descriptor area in these ways:

- It sets the COUNT field in the system-descriptor area to the number of values in the statement list. An error results if COUNT is greater than the number of item descriptors in the system-descriptor area.
- It sets the TYPE, LENGTH, NAME, SCALE, PRECISION, and NULLABLE fields in the system-descriptor area.

In Dynamic Server, if the column has an opaque data type, the database server sets the EXTTYPEID, EXTYPENAME, EXTYPELENGTH, EXTYPEOWNERLENGTH, and EXTYPEOWNERNAME fields of the item descriptor.

- It allocates memory for the DATA field for each item descriptor, based on the TYPE and LENGTH information.

After a DESCRIBE statement is executed, the SCALE and PRECISION fields contain the scale and precision of the column, respectively. If SCALE and PRECISION are set in the SET DESCRIPTOR statement, and TYPE is set to DECIMAL or MONEY, the LENGTH field is modified to adjust for the scale and precision of the decimal value. If TYPE is not set to DECIMAL or MONEY, the values for SCALE and PRECISION are not set, and LENGTH is unaffected.

You must modify system-descriptor-area information with SET DESCRIPTOR statements to show the address in memory that is to receive the described value.

You can change the data type to another compatible type. This change causes data conversion to take place when data values are fetched.

You can use the system-descriptor area in prepared statements that support a USING SQL DESCRIPTOR clause, such as EXECUTE, FETCH, OPEN, and PUT.

The following examples show the use of a system descriptor in a DESCRIBE statement. In the first example, the system descriptor is a quoted string; in the second example, it is an embedded variable name.

```
main()
{
  . . .
  EXEC SQL allocate descriptor 'desc1' with max 3;
  EXEC SQL prepare curs1 FROM 'select * from tab';
  EXEC SQL describe curs1 using sql descriptor 'desc1';
}
EXEC SQL describe curs1 using sql descriptor :desc1var;
```

Using the INTO sqllda Pointer Clause

Use the INTO *sqllda_pointer* clause to allocate memory for an **sqllda** structure and store its address in an **sqllda** pointer. The DESCRIBE statement fills in the allocated memory with descriptive information. Unlike the USING clause, the INTO clause creates new **sqllda** structures to store the output from DESCRIBE.

The DESCRIBE statement sets the **sqllda.sqlld** field to the number of values in the statement list. The **sqllda** structure also contains an array of data descriptors (**sqllvar** structures), one for each value in the statement list. After a DESCRIBE statement is executed, the **sqllda.sqllvar** structure has the **sqltype**, **sqlllen**, and **sqllname** fields set.

In Dynamic Server, if the column has an opaque data type, DESCRIBE...INTO sets the **sqlxid**, **sqltypename**, **sqltypelen**, **sqlownerlen**, and **sqlownername** fields of the item descriptor.

The DESCRIBE statement allocates memory for an **sqllda** pointer once it is declared in a program. The application program, however, must designate the storage area of the **sqllda.sqllvar.sqlldatafields**.

Describing a Collection Variable (IDS)

The DESCRIBE statement can provide information about a collection variable when you use the USING SQL DESCRIPTOR or INTO clause. You must issue the DESCRIBE statement *after* you open the Select or Insert cursor, because it the OPEN...USING statement specifies the name of the collection variable to use.

The next Informix ESQL/C code fragment dynamically selects the elements of the **:a_set** collection variable into a system-descriptor area called **desc1**:

```
EXEC SQL BEGIN DECLARE SECTION;
      client collection a_set;
      int i, set_count;
      int element_type, element_value;
EXEC SQL END DECLARE SECTION;

EXEC SQL allocate collection :a_set;
EXEC SQL allocate descriptor 'desc1';
EXEC SQL select set_col into :a_set from table1;
EXEC SQL prepare set_id from 'select * from table(?)'

EXEC SQL declare set_curs cursor for set_id;
```

```

EXEC SQL open set_curs using :a_set;
EXEC SQL describe set_id using sql descriptor 'desc1';

do
{
    EXEC SQL fetch set_curs using sql descriptor 'desc1';
    ...
    EXEC SQL get descriptor 'desc1' :set_count = count;
    for (i = 1; i <= set_count; i++)
    {
        EXEC SQL get descriptor 'desc1' value :i
            :element_type = TYPE;
        switch
        {
            case SQLINTEGER:
                EXEC SQL get descriptor 'desc1' value :i
                    :element_value = data;
                ...
        } /* end switch */
    } /* end for */
} while (SQLCODE == 0);

EXEC SQL close set_curs;
EXEC SQL free set_curs;
EXEC SQL free set_id;
EXEC SQL deallocate collection :a_set;
EXEC SQL deallocate descriptor 'desc1';

```

Related Information

Related statements: “ALLOCATE DESCRIPTOR ” on page 2-2, “DEALLOCATE DESCRIPTOR ” on page 2-284, “DECLARE ” on page 2-286, “DESCRIBE INPUT ,” “EXECUTE ” on page 2-348, “FETCH ” on page 2-367, “GET DESCRIPTOR ” on page 2-381, “OPEN ” on page 2-459, “PREPARE ” on page 2-468, “PUT ” on page 2-479, and “SET DESCRIPTOR ” on page 2-605

For a task-oriented discussion of the DESCRIBE statement, see the *IBM Informix Guide to SQL: Tutorial*.

For more information about how to use a system-descriptor area and **sqllda**, refer to the *IBM Informix ESQL/C Programmer's Manual*.

DESCRIBE INPUT

Use the DESCRIBE INPUT statement to return input parameter information before a prepared statement is executed.

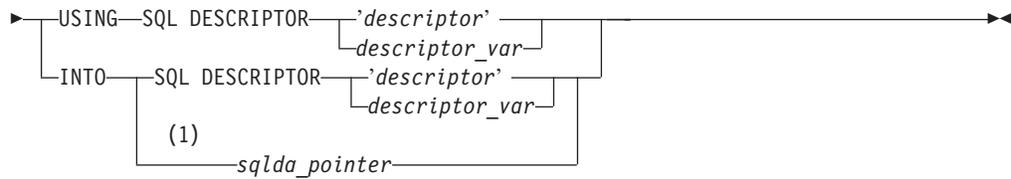
Only Dynamic Server supports this statement. Use this statement with ESQL/C.

Syntax

```

▶▶—DESCRIBE INPUT—statement_var—————▶
                    |
                    |—statement_id—|

```



Notes:

- 1 Informix extension

Element	Description	Restrictions	Syntax
<i>descriptor</i>	Name of a system-descriptor area	System-descriptor area must already be allocated	“Quoted String” on page 4-189
<i>descriptor_var</i>	Host variable specifying a system-descriptor area	Must contain the name of an allocated system-descriptor area	Language-specific rules for names
<i>sqlda_pointer</i>	Pointer to an sqlda structure	Cannot begin with dollar (\$) sign or colon (:). An sqlda structure is required if dynamic SQL is used.	See the sqlda structure in the <i>IBM Informix ESQL/C Programmer’s Manual</i> .
<i>statement_id</i>	Statement identifier for a prepared SQL statement	Must be defined in a previously executed PREPARE statement	“PREPARE ” on page 2-468; “PREPARE ” on page 2-468; “Identifier ” on page 5-21
<i>statement_var</i>	Host variable that contains the value of <i>statement_id</i>	Variable and <i>statement_id</i> both must be declared	Language-specific rules for names

Usage

The DESCRIBE INPUT and the DESCRIBE OUTPUT statements can return information about a prepared statement to an SQL Descriptor Area (**sqlda**):

- For a SELECT, EXECUTE FUNCTION (or PROCEDURE), INSERT, or UPDATE statement, the DESCRIBE statement (with no INPUT keyword) returns the number, data types, and size of the returned values, and the name of the column or expression.
- For a SELECT, EXECUTE FUNCTION, EXECUTE PROCEDURE, DELETE, INSERT, or UPDATE statement, the DESCRIBE INPUT statement returns all the input parameters of a prepared statement.

Tip: Dynamic Server versions earlier than 9.40 do not support the *INPUT* keyword. For compatibility with legacy applications, *DESCRIBE* without *INPUT* is supported. In new applications, you should use DESCRIBE INPUT statements to provide information about dynamic parameters in the WHERE clause, in subqueries, and in other syntactic contexts where the old form of DESCRIBE cannot provide information.

With this information, you can write code to allocate memory to hold retrieved values that you can display or process after they are fetched.

The *IFX_UPDDESC* environment variable does not need to be set before you can use DESCRIBE INPUT to obtain information about an UPDATE statement.

Describing the Statement Type

This statement takes a statement identifier from a PREPARE statement as input. After DESCRIBE INPUT executes, the *SQLCODE* field of the **sqlca** indicates the statement type (that is, the keyword with which the statement begins). If a

prepared object contains more than one SQL statement, DESCRIBE INPUT returns the type of the first statement in the prepared text.

SQLCODE is set to zero to indicate a SELECT statement *without* an INTO TEMP clause. This situation is the most common. For any other SQL statement, SQLCODE is set to a positive integer. You can compare the number with the named constants that are defined in the `sqltypes.h` header file.

The DESCRIBE and DESCRIBE INPUT statements use SQLCODE differently from other statements, under some circumstances returning a nonzero value after successful execution. You can revise standard error-checking routines to accommodate this behavior, if desired.

Checking for Existence of a WHERE Clause

If the DESCRIBE INPUT statement detects that a prepared object contains an UPDATE or DELETE statement without a WHERE clause, the database server sets the `sqlca.sqlwarn.sqlwarn4` variable to W.

When you do not specify a WHERE clause in either a DELETE or UPDATE statement, the database server performs the delete or update action on the entire table. Check the `sqlca.sqlwarn.sqlwarn4` variable after DESCRIBE INPUT executes to avoid unintended global changes to your table.

Describing a Statement with Dynamic Runtime Parameters

If the prepared statement specifies a set of parameters whose cardinality or data types must be supplied at runtime, you can describe these input values. If the prepared statement text includes one of the following statements, the DESCRIBE INPUT statement returns a description of each column or expression that is included in the list:

- EXECUTE FUNCTION (or EXECUTE PROCEDURE)
- INSERT or SELECT
- UPDATE or DELETE

The description includes the following information:

- The data type of the column, as defined in the table
- The length of the column, in bytes
- The name of the column or expression
- Information about *dynamic parameters* (parameters that are expressed as question (?) mark symbols within the prepared statement)

If the database server cannot infer the data type of an expression parameter, the DESCRIBE INPUT statement returns SQLUNKOWN as the data type.

You can specify a destination for the returned informations as a new or existing system-descriptor area, or as a pointer to an `sqlda` structure.

Using the SQL DESCRIPTOR Keywords

Specify INTO SQL DESCRIPTOR to create a new system-descriptor structure and store the description of a prepared statement list in that structure.

Use the USING SQL DESCRIPTOR clause to store the description of a prepared statement list in a previously allocated system-descriptor area. Executing the

DESCRIBE INPUT . . . USING SQL DESCRIPTOR statement updates an existing system-descriptor area in the following ways:

- It allocates memory for the DATA field for each item descriptor, based on the TYPE and LENGTH information.
- It sets the COUNT field in the system-descriptor area to the number of values in the statement list. An error results if COUNT is greater than the number of item descriptors in the system-descriptor area.
- It sets the TYPE, LENGTH, NAME, SCALE, PRECISION, and NULLABLE fields in the system-descriptor area.

For columns of opaque data types, the DESCRIBE INPUT statement sets the EXTTYPEID, EXTYPENAME, EXTYPELENGTH, EXTYPEOWNERLENGTH, and EXTYPEOWNERNAME fields of the item descriptor.

After a DESCRIBE INPUT statement executes, the SCALE and PRECISION fields contain the scale and precision of the column, respectively. If SCALE and PRECISION are set in the SET DESCRIPTOR statement, and TYPE is set to DECIMAL or MONEY, the LENGTH field is modified to adjust for the decimal scale and precision. If TYPE is not set to DECIMAL or MONEY, the values for SCALE and PRECISION are not set, and LENGTH is unaffected.

You must modify the system-descriptor-area information with the SET DESCRIPTOR statement to specify the address in memory that is to receive the described value. You can change the data type to another compatible type. This causes data conversion to take place when the data values are fetched.

You can also use the system-descriptor area in other statements that support a USING SQL DESCRIPTOR clause, such as EXECUTE, FETCH, OPEN, and PUT.

The following examples show the use of a system descriptor in a DESCRIBE statement. In the first example, the system descriptor is a quoted string; in the second example, it is an embedded variable name.

```
main()
{
. . .
EXEC SQL allocate descriptor 'desc1' with max 3;
EXEC SQL prepare curs1 FROM 'select * from tab';
EXEC SQL describe curs1 using sql descriptor 'desc1';
}
EXEC SQL describe curs1 using sql descriptor :desc1var;
```

A system-descriptor area conforms to the X/Open standards.

Using the INTO sqllda Pointer Clause

The INTO *sqllda_pointer* clause allocates memory for an **sqllda** structure and store its address in an **sqllda** pointer. The DESCRIBE INPUT statement fills in the allocated memory with descriptive information.

The DESCRIBE INPUT statement sets the **sqllda.sqlld** field to the number of values in the statement list. The **sqllda** structure also contains an array of data descriptors (**sqllvar** structures), one for each value in the statement list. After a DESCRIBE statement is executed, the **sqllda.sqllvar** structure has the **sqltype**, **sqlllen**, and **sqlname** fields set.

If the column has an opaque data type, DESCRIBE INPUT . . . INTO sets the **sqlxid**, **sqltypename**, **sqltypelen**, **sqlownerlen**, and **sqlownername** fields of the item descriptor.

The DESCRIBE INPUT statement allocates memory for an **sqlda** pointer once it is declared in a program. The application program, however, must designate the storage area of the **sqlda.sqlvar.sqldata** fields.

Describing a Collection Variable

The DESCRIBE INPUT statement can provide information about a collection variable if you use the INTO or USING SQL DESCRIPTOR clause.

You must execute the DESCRIBE INPUT statement *after* you open the Select or Insert cursor. Otherwise, DESCRIBE INPUT cannot get information about the collection variable because it is the OPEN . . . USING statement that specifies the name of the collection variable to use.

The next Informix ESQL/C program fragment dynamically selects the elements of the **:a_set** collection variable into a system-descriptor area called **desc1**:

```
EXEC SQL BEGIN DECLARE SECTION;
    client collection a_set;
    int i, set_count;
    int element_type, element_value;
EXEC SQL END DECLARE SECTION;

EXEC SQL allocate collection :a_set;
EXEC SQL allocate descriptor 'desc1';
EXEC SQL select set_col into :a_set from table1;
EXEC SQL prepare set_id from
    'select * from table(?)';

EXEC SQL declare set_curs cursor for set_id;
EXEC SQL open set_curs using :a_set;
EXEC SQL describe set_id using sql descriptor 'desc1';do
{
    EXEC SQL fetch set_curs using sql descriptor 'desc1';
    ...
    EXEC SQL get descriptor 'desc1' :set_count = count;
    for (i = 1; i <= set_count; i++)
    {
        EXEC SQL get descriptor 'desc1' value :i
            :element_type = TYPE;
        switch
        {
            case SQLINTEGER:
                EXEC SQL get descriptor 'desc1' value :i
                    :element_value = data;
                ...
        } /* end switch */
    } /* end for */
} while (SQLCODE == 0);

EXEC SQL close set_curs;
EXEC SQL free set_curs;
EXEC SQL free set_id;
EXEC SQL deallocate collection :a_set;
EXEC SQL deallocate descriptor 'desc1';
```

Related Information

Related statements: "ALLOCATE DESCRIPTOR " on page 2-2, "DEALLOCATE DESCRIPTOR " on page 2-284, "DECLARE " on page 2-286, "DESCRIBE" on page 2-311, "EXECUTE " on page 2-348, "FETCH " on page 2-367, "GET DESCRIPTOR " on page 2-381, "OPEN " on page 2-459, "PREPARE " on page 2-468, "PUT " on page 2-479, and "SET DESCRIPTOR " on page 2-605

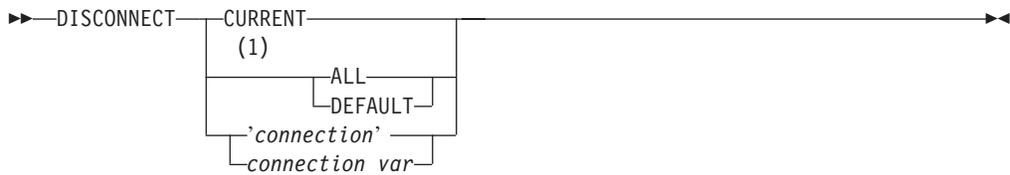
For a task-oriented discussion of the DESCRIBE statement, see the *IBM Informix Guide to SQL: Tutorial*.

For more information about how to use a system-descriptor area and **sqllda**, refer to the *IBM Informix ESQL/C Programmer's Manual*.

DISCONNECT

Use the DISCONNECT statement to terminate a connection between an application and a database server.

Syntax



Notes:

- 1 ESQL/C only

Element	Description	Restrictions	Syntax
<i>connection</i>	String that specifies a connection to terminate	Connection name that the CONNECT statement assigned	"Quoted String" on page 4-189
<i>connection_var</i>	Host variable that holds the name of a connection	Must be a fixed-length character data type	Language specific

Usage

DISCONNECT terminates a connection to a database server. If a database is open, it closes before the connection drops. Even if you made a connection to a specific database only, the connection to the database server is terminated by DISCONNECT. If DISCONNECT does not terminate the current connection, the connection context of the current environment is not changed.

DISCONNECT is not valid as statement text in a PREPARE statement.

In ESQL/C, if you disconnect with *connection* or *connection_var*, DISCONNECT generates an error if the specified connection is not a current or dormant connection.

DEFAULT Option

DISCONNECT DEFAULT disconnects the default connection.

The *default connection* is one of the following connections:

- A connection established by the CONNECT TO DEFAULT statement
- An implicit default connection established by the DATABASE or CREATE DATABASE statement

You can use DISCONNECT to drop the default connection. If the DATABASE statement does not specify a database server, as in the following example, the default connection is made to the default database server:

```
EXEC SQL database 'stores_demo';  
.  
.  
.  
EXEC SQL disconnect default;
```

If the DATABASE statement specifies a database server, as the following example shows, the default connection is made to that database server:

```
EXEC SQL database 'stores_demo@mydbsrvr';  
.  
.  
.  
EXEC SQL disconnect default;
```

In either case, the DEFAULT option of DISCONNECT disconnects this default connection. For more information, see “The DEFAULT Connection Specification” on page 2-88.

Specifying the CURRENT Keyword

The DISCONNECT CURRENT statement terminates the current connection. For example, the DISCONNECT statement in the following program fragment terminates the current connection to the database server **mydbsrvr**:

```
CONNECT TO 'stores_demo@mydbsrvr';  
.  
.  
.  
DISCONNECT CURRENT;
```

When a Transaction is Active

DISCONNECT generates an error during a transaction. The transaction remains active, and the application must explicitly commit it or roll it back. If an application terminates without issuing DISCONNECT (because of a system failure or program error, for example), active transactions are rolled back.

In an ANSI-compliant database, however, if no error is encountered while you exit from DB–Access in non-interactive mode without issuing the CLOSE DATABASE, COMMIT WORK, or DISCONNECT statement, the database server automatically commits any open transaction.

Disconnecting in a Thread-Safe Environment

If you issue the DISCONNECT statement in a thread-safe Informix ESQL/C application, keep in mind that an active connection can be disconnected only from within the thread in which it is active. Therefore, one thread cannot disconnect the active connection of another thread. The DISCONNECT statement generates an error if such an attempt is made.

Once a connection becomes dormant, however, any other thread can disconnect it unless an ongoing transaction is associated with the dormant connection that was

established with the WITH CONCURRENT TRANSACTION clause of CONNECT. If the dormant connection was not established with the WITH CONCURRENT TRANSACTION clause, DISCONNECT generates an error when it tries to disconnect the connection.

For an explanation of connections in a thread-safe Informix ESQL/C application, see "SET CONNECTION" on page 2-587.

Specifying the ALL Option

Use the keyword ALL to terminate all connections established by the application up to that time. For example, the following DISCONNECT statement disconnects the current connection as well as all dormant connections:

```
DISCONNECT ALL;
```

In Informix ESQL/C, the ALL keyword has the same effect on multithreaded applications that it has on single-threaded applications. Execution of the DISCONNECT ALL statement causes all connections in all threads to be terminated. However, the DISCONNECT ALL statement fails if any of the connections is in use or has an ongoing transaction associated with it. If either of these conditions is true, none of the connections is disconnected.

Related Information

Related statements: "CONNECT " on page 2-83, "DATABASE " on page 2-281, and "SET CONNECTION" on page 2-587

For information on multithreaded applications, see the *IBM Informix ESQL/C Programmer's Manual*.

DROP ACCESS_METHOD

Use the DROP ACCESS_METHOD statement to remove a previously defined primary or secondary access method from the database.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax

►►—DROP—ACCESS_METHOD—*access_method*—RESTRICT—►►

Element	Description	Restrictions	Syntax
<i>access_method</i>	Name of access method to drop	Must be registered in sysams system catalog table; cannot be a built-in access method	"Identifier " on page 5-21
<i>owner</i>	Owner of the access method	Must own the access method	" Owner Name " on page 5-45

Usage

The RESTRICT keyword is required. You cannot drop an access method if virtual tables or indexes exist that use the access method. You must be the owner of the access method or have DBA privileges to drop an access method.

If a transaction is in progress, the database server waits to drop the access method until the transaction is committed or rolled back. No other users can execute the access method until the transaction has completed.

Examples

Related Information

For this example, suppose that an access method was created by this statement:

```
CREATE SECONDARY ACCESS_METHOD T_tree
(
  am_getnext = ttree_getnext,
  am_unique,
  am_cluster,
  am_sptype = 'S'
);
```

The following statement drops this access method:

```
DROP ACCESS_METHOD T_tree RESTRICT;
```

Details of existing access methods can be found in the **sysams** system catalog table with the following query:

```
SELECT am_name FROM informix.sysams;
```

Related statements: “ALTER ACCESS_METHOD ” on page 2-5 and “CREATE ACCESS_METHOD ” on page 2-90

For a description of the RESTRICT keyword, see “Specifying RESTRICT Mode” on page 2-343. For more information on primary-access methods and secondary-access methods, see the *IBM Informix Virtual-Table Interface Programmer’s Guide* and the *IBM Informix Virtual-Index Interface Programmer’s Guide*

For a discussion of privileges, see the “GRANT ” on page 2-395 statement or the *IBM Informix Database Design and Implementation Guide*.

DROP AGGREGATE

Use the DROP AGGREGATE statement to drop a user-defined aggregate that you created with the CREATE AGGREGATE statement.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```
►► DROP AGGREGATE [ owner—, ] aggregate ◀◀
```

Element	Description	Restrictions	Syntax
<i>aggregate</i>	Name of the user-defined aggregate to be dropped	Must have been previously created with the CREATE AGGREGATE statement	“Identifier ” on page 5-21
<i>owner</i>	Owner of the aggregate	Must own the aggregate	“ Owner Name ” on page 5-45

Usage

Dropping a user-defined aggregate does not drop the support functions that you defined for the aggregate in the CREATE AGGREGATE statement. The database server does not track dependency of SQL statements on user-defined aggregates that you use in the statements. For example, you can drop a user-defined aggregate that is used in an SPL routine.

The following example drops the user-defined aggregate named **my_avg**:

```
DROP AGGREGATE my_avg
```

Related Information

Related statements: "CREATE AGGREGATE " on page 2-92

For information about how to invoke a user-defined aggregate, see the discussion of user-defined aggregates in the Expression segment. For a description of the **sysaggregates** system catalog table that holds information about user-defined aggregates, see the *IBM Informix Guide to SQL: Reference*. For a discussion of user-defined aggregates, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

DROP CAST

Use the DROP CAST statement to remove an existing cast from the database.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```
►►—DROP CAST—(—source_type—AS—target_type—)—————►
```

Element	Description	Restrictions	Syntax
<i>source_type</i>	Data type that the cast accepts as input	Must exist	"Identifier " on page 5-21; "Data Type " on page 4-21
<i>target_type</i>	Data type returned by the cast	Must exist	"Identifier " on page 5-21; "Data Type " on page 4-21

Usage

You must be owner of the cast or have the DBA privilege to use DROP CAST. Dropping a cast removes its definition from the **syscasts** system catalog table, so the cast cannot be invoked explicitly or implicitly. Dropping a cast has no effect on the user-defined function associated with the cast. Use the DROP FUNCTION statement to remove the user-defined function from the database.

Warning: Do not drop built-in casts, which user **informix** owns. These are required for automatic conversions between built-in data types.

A cast defined on a given data type can also be used on any DISTINCT types created from that source type. If you drop the cast, you can no longer invoke it for the DISTINCT types, but dropping a cast that is defined for a DISTINCT type has

no effect on casts for its source type. When you create a DISTINCT type, the database server automatically defines an explicit cast from the DISTINCT type to its source type and another explicit cast from the source type to the DISTINCT type. When you drop the DISTINCT type, the database server automatically drops these two casts.

Examples

A cast (like this one in the **superstores_demo** database) can be dropped with the DROP CAST statement:

```
DROP CAST (decimal(5,5) AS percent);
```

Details of existing casts can be found in the **syscasts** system catalog table the following SQL:

```
SELECT routine_name, class, argument_type, result_type FROM Syscasts;
```

Related Information

Related statements: “CREATE CAST ” on page 2-95 and “DROP FUNCTION ” on page 2-327.

For more information about data types, refer to the *IBM Informix Guide to SQL: Reference*.

DROP DATABASE

Use the DROP DATABASE statement to delete an entire database, including all system catalog tables, objects, and data.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax

```
►►—DROP DATABASE—| Database Name | (1) —————►►
```

Notes:

- 1 See “Database Name” on page 5-15

Usage

The DROP DATABASE statement is an extension to the ANSI/ISO standard, which does not provide syntax for the destruction of a database.

The following statement drops the **stores_demo** database:

```
DROP DATABASE stores_demo
```

You must have the DBA privilege or be user **informix** to run the DROP DATABASE statement successfully. Otherwise, the database server issues an error message and does not drop the database.

You cannot drop the current database or a database that is currently being used by another user. All the current users of the database must first execute the CLOSE DATABASE statement before DROP DATABASE can be successful.

The DROP DATABASE statement attempts to create an implicit connection to the database that you intend to drop. If a previous CONNECT statement has established an explicit connection to another database, and that connection is still your current connection, the DROP DATABASE statement fails with error -1811. In this case, you must first use the DISCONNECT statement to close the explicit connection before you can execute the DROP DATABASE statement.

The DROP DATABASE statement cannot appear in a multistatement PREPARE statement, nor within an SPL routine.

In a DROP DATABASE operation, the database server acquires a lock on each table in the database and holds the locks until the entire operation is complete. Configure your database server with enough locks to accommodate this fact.

For example, if the database to be dropped has 2500 tables, but fewer than 2500 locks were configured for your database server, the DROP DATABASE statement fails. For more information on how to configure the number of locks available to the database server, see the discussion of the LOCKS configuration parameter in your *IBM Informix Administrator's Reference*.

In DB–Access, use the DROP DATABASE statement with caution. DB–Access does not prompt you to verify that you want to delete the entire database.

In ESQL/C, you can use an unqualified database name in a program or host variable, or you can specify the fully-qualified *database@server* format. For example, the following statement drops the **stores_demo** database of a database server called **gibson95**:

```
EXEC SQL DROP DATABASE stores_demo@gibson95
```

If this statement executes successfully, the **gibson95** database server instance continues to exist, but the **stores_demo** database of that database server no longer exists. For more information, see “Database Name” on page 5-15.

The DROP DATABASE statement cannot appear in a multistatement PREPARE statement.

Related Information

Related statements: “CLOSE DATABASE ” on page 2-80, “CONNECT ” on page 2-83, “CREATE DATABASE ” on page 2-97, “DATABASE ” on page 2-281, and “DISCONNECT ” on page 2-320

DROP DUPLICATE

Use the DROP DUPLICATE statement to remove from the database all duplicate copies of a specified existing table that the CREATE DUPLICATE statement created in a specified dbslice or in specified dbspaces across coservers. The original table is not affected by the DROP DUPLICATE statement.

Only Extended Parallel Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```

>> DROP DUPLICATE OF TABLE owner.table

```

Element	Description	Restrictions	Syntax
<i>owner</i>	Owner of the duplicate table	Must own the duplicate table	" Owner Name " on page 5-45
<i>table</i>	Name of the table for which you want to remove all duplicates	Must exist and must be a duplicated table	"Identifier " on page 5-21

Usage

To drop all duplicate copies of a duplicated table and leave only the original table, enter the DROP DUPLICATE statement. Because duplicate tables are read-only, to update a duplicated table, you must first drop all duplicate copies.

Attached indexes on the copies of the duplicate table are also dropped when DROP DUPLICATE is successfully executed.

Related Information

"CREATE DUPLICATE " on page 2-102

DROP FUNCTION

Use the DROP FUNCTION statement to remove a user-defined function from the database. Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```

>> DROP owner.function (parameter_type)
    SPECIFIC FUNCTION Specific Name

```

Notes:

- 1 See "Specific Name " on page 5-73

Element	Description	Restrictions	Syntax
<i>function</i>	Name of the user-defined function to be dropped	Must exist (that is, be registered) in the database. If the name does not uniquely identify a function, you must enter one or more appropriate values for <i>parameter_type</i> .	"Identifier " on page 5-21
<i>parameter_type</i>	Data type of the parameter	The data type (or list of data types) must be the same data types (and specified in the same order) as those specified in the CREATE FUNCTION statement when the function was created	"Data Type " on page 4-21

Usage

Dropping a user-defined function removes the text and executable versions of the function from the database. (Make sure to keep a copy of the function text somewhere outside the database, in case you need to re-create a function after it is dropped.)

If you do not know if a UDR is a user-defined function or a user-defined procedure, you can drop the UDR by using the DROP ROUTINE statement.

To use the DROP FUNCTION statement, you must be the owner of the user-defined function or have the DBA privilege. To drop an external user-defined function, see also “Dropping an External Routine” on page 2-336.

You cannot drop an SPL function from within the same SPL function.

Dynamic Server can resolve a function by its *specific name*, if the function definition declared a specific name. If you use the specific name in this statement, you must also use the keyword SPECIFIC, as in the following example:

```
DROP SPECIFIC FUNCTION compare_point;
```

Otherwise, if the *function* name is not unique within the database, you must specify enough *parameter_type* information to disambiguate the name. If you use parameter data types to identify a user-defined function, they follow the function name, as in the following example:

```
DROP FUNCTION compare (int, int);
```

But the database server returns an error if it cannot resolve an ambiguous function name whose signature differs from that of another function only in an unnamed ROW-type parameter. (This error cannot be anticipated by the database server when the ambiguous *function* is defined.)

Dropping External Functions

A user-defined function (UDF) written in C language or in the Java language is called an *external function*. External functions must include the External Routine Reference clause that specifies a shared-object filename. By default, only users to whom the DBSA has granted the built-in EXTEND role can create or drop an external function. See the section “Granting the EXTEND Role (IDS)” on page 2-411 for additional information about this feature.

To remove the executable version of a C language routine from shared memory, call the **IFX_UNLOAD_MODULE** function. To replace the executable version of a C routine with another routine, call the **IFX_REPLACE_MODULE** function. Both of these built-in functions are described in “UDR Definition Routines” on page 6-54.

Examples

Most functions can be dropped using SQL statements similar to the following:

```
DROP FUNCTION best_month;
```

If you have more than one function with the same name, however, by using function overloading, the DROP FUNCTION statement must either specify the specific name of the function (if it has one), or the parameter list to uniquely

Usage

In a typical online transaction processing (OLTP) environment, concurrent applications are connected to the database server to perform DML operations. For every query, the optimizer chooses a plan that is based on existing indexes, statistics, and directives. After numerous OLTP transactions, however, the chosen plan might no longer be the best plan for query execution. In this case, dropping an inefficient index can sometimes improve performance.

You must be the owner of the *index* or have the DBA privilege to use the DROP INDEX statement. The following example drops the index **o_num_ix** that **joed** owns. The **stores_demo** database must be the current database:

```
DROP INDEX stores_demo:joed.o_num_ix;
```

You cannot use the DROP INDEX statement to drop a unique constraint, nor to drop an index that supports a constraint; you must use the ALTER TABLE ... DROP CONSTRAINT statement to drop the constraint. When you drop the constraint, the database server automatically drops any index that exists solely to support that constraint. If you attempt to use DROP INDEX to drop an index that is shared by a unique constraint, the database server renames the specified index in the **sysindexes** system catalog table, declaring a new name in this format:

```
[space]<tabid>_<constraint_id>
```

Here *tabid* and *constraint_id* are from the **sysables** and **sysconstraints** system catalog tables, respectively. The **sysconstraints.idxname** column is then updated to something like: " 121_13" (where quotation marks show the leading blank space). If this index is a unique index with only referential constraints sharing it, the index is downgraded to a duplicate index after it is renamed.

The ONLINE Keyword (IDS)

By default, DROP INDEX attempts to place an exclusive lock on the indexed table to prevent all other users from accessing the table while the index is being dropped. The DROP INDEX statement fails if another user already has a lock on the table, or is currently accessing the table at the Dirty Read isolation level.

The DBA can reduce the risk of non-exclusive access errors, and can increase the availability of the indexed table, by including the ONLINE keyword as the last specification of the DROP INDEX statement. The ONLINE keyword instructs the database server to drop the index while minimizing the duration of an exclusive locks. The index can be dropped while concurrent users are accessing the table.

After you issue the DROP INDEX ... ONLINE statement, the query optimizer does not consider using the specified index in subsequent query plans or cost estimates, and the database server does not support any other DDL operations on the indexed table, until after the specified index has been dropped. Query operations that were initiated prior to the DROP INDEX ... ONLINE statement, however, can continue to access the index until the queries are completed.

When no other users are accessing the index, the database server drops the index, and the DROP INDEX ... ONLINE statement terminates execution.

By default, the DROP INDEX ... ONLINE statement does not wait indefinitely for locks to be released. If one or more concurrent sessions hold locks on the table, the statement might fail with error -216 or -113 unless you first issue the SET LOCK MODE TO WAIT statement to specify an indefinite wait. Otherwise, DROP INDEX

... ONLINE uses the waiting period for locks that the DEADLOCK_TIMEOUT configuration parameter specifies, or that a previous SET LOCK MODE statement specified. To avoid locking errors, execute SET LOCK MODE TO WAIT (with no specified limit) before you attempt to drop an index online.

You cannot use the CREATE INDEX statement to declare a new index that has the same identifier until after the specified index has been dropped. No more than one CREATE INDEX ... ONLINE or DROP INDEX ... ONLINE statement can concurrently reference indexes on the same table.

The indexed table can be permanent or temporary, logged or unlogged, and fragmented or non-fragmented. You cannot specify the ONLINE keyword, however, when you drop a virtual index, a clustered index, or an R-tree index.

Examples

An index such as the one found in the *stores_demo* database can be dropped with:

```
DROP INDEX zip_ix;
```

If necessary, you can specify the index name as the fully qualified four-part object name (*database@instance:owner.indexname*), as in the following:

```
DROP INDEX stores_demo@prod:"informix".zip_ix ;
```

Details of existing functions can be found in the *sysprocedures* system catalog table, as in the following:

```
SELECT idxname FROM sysindices ;
```

Related Information

Related statements: “ALTER TABLE” on page 2-44, “CREATE INDEX ” on page 2-124, “CREATE TABLE ” on page 2-191. and “CREATE Temporary TABLE ” on page 2-231. For the effect of indexes on performance, see your *IBM Informix Performance Guide*.

For more information on virtual indexes, see the *IBM Informix Virtual-Index Interface Programmer’s Guide*.

DROP OPCLASS

Use the DROP OPCLASS statement to remove an existing operator class from the database.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```

>> DROP OPCLASS [ owner . ] opclass RESTRICT <<<

```

Element	Description	Restrictions	Syntax
<i>opclass</i>	Name of operator class to be dropped	Must have been created by a previous CREATE OPCLASS statement	“Identifier ” on page 5-21

Element	Description	Restrictions	Syntax
<i>owner</i>	Name of opclass owner	Must own the operator class	" Owner Name " on page 5-45

Usage

You must be the owner of the operator class or have the DBA privilege to use the DROP OPCLASS statement.

The RESTRICT keyword causes DROP OPCLASS to fail if the database contains indexes defined on the operator class that you plan to drop. Therefore, before you drop the operator class, you must use the DROP INDEX statement to drop any dependent indexes.

The following DROP OPCLASS statement drops an operator class called **abs_btree_ops**:

```
DROP OPCLASS abs_btree_ops RESTRICT
```

Related Information

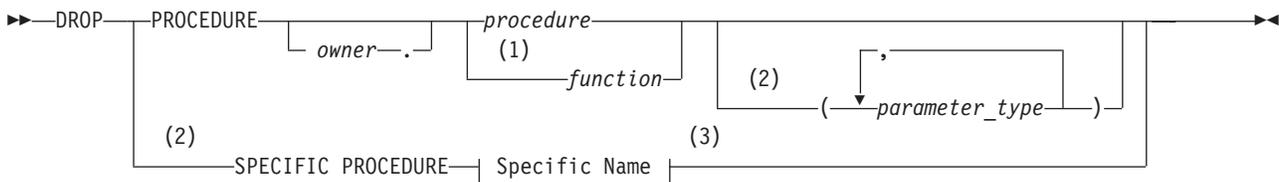
Related statement: "CREATE OPCLASS" on page 2-149, "DROP INDEX " on page 2-329

For information on how to create or extend an operator class, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

DROP PROCEDURE

Use the DROP PROCEDURE statement to drop a user-defined procedure from the database. This statement is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 Stored Procedure Language only
- 2 Dynamic Server only
- 3 See "Specific Name " on page 5-73

Element	Description	Restrictions	Syntax
<i>function</i>	Name of a procedure or SPL function to drop	Must exist (that is, be registered) in the database	"Identifier " on page 5-21
<i>owner</i>	Name of UDR owner	Must own the procedure or SPL function	" Owner Name " on page 5-45

Element	Description	Restrictions	Syntax
<i>parameter_type</i>	The data type of the parameter	The data type (or list of data types) must be the same types (and in the same order) as those specified when the procedure was created	"Identifier " on page 5-21; "Data Type " on page 4-21
<i>procedure</i>	Name of user-defined procedure to drop	Must exist (that is, be registered) in the database	"Database Object Name " on page 5-16

Usage

Dropping a user-defined procedure removes the text and executable versions of the procedure. You cannot drop an SPL procedure within the same SPL procedure.

To use the DROP PROCEDURE statement, you must be the owner of the procedure or have the DBA privilege.

In Extended Parallel Server, which does not support the DROP FUNCTION statement, use the DROP PROCEDURE statement to drop any SPL routine.

Extended Parallel Server does not support ALTER PROCEDURE, ALTER ROUTINE, or ALTER FUNCTION statements. To change the text of an SPL procedure, you must drop the procedure, modify it, and then re-create it. Make sure to keep a copy of the SPL procedure text somewhere outside the database, in case you need to re-create the procedure after it is dropped.

If the *function* or *procedure* name is not unique within the database, you must specify enough *parameter_type* information to disambiguate the name. If the database server cannot resolve an ambiguous UDR name whose signature differs from that of another UDR only in an unnamed ROW type parameter, an error is returned. (This error cannot be anticipated by the database server when the ambiguous *function* or *procedure* is defined.)

If you do not know whether a UDR is a user-defined procedure or a user-defined function, you can use the DROP ROUTINE statement. For more information, see "DROP ROUTINE " on page 2-335.

For compatibility with earlier Dynamic Server versions, you can use this statement to drop an SPL function that CREATE PROCEDURE created. You can include parameter data types after the name of the procedure to identify the procedure:

```
DROP PROCEDURE compare(int, int);
```

In Dynamic Server, if you use the specific name for the user-defined procedure, you must also use the keyword SPECIFIC, as in the following example:

```
DROP SPECIFIC PROCEDURE compare_point;
```

Dropping an External Procedure (IDS)

A user-defined procedure (UDR) written in C language or in the Java language is called an *external routine*. External routines must include the External Routine Reference clause that specifies a shared-object filename. By default, only users to whom the DBSA has granted the built-in EXTEND role can create or drop an external routine. See the section "Granting the EXTEND Role (IDS)" on page 2-411 for additional information about this feature.

To remove the executable version of a C language procedure from shared memory, call the `IFX_UNLOAD_MODULE` function. To replace the executable version of a C routine with another routine, call the `IFX_REPLACE_MODULE` function. Both of these built-in functions are described in “UDR Definition Routines” on page 6-54.

Related Information

Related statements: “CREATE PROCEDURE ” on page 2-153, “CREATE PROCEDURE FROM ” on page 2-163, “DROP FUNCTION ” on page 2-327, “DROP ROUTINE ” on page 2-335, “EXECUTE PROCEDURE ” on page 2-364, and “GRANT ” on page 2-395

For information on how to create user-defined routines, see *IBM Informix User-Defined Routines and Data Types Developer’s Guide*.

DROP ROLE

Use the DROP ROLE statement to remove a user-defined role from the database.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax

```

▶▶ DROP ROLE role
                └─'role'─┘
  
```

Element	Description	Restrictions	Syntax
<i>role</i>	Name of the role to be dropped	Must be registered in the local database. When a <i>role</i> name is enclosed in quotation marks, it is case sensitive.	“ Owner Name ” on page 5-45

Usage

Either the DBA or a user to whom the role was granted with the WITH GRANT OPTION keywords can issue the DROP ROLE statement. (Like a *user* name, a *role* is an authorization identifier, not a database object, so a *role* has no owner.)

After you drop a role, no user can grant or enable the dropped role, and any user who had been assigned the role loses its privileges (such as table-level privileges or routine-level privileges) when the role is dropped, unless the same privileges were granted to PUBLIC or to the user individually. If the dropped role was the default role of a user, the default role for that user becomes NULL.

The following statement drops the role **engineer**:

```
DROP ROLE engineer
```

You cannot use the DROP ROLE statement to drop a built-in role, such as the EXTEND or DBSECADM roles of Dynamic Server.

Related Information

Related statements: "CREATE ROLE " on page 2-164, "GRANT " on page 2-395, "REVOKE" on page 2-494, and "SET ROLE " on page 2-647

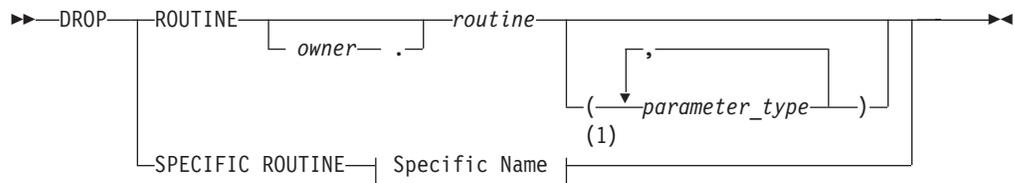
For a discussion of how to use roles, see the *IBM Informix Guide to SQL: Tutorial*.

DROP ROUTINE

Use the DROP ROUTINE statement to remove a user-defined routine (UDR) from the database.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 See "Specific Name " on page 5-73

Element	Description	Restrictions	Syntax
<i>owner</i>	Name of UDR owner	Must own the UDR	" Owner Name " on page 5-45
<i>parameter_type</i>	Data type of a parameter of <i>routine</i>	The data type (or list of data types) must be the same type (and specified in the same order) as in the UDR definition	"Identifier " on page 5-21; "Data Type " on page 4-21
<i>routine</i>	Name of the UDR to drop	The UDR must exist (that is, be registered) in the database	"Identifier " on page 5-21

Usage

Dropping a UDR removes the text and executable versions of the UDR from the database. If you do not know whether a UDR is a user-defined function or a user-defined procedure, this statement instructs the database server to drop the specified user-defined function or user-defined procedure.

To use the DROP ROUTINE statement, you must be the owner of the UDR or have the DBA privilege.

You cannot drop an SPL routine from within the same SPL routine.

Restrictions

To use this statement, the type of UDR cannot be ambiguous. The UDR that you specify must refer to either a user-defined function or a user-defined procedure. If either of the following conditions exist, the database server returns an error:

- The name (and parameters) that you specify apply to both a user-defined procedure and a user-defined function,
- The *specific* name that you specify applies to both a user-defined procedure and a user-defined function.

If the *routine* name is not unique within the database, you must specify enough *parameter_type* information to disambiguate the name. If the database server cannot resolve an ambiguous routine name whose signature differs from that of another routine only in an unnamed ROW type parameter, an error is returned. (This error cannot be anticipated by the database server when the ambiguous *routine* is defined.)

If you use parameter data types to identify a UDR, they follow the UDR name, as in the following example:

```
DROP ROUTINE compare(int, int);
```

If you use the specific name for the UDR, you must also include the keyword SPECIFIC, as in the following example:

```
DROP SPECIFIC ROUTINE compare_point;
```

Modifying SPL Routines (XPS)

Extended Parallel Server does not support the ALTER PROCEDURE, ALTER ROUTINE, or ALTER FUNCTION statements. To change the text of an SPL routine, you must drop the routine and then re-create it. Make sure to keep a copy of the SPL routine text somewhere outside the database, in case you need to re-create the procedure after you drop it.

Dropping an External Routine

A user-defined routine (UDR) written in C language or in the Java language is called an *external routine*. External routines must include the External Routine Reference clause that specifies a shared-object filename. By default, only users to whom the DBSA has granted the built-in EXTEND role can create or drop an external routine. See the section “Granting the EXTEND Role (IDS)” on page 2-411 for additional information about this feature.

To remove the executable version of a C language procedure from shared memory, call the **IFX_UNLOAD_MODULE** function. To replace the executable version of a C routine with another routine, call the **IFX_REPLACE_MODULE** function. Both of these built-in functions are described in “UDR Definition Routines” on page 6-54.

Related Information

Related statements: “CREATE FUNCTION ” on page 2-113, “CREATE PROCEDURE ” on page 2-153, “DROP FUNCTION ” on page 2-327, “DROP PROCEDURE ” on page 2-332, “EXECUTE FUNCTION ” on page 2-356, “EXECUTE PROCEDURE ” on page 2-364, and “GRANT ” on page 2-395

DROP ROW TYPE

Use the DROP ROW TYPE statement to remove an existing named ROW data type from the database.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```

▶▶ DROP ROW TYPE [ owner [, ] row_type ] RESTRICT

```

Element	Description	Restrictions	Syntax
<i>owner</i>	Name of owner of the ROW type	Must be the owner of <i>row_type</i>	"Owner Name " on page 5-45
<i>row_type</i>	Name of an existing named ROW data type to be dropped	Must exist. See also the Usage section that follows.	"Identifier " on page 5-21; "Data Type " on page 4-21

Usage

The DROP ROW TYPE statement removes the entry for the specified *row_type* from the **sysxdtypes** system catalog table. You must be the owner of the specified named ROW data type or have the DBA privilege to execute the DROP ROW TYPE statement.

You cannot drop a named ROW data type if its name is in use. You cannot drop a named ROW data type when any of the following conditions are true:

- Any existing tables or columns are using the named ROW data type.
- The named ROW data type is a supertype in an inheritance hierarchy.
- A view is defined on a column of the named ROW data type.

To drop a named ROW-type column from a table, use ALTER TABLE.

The DROP ROW TYPE statement cannot drop unnamed ROW data types.

The RESTRICT Keyword

The RESTRICT keyword is required with the DROP ROW TYPE statement. RESTRICT causes DROP ROW TYPE to fail if dependencies on *row_type* exist.

The DROP ROW TYPE statement fails and returns an error message if any of the following conditions is true:

- The named ROW data type is used for an existing table or column.
Check the **sysables** and **syscolumns** system catalog tables to find out whether any tables or data types use the named ROW data type.
- The named ROW data type is the supertype in an inheritance hierarchy.
Look in the **sysinherits** system catalog table to see which named ROW data types have child types.

The following statement drops the named ROW data type **employee_t**:

```
DROP ROW TYPE employee_t RESTRICT
```

Related Information

Related statement: "CREATE ROW TYPE " on page 2-168

For a description of the system catalog tables, see the *IBM Informix Guide to SQL: Reference*.

For a discussion of named ROW data types, see the *IBM Informix Guide to SQL: Tutorial*.

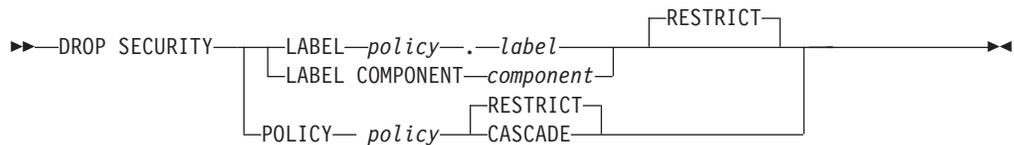
For information on how to create user-defined data types, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

DROP SECURITY

Use the DROP SECURITY statement to remove an existing security object from the current database. The object can be a security policy, or a security label, or a security label component.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Element	Description	Restrictions	Syntax
<i>component</i>	Security label component to drop	Must exist in the database	"Identifier " on page 5-21
<i>label</i>	Security label to drop	Must exist in the database as a label of the specified <i>policy</i>	"Identifier " on page 5-21
<i>policy</i>	Security policy to drop	Must exist in the database	"Identifier " on page 5-21

Usage

Only DBSECADM can issue this statement. When the DROP SECURITY statement executes successfully, the database server deletes any rows that reference the name or the numeric identifier of the specified object from the tables of the system catalog, including these tables:

- **sysecpolicies** for security policies
- **sysseclabels** for security labels
- **sysseclabelcomponents** for security label components.

The keyword or keywords that follow the SECURITY keyword identify the type of security object that is being dropped.

- SECURITY POLICY *policy* specifies a security policy
- SECURITY LABEL *policy.label* specifies a security label
- SECURITY LABEL COMPONENT *component* specifies a security label component.

There is no SQL statement that selectively drops some elements of a security label component without destroying the entire component object. To remove only a subset of the elements of a security label component from the database, DBSECADM can use the `DROP SECURITY LABEL COMPONENT` statement to drop the component, and then redefine the dropped component, using the `CREATE SECURITY LABEL COMPONENT` statement, but without including any elements that are no longer needed. (An alternative is to drop all the security labels that include the deprecated elements, and then use the `CREATE SECURITY LABEL` statement to redefine new labels with the same components as the dropped labels, but without those elements. In this case, the deprecated elements persist in the database, but no security label uses them as values for their component.)

Restrictions on Dropping Security Objects

By default, the `RESTRICT` keyword is in effect when any security object is dropped. Only a security *policy* can be dropped in `CASCADE` mode. DBSECADM cannot drop a security policy in `RESTRICT` mode if any of the following conditions are true:

- A table is protected by that security policy
- A security label depends on that security policy
- A user has been granted an exemption from a rule of that security policy.

A security policy cannot be dropped in `CASCADE` mode if the policy is protecting any table. When a security policy is successfully dropped in `CASCADE` mode, the following security objects are also dropped or revoked:

- All the security labels that are associated with the dropped security policy
- All the security labels that were dropped are also revoked from the users who hold those labels
- All the exemptions from the dropped security policy are revoked.

A security label cannot be dropped in `RESTRICT` mode, which is the only supported mode for dropping labels, if any of the following conditions are true:

- A column is protected by that security label
- A user holds that security label.

A security label component cannot be dropped in `RESTRICT` mode, which is the only supported mode for dropping components, if any security policy depends on that security label component.

Examples of Dropping Security Objects

The following statement instructs the database server to drop the security label **witty**:

```
DROP SECURITY LABEL witty;
```

The statement fails if any column is protected by the **witty** label, or if any user holds this label.

The next example instructs the database server to drop the security label component **adhesive** from the database:

```
DROP SECURITY LABEL COMPONENT adhesive;
```

The statement fails if any security policy depends on the **adhesive** security label component.

The following example instructs the database server to drop the **best** security policy in CASCADE mode:

```
DROP SECURITY POLICY best CASCADE;
```

This statement fails if that policy is currently protecting any table. If this statement succeeds, however, it has the following additional effects because of the CASCADE specification:

- All security labels associated with the **best** security policy are also dropped.
- All exemptions from the **best** security policy are revoked.
- All security labels that were dropped because the **best** security policy was dropped are revoked from all users who hold those labels.

Related Information

Related statements: “ALTER SECURITY LABEL COMPONENT ” on page 2-36, ALTER TABLE, “CREATE SECURITY LABEL ” on page 2-175, “CREATE SECURITY LABEL COMPONENT ” on page 2-177, “CREATE SECURITY POLICY ” on page 2-181, CREATE TABLE, GRANT EXEMPTION, GRANT SECURITY LABEL, “RENAME SECURITY ” on page 2-490, REVOKE EXEMPTION, and REVOKE SECURITY LABEL

For a discussion of LBAC security objects, see your IBM Informix Dynamic Server Security Guide

DROP SEQUENCE

Use the DROP SEQUENCE statement to remove a sequence from the database.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```

▶▶ DROP SEQUENCE [ owner . ] sequence

```

Element	Description	Restrictions	Syntax
<i>owner</i>	Name of sequence owner	Must own the sequence object	“ Owner Name ” on page 5-45
<i>sequence</i>	Name of a sequence	Must exist in the current database	“Identifier ” on page 5-21

Usage

This statement removes the *sequence* entry from the **syssequences** system catalog table. To drop a sequence, you must be its owner or have the DBA privilege on the database. In an ANSI-compliant database, you must qualify the name of the sequence with the name of its owner (*owner.sequence*) if you are not the owner.

If you drop a sequence, any synonyms for the name of the sequence are also dropped automatically by the database server.

You cannot use a synonym to specify the identifier of the *sequence* in the DROP SEQUENCE statement.

Examples

Related Information

Suppose you had a sequence created with something code similar to the following:

```
CREATE SEQUENCE Invoice_Numbers
  START 10000 INCREMENT 1 NOCYCLE ;
```

Such a sequence can be dropped using the following:

```
DROP SEQUENCE Invoice_Numbers;
```

Details of existing sequences can be found by joining the *syssequences* and *systables* system catalog tables as in the following:

```
SELECT t.tabname SeqName
  FROM Syssequences s, Systables t
  WHERE t.tabid = s.tabid ;
```

Related statements: “ALTER SEQUENCE ” on page 2-40, “CREATE SEQUENCE” on page 2-184, “RENAME SEQUENCE ” on page 2-491, “GRANT ” on page 2-395, “REVOKE” on page 2-494, “INSERT ” on page 2-430, “UPDATE” on page 2-689, and “SELECT” on page 2-528

DROP SYNONYM

Use the DROP SYNONYM statement to remove an existing synonym.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax

```
▶▶ DROP SYNONYM [owner—.] synonym ▶▶
```

Element	Description	Restrictions	Syntax
<i>owner</i>	Name of synonym owner	Must own the synonym	“ Owner Name ” on page 5-45
<i>synonym</i>	Name of synonym to be dropped	The synonym and the table, view, or sequence object to which the synonym points must exist in the database	“Identifier ” on page 5-21

Usage

This removes the entries for the *synonym* from the **systables**, **sys synonyms**, and **sys syn table** system catalog tables. You must own the *synonym* or have the DBA privilege to execute the DROP SYNONYM statement. Dropping a synonym has no effect on the table, view, or sequence object to which the synonym points.

The following statement drops the synonym **nj_cust**, which **cathyg** owns:

```
DROP SYNONYM cathyg.nj_cust
```

DROP SYNONYM is not the only DDL operation that can unregister a synonym. If a table, view, or sequence is dropped, any synonyms that are in the same database and that refer to that table, view, or sequence object are also dropped.

If a synonym refers to an external table or view that is dropped, however, the synonym remains in place until you explicitly drop it using DROP SYNONYM. You can create another table, view, or synonym in place of the dropped table or view and give the new database object the name of the dropped table or view. The old synonym then refers to the new database object. For a complete discussion of synonym chaining, see the CREATE SYNONYM statement.

Related Information

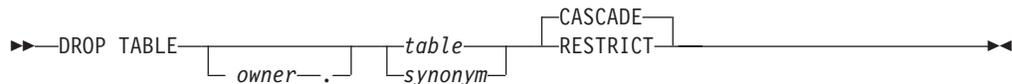
Related statement: "CREATE SYNONYM " on page 2-187

For a discussion of synonyms, see the *IBM Informix Guide to SQL: Tutorial*.

DROP TABLE

Use the DROP TABLE statement to remove a table with its associated indexes and data. This statement is an extension to the ANSI/ISO standard for SQL.

Syntax



Element	Description	Restrictions	Syntax
<i>owner</i>	Name of table owner	Must own the table	" Owner Name " on page 5-45
<i>synonym</i>	Local synonym for a table that is to be dropped	The synonym and its table must exist, and USERTABLENAME must not be set to 1	"Identifier " on page 5-21
<i>table</i>	Name of a table to drop	Must be registered in the systables system catalog table of the local database	"Identifier " on page 5-21

Usage

You must be the owner of the table or have the DBA privilege to use the DROP TABLE statement.

You cannot drop an Extended Parallel Server table that includes a dependent GK index unless that index is entirely dependent on the affected table.

If you issue a DROP TABLE statement, DB-Access does not prompt you to verify that you want to delete an entire table.

Effects of the DROP TABLE Statement

Use the DROP TABLE statement with caution. When you remove a table, you also delete the data stored in it, the indexes or constraints on the columns (including all the referential constraints placed on its columns), any local synonyms assigned to

it, any triggers created on it, and any access privileges granted on the table. You also drop all views based on the table and any violations and diagnostics tables associated with the table.

DROP TABLE does not remove any synonyms for the table that were created in an external database. To remove external synonyms for the dropped table, you must do so explicitly with the DROP SYNONYM statement.

You can prevent users from specifying a synonym in DROP TABLE statements by setting the USETABLENAME environment variable. If USETABLENAME is set, an error results if any user attempts to specify DROP TABLE *synonym*.

Specifying CASCADE Mode

The CASCADE keyword in DROP TABLE removes related database objects, including referential constraints built on the table, views defined on the table, and any violations and diagnostics tables associated with the table.

In Dynamic Server, if the table is the supertable in an inheritance hierarchy, CASCADE drops all of the subtables as well as the supertable.

The CASCADE mode is the default mode of the DROP TABLE statement. You can also specify this mode explicitly with the CASCADE keyword.

Specifying RESTRICT Mode

The RESTRICT keyword can control the drop operation for supertables, for tables that have referential constraints and views defined on them, or for tables that have violations and diagnostics tables associated with them. Using the RESTRICT option causes the drop operation to fail and an error message to be returned if any of the following conditions are true:

- Existing referential constraints reference *table*.
- Existing views are defined on *table*.
- Any violations tables or diagnostics tables are associated with *table*.
- For Dynamic Server, the *table* is the supertable in an inheritance hierarchy.

Dropping a Table That Contains Opaque Data Types (IDS)

Some opaque data types require special processing when they are deleted. For example, if an opaque type contains spatial or multi-representational data, it might provide a choice of how to store the data: inside the internal structure or, for large objects, in a smart large object.

The database server removes opaque types by calling a user-defined support function called **destroy()**. When you execute the DROP TABLE statement on a table whose rows contain an opaque type, the database server automatically invokes the **destroy()** function for the type. The **destroy()** function can perform certain operations on columns of the opaque data type before the table is dropped. For more information about the **destroy()** support function, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

Tables That Cannot Be Dropped

Restrictions exist on the types of tables that you can drop:

- You cannot drop any system catalog table.
- You cannot drop a table that is not in the current database.

- You cannot drop a violations table or (in Dynamic Server) a diagnostics table. Before you can drop such a table, you must first issue a STOP VIOLATIONS TABLE statement on the base table with which the violations and diagnostics tables are associated.
- With Extended Parallel Server, you cannot drop a table specified in the FROM clause of a SELECT statement that defines a GK index.

The following example removes two tables in the current database. Both are owned by **joed**, the current user. Neither table has an associated violations or diagnostics table, nor a referential constraint or view defined on it.

```
DROP TABLE customer;
DROP TABLE stores_demo@acctng:joed.state;
```

Related Information

Related statements: “CREATE TABLE ” on page 2-191, “DROP DATABASE ” on page 2-325, and “TRUNCATE (IDS) ” on page 2-677.

For a discussion of the data integrity of tables, see the *IBM Informix Guide to SQL: Tutorial*. For a discussion of how to create a table, see the *IBM Informix Database Design and Implementation Guide*.

DROP TRIGGER

Use the DROP TRIGGER statement to remove a trigger definition from the database.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax

```

▶▶ DROP TRIGGER [ owner—, ] trigger

```

Element	Description	Restrictions	Syntax
<i>owner</i>	Name of the owner of the trigger	Must own the trigger	“ Owner Name ” on page 5-45
<i>trigger</i>	Name of the trigger to drop	The trigger must exist in the local database	“Identifier ” on page 5-21

Usage

You must be the owner of the trigger or have the DBA privilege to drop the trigger. Dropping a trigger removes the text of the trigger definition and the executable trigger from the database. The row describing the specified trigger is deleted from the **systriggers** system catalog table.

In Dynamic Server, dropping an INSTEAD OF trigger on a complex view (a view with columns from more than one table) revokes any privileges on the view that the owner of the trigger received automatically when creating the trigger, and also revokes any privileges that the owner of the trigger granted to other users. (Dropping a trigger on a simple view does not revoke any privileges.)

The following statement drops the **items_pct** trigger:

```
DROP TRIGGER items_pct;
```

If a DROP TRIGGER statement appears inside an SPL routine that is called by a data manipulation (DML) statement, the database server returns an error.

When multiple triggers are defined on the same table or view for the same triggering event, the order in which the triggers execute is not guaranteed. If you have a preferred sequence of execution, but the triggers are executing in some other sequence, you might wish to drop all of the triggers except the one that you want to run first, and then re-create the other triggers in the relative order in which you want them to execute, so that they are listed in the system catalog in the intended order of execution.

Related Information

Related statements: “CREATE TRIGGER” on page 2-239, “DROP TABLE ” on page 2-342

DROP TYPE

Use the DROP TYPE statement to remove a user-defined distinct or opaque data type from the database. (You cannot use this to drop a built-in data type.)

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```
►► DROP TYPE owner . data_type RESTRICT ◀◀
```

Element	Description	Restrictions	Syntax
<i>data_type</i>	Name of distinct or opaque data type to be removed	Must be an existing user-defined distinct or opaque type in the local database; cannot be a built-in data type	“Identifier ” on page 5-21
<i>owner</i>	Name of data type owner	Must own the data type	“ Owner Name ” on page 5-45

Usage

To drop a distinct or opaque data type with the DROP TYPE statement, you must be the owner of the data type or have the DBA privilege. When you use this statement, you remove the data type definition from the database (in the **sysxdtypes** system catalog table). In general, this statement does not remove any definitions of casts or of support functions associated with that data type.

Important: When you drop a distinct type, the database server automatically drops the two explicit casts between the distinct type and the type on which it is based.

The DROP TYPE statement fails with an error if you attempt to drop a built-in data type, such as the built-in opaque BOOLEAN or LVARCHAR type, or the built-in distinct IDSSECURITYLABEL type.

You cannot drop a distinct or opaque type if the database contains any casts, columns, or user-defined functions whose definitions reference the data type.

The following statement drops the **new_type** data type:

```
DROP TYPE new_type RESTRICT;
```

Related Information

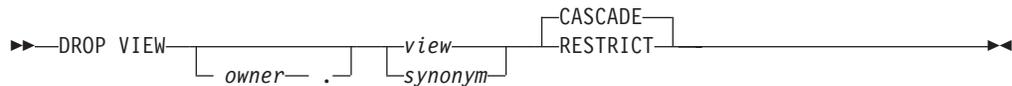
Related statements: “CREATE DISTINCT TYPE ” on page 2-100, “CREATE OPAQUE TYPE ” on page 2-145, “CREATE ROW TYPE ” on page 2-168, “DROP ROW TYPE ” on page 2-336, and “CREATE TABLE ” on page 2-191

DROP VIEW

Use the DROP VIEW statement to remove a view from the database.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax



Element	Description	Restrictions	Syntax
<i>owner</i>	Name of view owner	Must own the view	“ Owner Name ” on page 5-45
<i>synonym</i>	Synonym for a view that this statement drops	The <i>synonym</i> and the view to which it points must exist in the local database	“Identifier ” on page 5-21
<i>view</i>	Name of a view to drop	Must exist in sysables	“Identifier ” on page 5-21

Usage

To drop a view, you must be the owner or have the DBA privilege.

When you drop a view, you also drop any other views and INSTEAD OF triggers whose definitions depend on that view. (You can also specify this default behavior explicitly with the CASCADE keyword.)

When you include the RESTRICT keyword in the DROP VIEW statement, the drop operation fails if any other existing views are defined on *view*; otherwise, these dependent views would be unregistered by the DROP VIEW operation.

You can query the **sysdepend** system catalog table to determine which views, if any, depend on another view.

The following statement drops the view that is named **cust1**:

```
DROP VIEW cust1
```

Related Information

Related statements: “CREATE VIEW ” on page 2-274 and “DROP TABLE ” on page 2-342

For a discussion of views, see the *IBM Informix Guide to SQL: Tutorial*.

DROP XADATASOURCE

Use the DROP XADATASOURCE statement to drop a previously defined XA-compliant data source from the system catalog of the database.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

►►—DROP—XADATASOURCE—*xa_source*—RESTRICT—◄◄

Element	Description	Restrictions	Syntax
<i>xa_source</i>	The XA-compliant data source to drop	Must be present in the sysxadatasources system catalog table	"Identifier " on page 5-21

Usage

The RESTRICT keyword is required. You must be the owner of the XA data source or hold DBA privileges to drop an access method.

The following statement drops the XA data source instance called **NewYork** that is owned by user **informix**.

```
DROP XADATASOURCE informix.NewYork RESTRICT;
```

You cannot drop an access method if it is being used in a transaction that is currently open. If an XA data source has been registered with a transaction that is not complete, you can drop the data source only after the database is closed or the session ends.

Related Information

Related statements: "CREATE XADATASOURCE " on page 2-279, "CREATE XADATASOURCE TYPE " on page 2-280, "DROP XADATASOURCE TYPE"

For a description of the RESTRICT keyword, see "Specifying RESTRICT Mode" on page 2-343. For information on how to use XA data sources, see the *IBM Informix Datablade API Programmer's Guide*.

DROP XADATASOURCE TYPE

Use the DROP XADATASOURCE TYPE statement to drop a previously defined XA-compliant data source type from the database.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

►►—DROP—XADATASOURCE TYPE—*xa_type*—RESTRICT—►►

Element	Description	Restrictions	Syntax
<i>xa_type</i>	Name of the XA data source type to be dropped	Must be present in the sysxasourcetypes system catalog table	"Identifier " on page 5-21

Usage

The RESTRICT keyword is required. You cannot unregister an XA data source type if virtual tables or indexes exist that use the data source. You must be user **informix** or have DBA privileges to drop an XA data source type.

The following statement drops an XA data source type called **MQSeries** owned by user **informix**:

```
DROP XADATASOURCE TYPE informix.MQSeries RESTRICT;
```

You cannot drop an XA data source type until after all the XA data source instances that use that data source type have been dropped.

Related Information

Related statements: "CREATE XADATASOURCE " on page 2-279, "CREATE XADATASOURCE TYPE " on page 2-280, "DROP XADATASOURCE" on page 2-347

For a description of the RESTRICT keyword, see "Specifying RESTRICT Mode" on page 2-343. For information on how to use XA data sources, see the *IBM Informix Datablade API Programmer's Guide*.

EXECUTE

Use the EXECUTE statement to run a previously prepared statement or a multiple-statement prepared object.

Only Dynamic Server supports this statement. Use this statement with Informix ESQL/C.

Syntax

►►—EXECUTE—*stmt_id*—*stmt_id_var*—(1) INTO Clause—►►

►►—(2) USING Clause—►►

Notes:

- 1 See "INTO Clause" on page 2-350
- 2 See "USING Clause" on page 2-353

Element	Description	Restrictions	Syntax
<i>stmt_id</i>	Identifier of a prepared SQL statement	Must have been declared in a previous PREPARE statement	"Identifier " on page 5-21
<i>stmt_id_var</i>	Host variable containing the identifier of a prepared statement	Must exist and must contain a statement identifier that a previous PREPARE statement declared, and must be of a character data type	"PREPARE " on page 2-468

Usage

The EXECUTE statement passes a prepared SQL statement to the database server for execution. The following example shows an EXECUTE statement within an Informix ESQL/C program:

```
EXEC SQL PREPARE del_1 FROM
    'DELETE FROM customer WHERE customer_num = 119';
EXEC SQL EXECUTE del_1;
```

Once prepared, an SQL statement can be executed as often as needed.

After you release the database server resources (using a FREE statement), you cannot use the statement identifier with a DECLARE cursor or with the EXECUTE statement until you prepare the statement again.

If the statement contained question mark (?) placeholders, use the USING clause to provide specific values for them before execution. For more information, see the "USING Clause" on page 2-353.

You can execute any prepared statement except those in the following list:

- A prepared SELECT statement that returns more than one row
When you use a prepared SELECT statement to return multiple rows of data, you must use a cursor to retrieve the data rows. As an alternative, you can EXECUTE a prepared SELECT INTO TEMP statement to achieve the same result.
For more information on cursors, see "DECLARE " on page 2-286.
- A prepared EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement for an SPL function that returns more than one row
When you prepare an EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement to invoke an SPL function that returns multiple rows, you must use a cursor to retrieve the data rows.
For more information on how to execute a SELECT or an EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement, see "PREPARE " on page 2-468.

If you create or drop a trigger after you prepare a triggering INSERT, DELETE, or UPDATE statement, the prepared statement returns an error when you execute it.

Scope of Statement Identifiers

A program can consist of one or more source-code files. By default, the scope of reference of a statement identifier is global to the program. A statement identifier created in one file can be referenced from another file.

In a multiple-file program, if you want to limit the scope of reference of a statement identifier to the file in which it is executed, you can preprocess all the files with the **-local** command-line option.

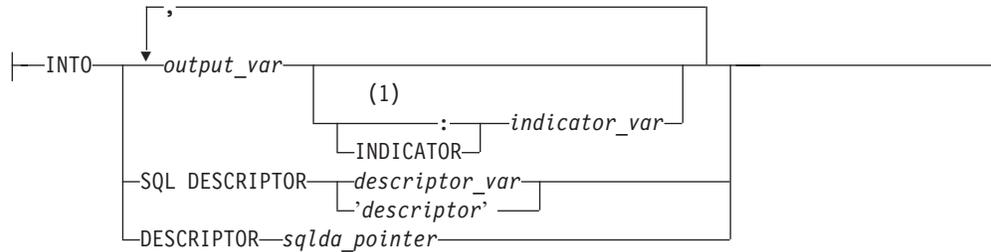
INTO Clause

Use the INTO clause to save the returned values of these SQL statements:

- A prepared singleton SELECT statement that returns only one row of column values for the columns in the select list
- A prepared EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement for an SPL function that returns only one set of values

The INTO clause of the EXECUTE statement has the following syntax:

INTO Clause:



Notes:

1 Informix extension

Element	Description	Restrictions	Syntax
<i>descriptor</i>	Quoted string that identifies a system-descriptor area	Must already be allocated. Use single (') quotation marks	"Quoted String" on page 4-189
<i>descriptor_var</i>	Host variable that identifies a system-descriptor area	System-descriptor area must already be allocated	Language specific
<i>indicator_var</i>	Host variable that receives a return code if corresponding <i>parameter_var</i> is NULL value, or if truncation occurs	Cannot be DATETIME or INTERVAL data type	Language specific
<i>output_var</i>	Host variable whose contents replace a question-mark (?) placeholder in a prepared statement	Must be a character data type	Language specific
<i>sqllda_pointer</i>	Pointer to an sqllda structure that defines data type and memory location of values to replace a question-mark (?) placeholder in a prepared object	Cannot begin with a dollar sign (\$) or a colon (:) symbol. An sqllda structure is required with dynamic SQL	"DESCRIBE INPUT" on page 2-315

This closely resembles the syntax of the "USING Clause" on page 2-353.

The INTO clause provides a concise and efficient alternative to more complicated and lengthy syntax. In addition, by placing values into variables that can be displayed, the INTO clause simplifies and enhances your ability to retrieve and display data values. For example, if you use the INTO clause, you do not need to use a cursor to retrieve values from a table.

You can store the returned values in output variables, in output SQL descriptors, or in output **sqllda** pointers.

Restrictions with the INTO Clause

If you execute a prepared SELECT statement that returns more than one row, or execute a prepared EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement for an SPL function that returns more than one group of return values, you receive an error message. In addition, if you prepare and declare a statement and then attempt to execute that statement, you receive an error message.

You cannot select a NULL value from a table column and place that value into an output variable. If you know in advance that a table column contains a NULL value, after you select the data, check the indicator variable that is associated with the column to determine if the value is NULL.

To use the INTO clause with the EXECUTE statement

1. Declare the output variables that the EXECUTE statement uses.
2. Use PREPARE to prepare your SELECT statement or to prepare your EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement.
3. Use the EXECUTE statement, with the INTO clause, to execute your SELECT statement or to execute your EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement.

Replacing Placeholders with Parameters

You can specify any of the following items to replace the question-mark placeholders in a prepared statement before you execute it:

- A host variable name (if the number and data type of the parameters are known at compile time)
- A system descriptor that identifies a system descriptor area
- A descriptor that is a pointer to an `sqllda` structure

Sections that follow describe each of these options for specifying parameters.

Saving Values In Host or Program Variables

If you know the number of return values to be supplied at runtime and their data types, you can define the values that the SELECT or EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement returns as host variables in your program. Use these host variables with the INTO keyword, followed by the names of the variables. These variables are matched with the return values in a one-to-one correspondence, from left to right.

You must supply one variable name for each value that the SELECT or EXECUTE FUNCTION (or EXECUTE PROCEDURE) returns. The data type of each variable must be compatible with the corresponding returned value from the prepared statement.

Saving Values in a System-Descriptor Area

If you do not know the number of return values to be supplied at runtime or their data types, you can associate output values with a system-descriptor area. A system-descriptor area describes the data type and memory location of one or more values.

A system-descriptor area conforms to the X/Open standards.

To specify a system-descriptor area as the location of output values, use the INTO SQL DESCRIPTOR clause of the EXECUTE statement. Each time that the EXECUTE statement is run, the values that the system-descriptor area describes are stored in the system-descriptor area.

The following example shows how to use the system-descriptor area to execute prepared statements in IBM Informix ESQL/C:

```
EXEC SQL allocate descriptor 'desc1';
...
sprintf(sel_stmt, "%s %s %s",
        "select fname, lname from customer",
        "where customer_num =",
        cust_num);
EXEC SQL prepare sel1 from :sel_stmt;
EXEC SQL execute sel1 into sql descriptor 'desc1';
```

The COUNT field corresponds to the number of values that the prepared statement returns. The value of COUNT must be less than or equal to the value of the occurrences that were specified when the system-descriptor area was allocated with the ALLOCATE DESCRIPTOR statement.

You can obtain the value of a field with the GET DESCRIPTOR statement and set the value with the SET DESCRIPTOR statement.

For more information, refer to the discussion of the system-descriptor area in the *IBM Informix ESQL/C Programmer's Manual*.

Saving Values in an sqlda Structure (ESQL/C)

If you do not know the number of output values to be returned at runtime or their data types, you can associate output values from an **sqlda** structure. An **sqlda** structure lists the data type and memory location of one or more return values. To specify an **sqlda** structure as the location of return values, use the INTO DESCRIPTOR clause of the EXECUTE statement. Each time the EXECUTE statement is run, the database server places the returns values that the **sqlda** structure describes into the **sqlda** structure.

The next example uses an **sqlda** structure to execute a prepared statement:

```
struct sqlda *pointer2;
...
sprintf(sel_stmt, "%s %s %s",
        "select fname, lname from customer",
        "where customer_num =",
        cust_num);
EXEC SQL prepare sel1 from :sel_stmt;
EXEC SQL describe sel1 into pointer2;
EXEC SQL execute sel1 into descriptor pointer2;
```

The **sqlda.sqld** value specifies the number of output values that are described in occurrences of **sqlvar**. This number must correspond to the number of values that the SELECT or EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement returns.

For more information, refer to the **sqlda** discussion in the *IBM Informix ESQL/C Programmer's Manual*.

This example uses the INTO clause with an EXECUTE statement in Informix ESQL/C:

```
EXEC SQL prepare sell from 'select fname, lname from customer
    where customer_num =123';
EXEC SQL execute sell into :fname, :lname using :cust_num;
```

The next example uses the INTO clause to return multiple rows of data:

```
EXEC SQL BEGIN DECLARE SECTION;
int customer_num =100;
char fname[25];
EXEC SQL END DECLARE SECTION;

EXEC SQL prepare sell from 'select fname from customer
    where customer_num=?';
for ( ;customer_num < 200; customer_num++)
{
    EXEC SQL execute sell into :fname using customer_num;
    printf("Customer number is %d\n", customer_num);
    printf("Customer first name is %s\n\n", fname);
}
```

The sqlca Record and EXECUTE

After an EXECUTE statement, the **sqlca** can reflect two results:

- The **sqlca** can reflect an error within the EXECUTE statement.
For example, when an UPDATE ...WHERE statement in a prepared statement processes zero rows, the database server sets **sqlca** to 100.
- The **sqlca** can reflect the success or failure of the executed statement.

Returned SQLCODE Values with EXECUTE

If a prepared statement fails to access any rows when it executes, the database server returns the **SQLCODE** value of zero (0).

For a multistatement prepared object, however, if any statement in the following list fails to access rows, the returned **SQLCODE** value is SQLNOTFOUND (= 100):

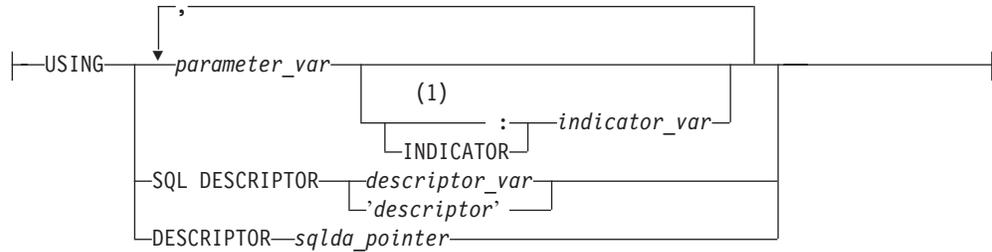
- INSERT INTO *table* SELECT ... WHERE
- SELECT...WHERE ... INTO TEMP
- DELETE ... WHERE
- UPDATE ... WHERE

In an ANSI-compliant database, if you prepare and execute any of the statements in the preceding list, and no rows are returned, the returned **SQLCODE** value is SQLNOTFOUND (= 100).

USING Clause

Use the USING clause to specify the values that are to replace question-mark (?) placeholders in the prepared statement. Providing values in the EXECUTE statement that replace the question-mark (?) placeholders in the prepared statement is sometimes called *parameterizing* the prepared statement.

USING Clause:



Notes:

1 Informix extension

Element	Description	Restrictions	Syntax
<i>descriptor</i>	Quoted string that identifies a system-descriptor area	System-descriptor area must already be allocated. Use single (') quotation marks.	“Quoted String” on page 4-189
<i>descriptor_var</i>	Host variable that identifies a system-descriptor area	System-descriptor area must already be allocated	Language specific
<i>indicator_var</i>	Host variable that receives a return code if corresponding <i>parameter_var</i> is NULL value, or if truncation occurs	Cannot be DATETIME or INTERVAL data type	Language specific
<i>parameter_var</i>	Host variable whose contents replace a question-mark (?) placeholder in a prepared statement	Must be a character data type	Language specific
<i>sqllda_pointer</i>	Pointer to an sqllda structure that defines data type and memory location of values to replace question-mark (?) placeholder in a prepared object	Cannot begin with a dollar sign (\$) or a colon (:). An sqllda structure is required with dynamic SQL	“DESCRIBE INPUT ” on page 2-315

This closely resembles the syntax of the “INTO Clause” on page 2-350.

If you know the number of parameters to be supplied at runtime and their data types, you can define the parameters that are needed by the EXECUTE statement as host variables in your program.

If you do not know the number of parameters to be supplied at runtime or their data types, you can associate input values from a system-descriptor area or an **sqllda** structure. Both of these descriptor structures describe the data type and memory location of one or more values to replace question-mark (?) placeholders.

Supplying Parameters Through Host or Program Variables

You pass parameters to the database server by opening the cursor with the USING keyword, followed by the names of the variables. These variables are matched with prepared statement question-mark (?) placeholders in a one-to-one correspondence, from left to right. You must supply one storage- parameter variable for each placeholder. The data type of each variable must be compatible with the corresponding value that the prepared statement requires.

The following example executes the prepared UPDATE statement in Informix ESQL/C:

```
stcopy ("update orders set order_date = ?
       where po_num = ?", stm1);
EXEC SQL prepare statement_1 from :stm1;
EXEC SQL execute statement_1 using :order_date, :po_num;
```

Supplying Parameters Through a System Descriptor

You can create a system-descriptor area that describes the data type and memory location of one or more values and then specify the descriptor in the USING SQL DESCRIPTOR clause of the EXECUTE statement.

Each time that the EXECUTE statement is run, the values that the system-descriptor area describes are used to replace question-mark (?) placeholders in the PREPARE statement.

The COUNT field corresponds to the number of dynamic parameters in the prepared statement. The value of COUNT must be less than or equal to the number of item descriptors that were specified when the system-descriptor area was allocated with the ALLOCATE DESCRIPTOR statement.

The following example shows how to use system descriptors to execute a prepared statement in Informix ESQL/C:

```
EXEC SQL execute prep_stmt using sql descriptor 'desc1';
```

Supplying Parameters Through an sqlda Structure (ESQL/C)

You can specify the **sqlda** pointer in the USING DESCRIPTOR clause of the EXECUTE statement.

Each time the EXECUTE statement is run, the values that the descriptor structure describes are used to replace question-mark (?) placeholders in the PREPARE statement.

The **sqlda.sqld** value specifies the number of input values that are described in occurrences of **sqlvar**. This number must correspond to the number of dynamic parameters in the prepared statement.

The following example shows how to use an **sqlda** structure to execute a prepared statement in Informix ESQL/C:

```
EXEC SQL execute prep_stmt using descriptor pointer2;
```

Related Information

Related statements: "ALLOCATE DESCRIPTOR " on page 2-2, "DEALLOCATE DESCRIPTOR " on page 2-284, "DECLARE " on page 2-286, "EXECUTE IMMEDIATE" on page 2-360, "FETCH " on page 2-367, "GET DESCRIPTOR " on page 2-381, "PREPARE " on page 2-468, "PUT " on page 2-479, and "SET DESCRIPTOR " on page 2-605

For a task-oriented discussion of the EXECUTE statement, see the *IBM Informix Guide to SQL: Tutorial*.

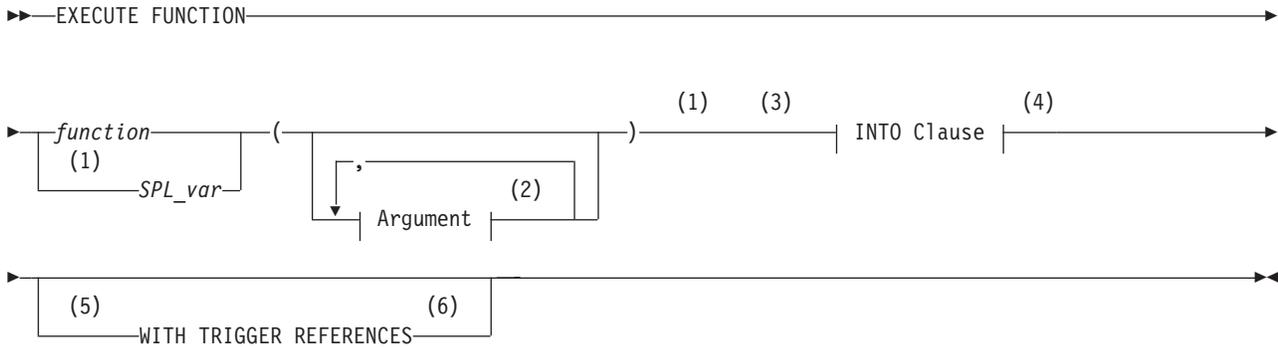
For more information about concepts that relate to the EXECUTE statement, see the *IBM Informix ESQL/C Programmer's Manual*.

EXECUTE FUNCTION

Use the EXECUTE FUNCTION statement to invoke a user-defined function or a built-in routine that returns a value.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 Stored Procedure Language only
- 2 See "Arguments " on page 5-1
- 3 ESQL/C only
- 4 See "INTO Clause" on page 2-358
- 5 Trigger functions only
- 6 See "The WITH TRIGGER REFERENCES Keywords" on page 2-366

Element	Description	Restrictions	Syntax
<i>function</i>	Name of a user-defined function to execute	Must be registered in the database	"Database Object Name " on page 5-16
<i>SPL_var</i>	Variable that contains the name of an SPL routine to be executed	Must be a CHAR, VARCHAR, NCHAR, or NVARCHAR data type that contains the non-NULL name of an existing SPL function	"Identifier " on page 5-21

Usage

The EXECUTE FUNCTION statement invokes a user-defined function (UDF), with arguments, and specifies where the results are to be returned.

An external C or Java language function returns exactly one value.

An SPL function can return one or more values.

You cannot use the EXECUTE FUNCTION statement to invoke any type of user-defined procedure that returns no value. Instead, use the EXECUTE PROCEDURE or EXECUTE ROUTINE statement to execute procedures.

You must have the Execute privilege on the user-defined function.

For more information, see “GRANT ” on page 2-395.

In ANSI/ISO-compliant databases that support implicit transactions, the EXECUTE FUNCTION statement does not, by default, begin a new transaction. SQL statements within the invoked function, however, can begin a new transaction.

Negator Functions and Their Companions

If a UDF that returns a BOOLEAN value has a companion function, any user who executes the function must have the Execute privilege on both the function and on its companion. For example, if a function has a negator function, any user who executes the function must have the Execute privilege on both the function and its negator. In addition, the companion function must have the same owner as its negator function.

For information on how to designate a UDF as the companion to its negator function, see “ NEGATOR” on page 5-63.

How the EXECUTE FUNCTION Statement Works

For a user-defined function to be executed with the EXECUTE FUNCTION statement, the following conditions must exist:

- The qualified function name or the function signature (the function name with its parameter list) must be unique within the name space or database.
- The function must exist.
- The function must not have any OUT nor INOUT parameters.

If EXECUTE FUNCTION specifies fewer arguments than the user-defined function expects, the unspecified arguments are said to be *missing*. Missing arguments are initialized to their corresponding parameter default values, if these were defined. The syntax of specifying default values for parameters is described in “Routine Parameter List” on page 5-67.

EXECUTE FUNCTION returns an error under the following conditions:

- It specifies more arguments than the user-defined function expects.
- One or more arguments are missing and do not have default values.
Unlike Dynamic Server, Extended Parallel Server can invoke a function with missing arguments, but the function might fail if a missing argument causes another error (for example, an undefined value for a variable).
- The fully qualified function name or the function signature is not unique.
- No function with the specified name or signature that you specify is found.
- You use it to attempt to execute a user-defined procedure.

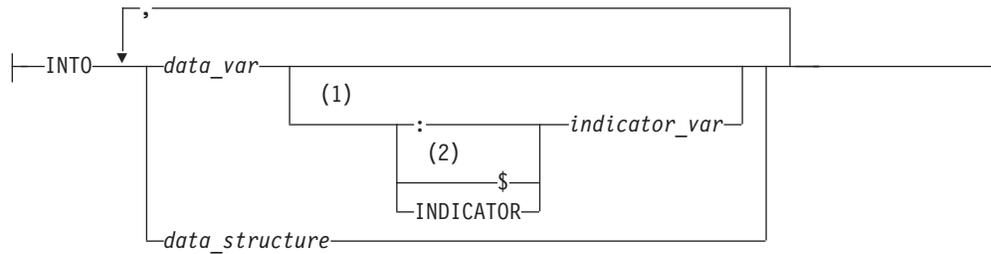
If the *function* name is not unique within the database, you must specify enough *parameter_type* information to disambiguate the name. See the section “Arguments ” on page 5-1 for additional information about how to specify parameters when invoking a function.

In Dynamic Server, the *specific name* of an external UDR is valid in some DDL statements, but is not valid in contexts where you invoke the function.

If Dynamic Server cannot resolve an ambiguous function name whose signature differs from that of another routine only in an unnamed-ROW type parameter, an error is returned. (This error cannot be anticipated by the database server when the ambiguous function is defined.)

INTO Clause

INTO Clause:



Notes:

- 1 ESQL/C only
- 2 Informix extension

Element	Description	Restrictions	Syntax
<i>data_structure</i>	Structure that was declared as a host variable	Individual elements of structure must be compatible with the data types of the returned values	Language specific
<i>data_var</i>	Variable to receive the value that a user-defined function returns	See "Data Variables."	Language specific
<i>indicator_var</i>	Program variable to store a return code if the corresponding <i>data_var</i> receives a NULL value	Use an indicator variable if the value of the corresponding <i>data_var</i> might be NULL	Language specific

You must include an INTO clause with EXECUTE FUNCTION to specify the variables that receive the values that a user-defined function returns. If the function returns more than one value, the values are returned into the list of variables in the order in which you specify them.

If the EXECUTE FUNCTION statement stands alone (that is, it is not part of a DECLARE statement and does not use the INTO clause), it must execute a noncursor function. A noncursor function returns only one row of values. The following example shows a SELECT statement in IBM Informix ESQ/C:

```
EXEC SQL execute function
      cust_num(fname, lname, company_name) into :c_num;
```

Data Variables

If you issue EXECUTE FUNCTION within an Informix ESQ/C program, *data_var* must be a host variable. Within an SPL routine, *data_var* must be an SPL variable.

If you issue EXECUTE FUNCTION within a CREATE TRIGGER statement, *data_var* must be a column name in the triggering table or in another table.

INTO Clause with Indicator Variables (ESQ/C)

You should use an indicator variable if the possibility exists that data returned from the user-defined function is NULL. For more information about indicator variables, see the *IBM Informix ESQ/C Programmer's Manual*.

INTO Clause with Cursors

If EXECUTE FUNCTION calls a UDF that returns more than one row of values, it must execute a cursor function. A cursor function can return one or more rows of values and must be associated with a Function cursor to execute.

If the SPL function returns more than one row or a collection data type, you must access the rows or collection elements with a cursor.

To return more than one row of values, an external function (one written in the C or Java language) must be defined as an iterator function. For more information on iterator functions, see the *IBM Informix Datablade API Programmer's Guide*.

In an SPL routine, if a SELECT returns more than one row, you must use the FOREACH statement to access the rows individually. The INTO clause of the SELECT statement can store the fetched values. For more information, see "FOREACH " on page 3-22.

To return more than one row of values, an SPL function must include the WITH RESUME keywords in its RETURN statement. For more information on how to write SPL functions, see the *IBM Informix Guide to SQL: Tutorial*.

In an IBM Informix ESQL/C program, the DECLARE statement can declare a Function cursor and the FETCH statement can return rows individually from the cursor. You can put the INTO clause in the EXECUTE FUNCTION or in the FETCH statement, but you cannot put it in both. The following IBM Informix ESQL/C code examples show different ways you can use the INTO clause:

- Using the INTO clause in the EXECUTE FUNCTION statement:

```
EXEC SQL declare f_curs cursor for
    execute function get_orders(customer_num)
    into :ord_num, :ord_date;
EXEC SQL open f_curs;
while (SQLCODE == 0)
    EXEC SQL fetch f_curs;
EXEC SQL close f_curs;
```

- Using the INTO clause in the FETCH statement:

```
EXEC SQL declare f_curs cursor for
    execute function get_orders(customer_num);
EXEC SQL open f_curs;
while (SQLCODE == 0)
    EXEC SQL fetch f_curs into :ord_num, :ord_date;
EXEC SQL close f_curs;
```

Alternatives to PREPARE ... EXECUTE FUNCTION ... INTO

In ESQL/C, you cannot prepare an EXECUTE FUNCTION statement that includes the INTO clause. For similar functionality, however, follow these steps:

1. Prepare the EXECUTE FUNCTION statement with no INTO clause.
2. Declare a Function cursor for the prepared statement.
3. Open the cursor.
4. Execute the FETCH statement with an INTO clause to fetch the returned values into program variables.

Alternatively, you can do the following:

1. Declare a cursor for the EXECUTE FUNCTION statement without first preparing the statement, and include the INTO clause in the EXECUTE FUNCTION when you declare the cursor.
2. Open the cursor.
3. Fetch the returned values from the cursor without using the INTO clause of the FETCH statement.

Dynamic Routine-Name Specification of SPL Functions

Dynamic routine-name specification simplifies the writing of an SPL function that calls another SPL routine whose name is not known until runtime. To specify the name of an SPL routine in the EXECUTE FUNCTION statement, instead of listing the explicit name of an SPL routine, you can use an SPL variable to hold the routine name. For more information about how to execute SPL functions dynamically, see the *IBM Informix Guide to SQL: Tutorial*.

Related Information

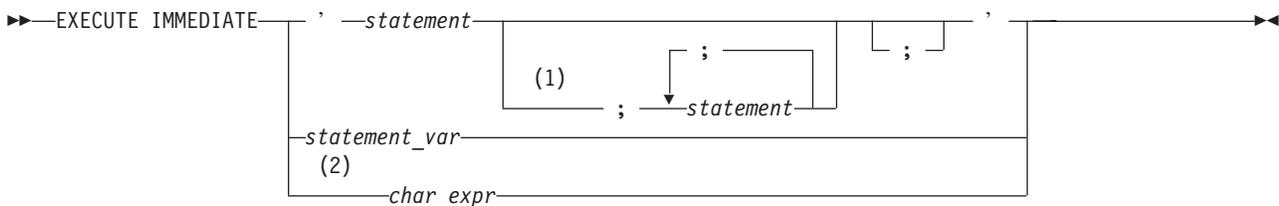
“Receiving Input from the Called UDR” on page 3-4, “CREATE FUNCTION ” on page 2-113, “CREATE FUNCTION FROM ” on page 2-122, “DROP FUNCTION ” on page 2-327, “DROP ROUTINE ” on page 2-335, “EXECUTE PROCEDURE ” on page 2-364, “FOREACH ” on page 3-22

EXECUTE IMMEDIATE

Use the EXECUTE IMMEDIATE statement to perform tasks equivalent to what the PREPARE, EXECUTE, and FREE statements accomplish, but as a single operation.

Use this Dynamic SQL statement with Informix ESQL/C and SPL.

Syntax



Notes:

- 1 ESQL/C only
- 2 SPL only

Element	Description	Restrictions	Syntax
<i>char_expr</i>	Expression that evaluates to a character data type	Must evaluate to a CHAR, LVARCHAR, NCHAR, NVARCHAR, or VARCHAR data type	“Expression” on page 4-40
<i>statement</i>	Text of a valid SQL statement	See the same sections that are listed below for <i>statement_var</i>	See this chapter.

Element	Description	Restrictions	Syntax
<i>statement _var</i>	Variable containing <i>statement</i> or (in ESQL/C) a semicolon-separated list of statements	Must be a previously declared variable of type CHAR, NCHAR, NVARCHAR, or VARCHAR (or in SPL, LVARCHAR). See “EXECUTE IMMEDIATE and Restricted Statements” and “Restrictions on Valid Statements” on page 2-362.	Language specific

Usage

The EXECUTE IMMEDIATE statement dynamically executes a single SQL statement (or in ESQL/C routines, a semicolon-separated list of SQL statements) that is constructed during program execution. For example, you can obtain the name of a database from program input, construct the DATABASE statement as a program variable, and then use EXECUTE IMMEDIATE to execute the statement, which opens the specified database.

Within ESQL/C routines, the statement text specified by the variable or quoted string can include more than one SQL statement, if consecutive statements are separated by a semicolon (;) delimiter. In SPL routines, however, only one statement can be included. The *statement* cannot be an SPL statement, but can be any SQL statement that is not listed in the sections “EXECUTE IMMEDIATE and Restricted Statements” or “Restrictions on Valid Statements” on page 2-362.

The specification that follows the IMMEDIATE keyword, if valid, is parsed and executed; then all data structures and memory resources are released immediately. Unless you use EXECUTE IMMEDIATE, these operations would otherwise require separate PREPARE, EXECUTE, and FREE statements.

The session environment values (such as the EXTDIRECTIVES, OPTCOMPIND, or USELASTCOMMITTED settings of the ESQL/C or SPL routine that issues the EXECUTE IMMEDIATE statement) override the corresponding ONCONFIG parameter values, if these are different.

The maximum length of an EXECUTE IMMEDIATE statement is 64 kilobytes in ESQL/C routines, or 32 kilobytes in UDRs written in the SPL language.

In ANSI/ISO-compliant databases that support implicit transactions, the EXECUTE IMMEDIATE statement does not, by default, begin a new transaction. Execution of the specified SQL statement text, however, can begin a new transaction.

EXECUTE IMMEDIATE and Restricted Statements

The EXECUTE IMMEDIATE statement cannot execute the following SQL statements.

- CLOSE
- CONNECT
- DECLARE
- DISCONNECT
- EXECUTE
- EXECUTE FUNCTION
- EXECUTE PROCEDURE
- FETCH

- FLUSH
- FREE
- GET DESCRIPTOR
- GET DIAGNOSTICS
- OPEN
- OUTPUT
- PREPARE
- PUT
- SELECT
- SET AUTOFREE
- SET CONNECTION
- SET DEFERRED_PREPARE
- SET DESCRIPTOR
- WHENEVER

For EXECUTE PROCEDURE, this restriction applies only to calls that return one or more values.

The only form of the SELECT statement that EXECUTE IMMEDIATE supports as statement text is SELECT ... INTO TEMP *table*. For the syntax of the INTO TEMP *table* clause in SELECT statements, see “INTO Table Clauses” on page 2-574.

In addition, ESQL/C cannot use the EXECUTE IMMEDIATE statement to execute the following statements in text that contains multiple SQL statements that are separated by semicolons:

- CLOSE DATABASE
- CREATE DATABASE
- DATABASE
- DROP DATABASE
- SELECT (except SELECT INTO TEMP)

The EXECUTE IMMEDIATE statement cannot process SQL statement text that includes question mark (?) symbols as placeholders. Use the PREPARE statement and either a cursor or the EXECUTE statement to execute a dynamically constructed SELECT statement.

(In SPL routines, the EXECUTE IMMEDIATE statement can execute only a single SQL statement. If the specification that immediately follows the IMMEDIATE keyword evaluates to a list of multiple SQL statements, or by a NULL value, or text that is not a valid SQL statement, the database server issues a runtime error.)

Restrictions on Valid Statements

The following restrictions apply to the statements contained in the character expression, quoted string, or statement variable that immediately follows the EXECUTE IMMEDIATE keywords:

- The SQL statement cannot contain a host-language comment.
- Names of host-language variables are not recognized as such in prepared text.

The only identifiers that you can use are names registered in the system catalog of the current database, such as table names and column names.

- The statement cannot reference a host-variable list or a descriptor; it must not contain any question-mark (?) placeholders, which are allowed with a PREPARE statement.
- The text must not include any embedded SQL statement prefix, such as the dollar sign (\$) or the keywords EXEC SQL.
Although it is not required, the SQL statement terminator (;) can be included in the statement text.
- In Dynamic Server, a SELECT or INSERT statement specified within the EXECUTE IMMEDIATE statement cannot contain a Collection-Derived Table clause.
EXECUTE IMMEDIATE cannot process input host variables, which are required for a collection variable. Use the EXECUTE statement or a cursor to process prepared accesses to collection variables.

Handling Exceptions from EXECUTE IMMEDIATE Statements

If the Informix ESQL/C parser detects a syntax error when the EXECUTE IMMEDIATE statement is compiled, it issues a compilation error, and no executable UDR is produced until the syntax is corrected and recompiled. If the parser accepts the EXECUTE IMMEDIATE syntax and the UDR compiles successfully, but an exception occurs during a call to the UDR when the EXECUTE IMMEDIATE statement is executing, the database server issues an error at runtime. Runtime errors can be trapped by the WHENEVER statement, or by some other exception-handling mechanism in the program logic of the UDR.

For routines written in the SPL language, SQL expressions are evaluated at runtime, not when the routine is compiled or optimized. If an expression that follows the IMMEDIATE keyword specifies invalid SQL statement text, Dynamic Server issues a runtime exception, rather than a compilation error. After any runtime error condition in an SPL routine, program control passes to the ON EXCEPTION statement block (if this is defined); otherwise, execution of the UDR terminates abnormally, and an error is returned to the calling context. For information on how to handle runtime errors in SPL routines, see the descriptions of the SPL statement "ON EXCEPTION " on page 3-38. (See also the built-in SQL function **SQLCODE**.)

Examples of the EXECUTE IMMEDIATE Statement

The following ESQL/C examples show EXECUTE IMMEDIATE statements in Informix ESQL/C. Both examples use host variables that contain a CREATE DATABASE statement. The first example uses the SQL statement terminator (;) inside the quoted string.

```
sprintf(cdb_text1, "create database ", usr_db_id);
EXEC SQL execute immediate :cdb_text;

sprintf(cdb_text2, "create database %s", usr_db_id);
EXEC SQL execute immediate :cdb_text;
```

Both examples use host variables that contain a CREATE DATABASE statement. The first example uses the SQL statement terminator (;) inside the quoted string.

The next example shows an SPL program fragment that declares local SPL variables and assigns to them portions of the text of two DDL statements. It then issues an EXECUTE IMMEDIATE statement to drop a table called **DYN_TAB**, specifying the DROP TABLE statement text in an SPL variable. The second EXECUTE IMMEDIATE statement in this example creates a table of the same

Element	Description	Restrictions	Syntax
<i>procedure</i>	User-defined procedure to execute	Must exist	"Database Object Name " on page 5-16
<i>SPL_var</i>	Variable that contains the name of the SPL routine to execute	Must be a character data type that contains the non-NULL name of an SPL routine.	"Identifier " on page 5-21

Usage

The EXECUTE PROCEDURE statement invokes a user-defined procedure and specifies its arguments.

For compatibility with earlier Dynamic Server versions, you can use the EXECUTE PROCEDURE statement to execute an SPL function that the CREATE PROCEDURE statement defined.

In Extended Parallel Server, use the EXECUTE PROCEDURE statement to execute any SPL routine. Extended Parallel Server does not support the EXECUTE FUNCTION statement.

If the EXECUTE PROCEDURE statement returns more than one row, the result set must be processed within a FOREACH loop of an SPL routine, or else accessed through a cursor of an ESQL/C routine.

In ANSI/ISO-compliant databases that support implicit transactions, the EXECUTE PROCEDURE statement does not, by default, begin a new transaction. SQL statements within the invoked procedure, however, can begin a new transaction.

Causes of Errors

EXECUTE PROCEDURE returns an error under the following conditions:

- It has more arguments than the called procedure expects.
- One or more arguments are missing and do not have default values. Unlike Dynamic Server, Extended Parallel Server can invoke procedures with missing arguments, but the procedure might fail if a missing argument causes another error (for example, an undefined value for a variable).
- The fully qualified procedure name or the routine signature is not unique.
- No procedure with the specified name or signature is found.

If the *procedure* name is not unique within the database, you must specify enough *parameter_type* information to disambiguate the name. See "Arguments " on page 5-1 for additional information about how to specify parameters when invoking a procedure. (In Dynamic Server, the *specific name* of an external UDR is valid in DDL statements, but is not valid in contexts where you invoke the procedure.)

Using the INTO Clause

Use the INTO clause to specify where to store the values that the SPL function returns.

If an SPL function returns more than one value, the values are returned into the list of variables in the order in which you specify them. If an SPL function returns more than one row or a collection data type, you must access the rows or collection elements with a cursor.

You cannot prepare an EXECUTE PROCEDURE statement that has an INTO clause. For more information, see “Alternatives to PREPARE ... EXECUTE FUNCTION ... INTO” on page 2-359.

The WITH TRIGGER REFERENCES Keywords

You must include the WITH TRIGGER REFERENCES keywords when you use the EXECUTE PROCEDURE statement to invoke a trigger procedure.

A trigger procedure is an SPL routine that EXECUTE PROCEDURE can invoke only from the FOR EACH ROW section of the Action clause of a trigger definition. Such procedures must include the REFERENCING clause and the FOR clause in the CREATE PROCEDURE statement that defined the procedure. This REFERENCING clause declares names for correlated variables that the procedure can use to reference the *old* column value in the row when the trigger event occurred, or the *new* value of the column after the row was modified by the trigger. The FOR clause specifies the table or view on which the trigger is defined.

Example of Invoking a Trigger Procedure

The following example defines three tables and a trigger procedure that references one of these tables in its FOR clause:

```
CREATE TABLE tab1 (col1 INT,col2 INT);
CREATE TABLE tab2 (col1 INT);
CREATE TABLE temptab1
    (old_col1 INTt, new_col1 INT, old_col2 INT, new_col2 INT);

/* The following procedure is invoked from an INSERT trigger in this example.
*/
CREATE PROCEDURE proc1()
REFERENCING OLD AS o NEW AS n FOR tab1;

IF (INSERTING) THEN -- INSERTING Boolean operator
    LET n.col1 = n.col1 + 1; -- You can modify new values.
    INSERT INTO temptab1 VALUES(0,n.col1,1,n.col2);
END IF

IF (UPDATING) THEN -- UPDATING Boolean operator
    -- you can access relevant old and new values.
    INSERT INTO temptab1 values(o.col1,n.col1,o.col2,n.col2);
END IF

if (SELECTING) THEN -- SELECTING Boolean operator
    -- you can access relevant old values.
    INSERT INTO temptab1 VALUES(o.col1,0,o.col2,0);
END IF

if (DELETING) THEN -- DELETING Boolean operator
    DELETE FROM temptab1 WHERE temptab1.col1 = o.col1;
END IF

END PROCEDURE;
```

This example illustrates that the triggered action can be a different DML operation from the triggering event. Although this procedure inserts a row when an Insert trigger calls it, and deletes a row when a Delete trigger calls it, it also performs INSERT operations if it is called by a Select trigger or by an Update trigger.

The **proc1()** trigger procedure in this example uses Boolean conditional operators that are valid only in trigger routines. The INSERTING operator returns true only if the procedure is called from the FOR EACH ROW action of an INSERT trigger. This procedure can also be called from other triggers whose trigger event is an

UPDATE, SELECT, or DELETE. statement, because the UPDATING, SELECTING and DELETING operators return true (t) if the procedure is invoked in the triggered action of the corresponding type of triggering event.

The following statement defines an Insert trigger on **tab1** that calls **proc1()** from the FOR EACH ROW section as its triggered action, and perform an INSERT operation that activates this trigger:

```
CREATE TRIGGER ins_trig_tab1 INSERT ON tab1 REFERENCING NEW AS post
  FOR EACH ROW(EXECUTE PROCEDURE proc1() WITH TRIGGER REFERENCES);
```

Note that the REFERENCING clause of the trigger declares a correlation name for the NEW value that is different from the correlation name that the trigger procedure declared. These names do not need to match, because the correlation name that was declared in the trigger procedure has that procedure as its scope of reference. The following statement activates the **ins_trig_tab1** trigger, which executes the **proc1()** procedure.

```
INSERT INTO tab1 VALUES (111,222);
```

Because the trigger procedure increments the new value of **col1** by 1, the value inserted is (112, 222), rather than the value that the triggering event specified.

Dynamic Routine-Name Specification of SPL Procedures

Dynamic routine-name specification simplifies the writing of an SPL routine that calls another SPL routine whose name is not known until runtime. To specify the name of an SPL routine in the EXECUTE PROCEDURE statement, instead of listing the explicit name of an SPL routine, you can use an SPL variable to hold the routine name.

If the SPL variable names an SPL routine that returns a value (an SPL function), include the INTO clause of EXECUTE PROCEDURE to specify a *receiving variable* (or variables) to hold the value (or values) that the SPL function returns. For more information on how to execute SPL procedures dynamically, see the *IBM Informix Guide to SQL: Tutorial*.

Related Information

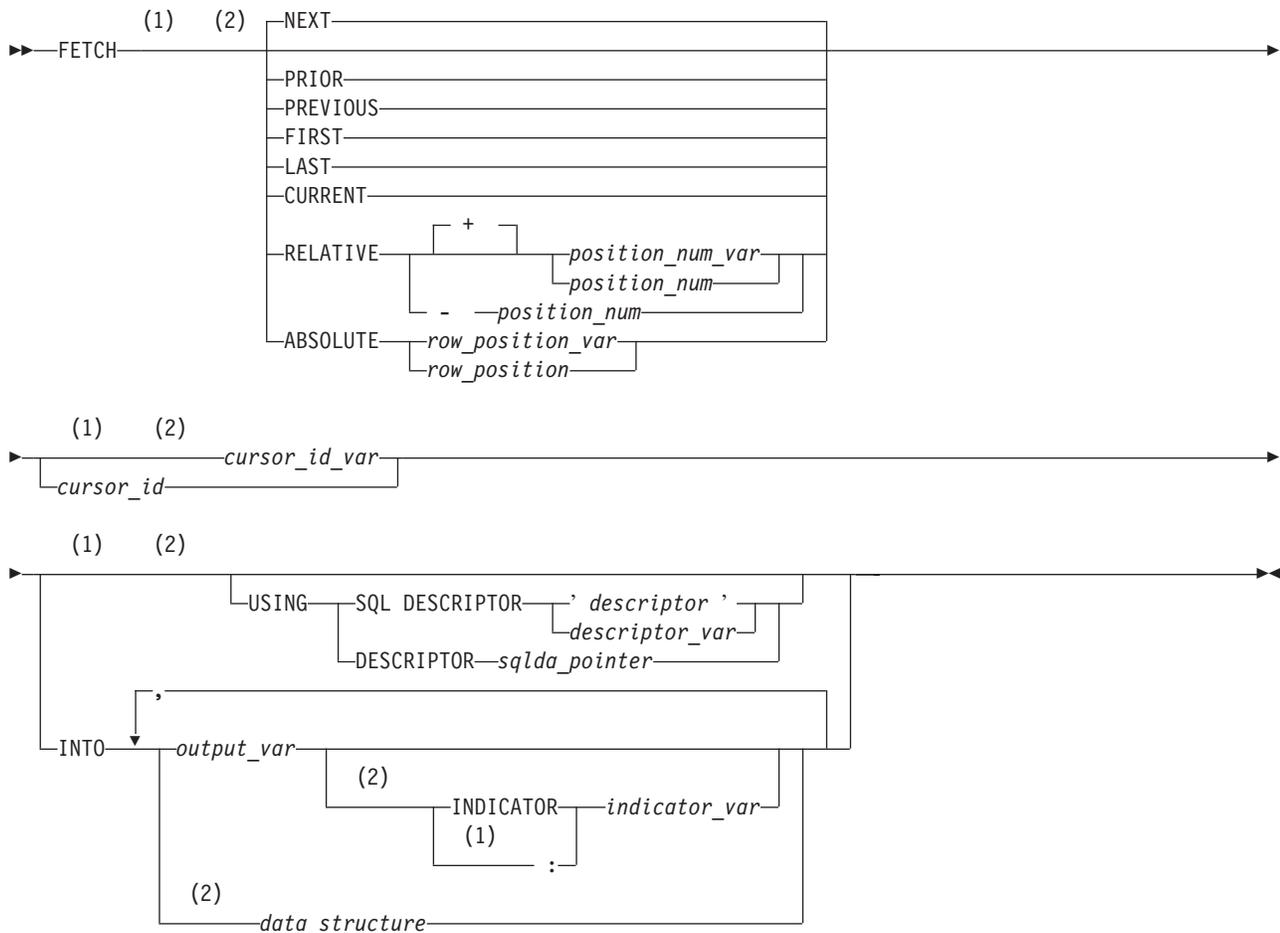
Related statements: "CREATE FUNCTION " on page 2-113, "CREATE PROCEDURE " on page 2-153, "CREATE TRIGGER" on page 2-239, "EXECUTE FUNCTION " on page 2-356, "GRANT " on page 2-395, "Receiving Input from the Called UDR" on page 3-4, "FOREACH " on page 3-22, and "LET" on page 3-32

FETCH

Use the FETCH statement to move a cursor to a new row in the active set and to retrieve the row values from memory.

Use this statement with Informix ESQL/C and with SPL.

Syntax



Notes:

- 1 Informix extension
- 2 ESQL/C only

Element	Description	Restrictions	Syntax
<i>cursor_id</i>	Cursor to retrieve rows	Must be open	"Identifier " on page 5-21
<i>cursor_id_var</i>	Host variable storing <i>cursor_id</i>	Must be character data type	Language specific
<i>data_structure</i>	Structure as a host variable	Must store fetched values	Language specific
<i>descriptor</i>	System-descriptor area	Must have been allocated	"Quoted String" on page 4-189
<i>descriptor_var</i>	Host variable storing <i>descriptor</i>	Must be allocated	Language specific
<i>indicator_var</i>	Host variable for return code if <i>output_var</i> can be NULL value	See "Using Indicator Variables" on page 2-371.	Language specific
<i>output_var</i>	Host variable for fetched value	Must store value from row	Language specific
<i>position_num</i>	Position relative to current row	Value 0 fetches current row	"Literal Number " on page 4-185
<i>position_num_var</i>	Host variable (= <i>position_num</i>)	Value 0 fetches current row	Language specific
<i>row_position</i>	Ordinal position in active set	Must be an integer >1	"Literal Number " on page 4-185
<i>row_position_var</i>	Host variable (= <i>row_position</i>)	Must be 1 or greater	Language specific

Element	Description	Restrictions	Syntax
<i>sqlda_pointer</i>	Pointer to an sqlda structure	Cannot begin with \$ nor :	See ESQL/C .

Usage

Except as noted, sections that follow describe how to use the **FETCH** statement in Informix ESQL/C routines. For information about the more restricted syntax and semantics of the **FETCH** statement in SPL routines, see "Fetching from Dynamic Cursors in SPL Routines" on page 2-376.

How the database server creates, stores, and fetches members of the active set of rows depends on whether the cursor was declared as a sequential cursor or as a scroll cursor. All cursors that the **FETCH** statement can reference in SPL routines are sequential cursors.

In X/Open mode, if a cursor-direction value (such as **NEXT** or **RELATIVE**) is specified, a warning message is issued, indicating that the statement does not conform to X/Open standards.

FETCH with a Sequential Cursor

A sequential cursor can fetch only the next row in sequence from the active set. The only option available is the default option, **NEXT**. A sequential cursor can read through a table only once each time the table is opened. The following Informix ESQL/C example illustrates the **FETCH** statement with a sequential cursor:

```
EXEC SQL FETCH seq_curs INTO :fname, :lname;
EXEC SQL FETCH NEXT seq_curs INTO :fname, :lname;
```

When the program opens a sequential cursor, the database server processes the query to the point of locating or constructing the first row of data. The goal of the database server is to tie up as few resources as possible.

Because the sequential cursor can retrieve only the next row, the database server can frequently create the active set one row at a time.

On each **FETCH** operation, the database server returns the contents of the current row and locates the next row. This one-row-at-a-time strategy is not possible if the database server must create the entire active set to determine which row is the first row (as would be the case if the **SELECT** statement included an **ORDER BY** clause).

FETCH with a Scroll Cursor

These Informix ESQL/C examples illustrate the **FETCH** statement with a scroll cursor:

```
EXEC SQL fetch previous q_curs into :orders;
EXEC SQL fetch last q_curs into :orders;
EXEC SQL fetch relative -10 q_curs into :orders;
printf("Which row? ");
scanf("
EXEC SQL fetch absolute :row_num q_curs into :orders;
```

A scroll cursor can fetch any row in the active set, either by specifying an absolute row position or a relative offset. Use the following cursor-position options to specify a particular row that you want to retrieve.

Keyword**Effect**

NEXT Retrieves next row in active set

PREVIOUS

Retrieves previous row in active set

PRIOR

Retrieves previous row in active set (Synonymous with PREVIOUS.)

FIRST Retrieves the first row in active set

LAST Retrieves the last row in active set

CURRENT

Retrieves the current row in active set (the same row as returned by the previous FETCH statement from the scroll cursor)

RELATIVE

Retrieves *n*th row, relative to the current cursor position in the active set, where *position_num* (or *position_num_var*) supplies *n*. A negative value indicates the *n*th row prior to the current cursor position. If *position_num* = 0, the current row is fetched.

ABSOLUTE

Retrieves *n*th row in active set, where *row_position_var* (or *row_position*) = *n*. Absolute row positions are numbered from 1.

Tip: Do not confuse row-position values with **rowid** values. A **rowid** value is based on the position of a row in its table and remains valid until the table is rebuilt. A row-position value (a value that the ABSOLUTE keyword retrieves) is the relative position of the row in the current active set of the cursor; the next time the cursor is opened, different rows might be selected.

How the Database Server Implements Scroll Cursors

Because it cannot anticipate which row the program will ask for next, the database server must retain all the rows in the active set until the scroll cursor closes. When a scroll cursor opens, the database server implements the active set as a temporary table, although it might not populate this table immediately.

The first time a row is fetched, the database server copies it into the temporary table as well as returning it to the program.

When a row is fetched for the second time, it can be taken from the temporary table. This scheme uses the fewest resources, in case the program abandons the query before it fetches all the rows. Rows that are never fetched are usually not copied from the database, or are saved in a temporary table.

Specifying Where Values Go in Memory

Each value from the select list of the query or the output of the executed user-defined function must be returned into a memory location. You can specify these destinations in one of the following ways:

- Use the INTO clause of a SELECT statement.
- Use the INTO clause of an EXECUTE Function (or EXECUTE PROCEDURE) statement.
- Use the INTO clause of a FETCH statement.
- Use a system-descriptor area.

- Use an `sqlda` structure.

Using the INTO Clause

If you associate a SELECT or EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement with a Function cursor, the statement can contain an INTO clause to specify variables to receive the returned values. You can use this method only when you write the SELECT, EXECUTE FUNCTION, or EXECUTE PROCEDURE statement as part of the cursor declaration; see “DECLARE ” on page 2-286. In this case, the FETCH statement cannot contain an INTO clause.

The following example uses the INTO clause of the SELECT statement to specify program variables in Informix ESQL/C:

```
EXEC SQL declare ord_date cursor for
      select order_num, order_date, po_num
         into :o_num, :o_date, :o_po;
EXEC SQL open ord_date;
EXEC SQL fetch next ord_date;
```

If you prepare a SELECT statement, the SELECT *cannot* include the INTO clause so you must use the INTO clause of the FETCH statement.

When you create a SELECT statement dynamically, you cannot use an INTO clause because you cannot name host variables in a prepared statement.

If you are certain of the number and data type of values in the projection list, you can use an INTO clause in the FETCH statement. If user input generated the query, however, you might not be certain of the number and data type of values that are being selected. In this case, you must use either a system descriptor or else a pointer to an `sqlda` structure.

Using Indicator Variables

Use an indicator variable if the returned data might be null.

The *indicator_var* parameter is optional, but use an indicator variable if the possibility exists that the value of *output_var* is NULL.

If you specify the indicator variable without the INDICATOR keyword, you cannot put a blank space between *output_var* and *indicator_var*.

For information about rules for placing a prefix before the *indicator_var*, see the *IBM Informix ESQL/C Programmer's Manual*.

The host variable cannot be a DATETIME or INTERVAL data type.

When the INTO Clause of FETCH is Required

When SELECT or EXECUTE FUNCTION (or EXECUTE PROCEDURE) omits the INTO clause, you must specify a data destination when a row is fetched.

For example, to dynamically execute a SELECT or EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement, the SELECT or EXECUTE FUNCTION (or EXECUTE PROCEDURE) cannot include its INTO clause in the PREPARE statement. Therefore, the FETCH statement must include an INTO clause to retrieve data into a set of variables. This method lets you store different rows in different memory locations.

You can fetch into a program-array element only by using an INTO clause in the FETCH statement. If you use a program array, you must list both the array name and a specific element of the array in *data_structure*. When you are declaring a cursor, do not refer to an array element within the SQL statement.

Tip: If you are certain of the number and data type of values in the select list of the Projection clause, you can use an INTO clause in the FETCH statement.

In the following Informix ESQL/C example, a series of complete rows is fetched into a program array. The INTO clause of each FETCH statement specifies an array element as well as the array name:

```
EXEC SQL BEGIN DECLARE SECTION;
    char wanted_state[2];
    short int row_count = 0;
    struct customer_t{
    {
        int    c_no;
        char   fname[15];
        char   lname[15];
    } cust_rec[100];
EXEC SQL END DECLARE SECTION;

main()
{
    EXEC SQL connect to 'stores_demo';
    printf("Enter 2-letter state code: ");
    scanf ("%s", wanted_state);
    EXEC SQL declare cust cursor for
        select * from customer where state = :wanted_state;
    EXEC SQL open cust;
    EXEC SQL fetch cust into :cust_rec[row_count];
    while (SQLCODE == 0)
    {
        printf("\n%s %s", cust_rec[row_count].fname,
            cust_rec[row_count].lname);
        row_count++;
        EXEC SQL fetch cust into :cust_rec[row_count];
    }
    printf ("\n");
    EXEC SQL close cust;
    EXEC SQL free cust;
}
```

Using a System-Descriptor Area (X/Open)

You can use a system-descriptor area to store output values when you do not know the number of return values or their data types that a SELECT or EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement returns at runtime. A system-descriptor area describes the data type and memory location of one or more return values, and conforms to the X/Open standards.

The keywords USING SQL DESCRIPTOR introduce the name of the system-descriptor area into which you fetch the contents of a row or the return values of a user-defined function. You can then use the GET DESCRIPTOR statement to transfer the values that the FETCH statement returns from the system-descriptor area into host variables.

The following example shows a valid FETCH...USING SQL DESCRIPTOR statement:

```

EXEC SQL allocate descriptor 'desc';
...
EXEC SQL declare selcurs cursor for
    select * from customer where state = 'CA';
EXEC SQL describe selcurs using sql descriptor 'desc';
EXEC SQL open selcurs;
while (1)
{
    EXEC SQL fetch selcurs using sql descriptor 'desc';
}

```

You can also use an **sqlda** structure to supply parameters dynamically.

Using sqlda Structures

You can use a pointer to an **sqlda** structure to store output values when you do not know the number of values or their data types that a SELECT or EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement returns.

This structure contains data descriptors that specify the data type and memory location for one selected value. The keywords USING DESCRIPTOR introduce the name of the pointer to the **sqlda** structure.

Tip: If you know the number and data types of all values in the select list, you can use an INTO clause in the FETCH statement. For more information, see “When the INTO Clause of FETCH is Required” on page 2-371.

To specify an sqlda structure as the location of parameters

1. Declare an **sqlda** pointer variable.
2. Use the DESCRIBE statement to fill in the **sqlda** structure.
3. Allocate memory to hold the data values.
4. Use the USING DESCRIPTOR clause of FETCH to specify the **sqlda** structure as the location into which you fetch the returned values.

The following example shows a FETCH USING DESCRIPTOR statement:

```

struct sqlda *sqlda_ptr;
...
EXEC SQL declare selcurs2 cursor for
    select * from customer where state = 'CA';
EXEC SQL describe selcurs2 into sqlda_ptr;
...
EXEC SQL open selcurs2;
while (1)
{
    EXEC SQL fetch selcurs2 using descriptor sqlda_ptr;
}
...

```

The **sqlld** value specifies the number of output values that are described in occurrences of the **sqlvar** structures of the **sqlda** structure. This number must correspond to the number of values returned from the prepared statement.

Fetching a Row for Update

The FETCH statement does not ordinarily lock a row that is fetched. Thus, another process can modify (update or delete) the fetched row immediately after your program receives it. A fetched row is locked in the following cases:

- When you set the isolation level to Repeatable Read, each row that you fetch is locked with a read lock until the cursor closes or until the current transaction ends. Other programs can also read the locked rows.

- When you set the isolation level to Cursor Stability, the current row is locked.
- In an ANSI-compliant database, an isolation level of Repeatable Read is the default; you can set it to something else.
- When you are fetching through an update cursor (one that is declared FOR UPDATE), each row you fetch is locked with a promotable lock. Other programs can read the locked row, but no other program can place a promotable or write lock; therefore, the row is unchanged if another user tries to modify it using the WHERE CURRENT OF clause of an UPDATE or DELETE statement.

When you modify a row, the lock is upgraded to a write lock and remains until the cursor is closed or the transaction ends. If you do not modify the row, the behavior of the database server depends on the isolation level you have set. The database server releases the lock on an unchanged row as soon as another row is fetched, unless you are using Repeatable Read isolation. (See “SET ISOLATION” on page 2-631.)

Important: You can hold locks on additional rows even when Repeatable Read isolation is not in use or is unavailable. Update the row with unchanged data to hold it locked while your program is reading other rows. You must evaluate the effect of this technique on performance in the context of your application, and you must be aware of the increased potential for deadlock.

When you use explicit transactions, be sure that a row is both fetched and modified within a single transaction; that is, both the FETCH statement and the subsequent UPDATE or DELETE statement must fall between a BEGIN WORK statement and the next COMMIT WORK statement.

Fetching from a Collection Cursor (IDS)

A Collection cursor allows you to access the individual elements of an Informix ESQL/C collection variable. To declare a Collection cursor, use the DECLARE statement and include the Collection-Derived Table segment in the SELECT statement that you associate with the cursor. After you open the collection cursor with the OPEN statement, the cursor allows you to access the elements of the collection variable.

To fetch elements, one at a time, from a Collection cursor, use the FETCH statement and the INTO clause. The FETCH statement identifies the Collection cursor that is associated with the collection variable. The INTO clause identifies the host variable that holds the element value that is fetched from the Collection cursor. The data type of the host variable in the INTO clause must match the element type of the collection.

Suppose you have a table called **children** with the following structure:

```
CREATE TABLE children
(
  age          SMALLINT,
  name         VARCHAR(30),
  fav_colors   SET(VARCHAR(20) NOT NULL),
)
```

The following Informix ESQL/C code fragment shows how to fetch elements from the **child_colors** collection variable:

```
EXEC SQL BEGIN DECLARE SECTION;
client collection child_colors;
varchar one_favorite[21];
char child_name[31] = "marybeth";
```

```

EXEC SQL END DECLARE SECTION;
EXEC SQL allocate collection :child_colors;
/* Get structure of fav_colors column for untyped
 * child_colors collection variable */
EXEC SQL select fav_colors into :child_colors
    from children
    where name = :child_name;
/* Declare select cursor for child_colors collection
 * variable */
EXEC SQL declare colors_curs cursor for
    select * from table(:child_colors);
EXEC SQL open colors_curs;
do
{
    EXEC SQL fetch colors_curs into :one_favorite;
    ...
} while (SQLCODE == 0);
EXEC SQL close colors_curs;
EXEC SQL free colors_curs;
EXEC SQL deallocate collection :child_colors;

```

After you fetch a collection element, you can modify the element with the UPDATE or DELETE statements. For more information, see the UPDATE and DELETE statements in this document. You can also insert new elements into the collection variable with an INSERT statement. For more information, see the INSERT statement.

Checking the Result of FETCH

You can use the **SQLSTATE** variable to check the result of each FETCH statement. The database server sets the **SQLSTATE** variable after each SQL statement. If a row is returned successfully, the **SQLSTATE** variable contains the value 00000. If no row is found, the database server sets the **SQLSTATE** code to 02000, which indicates no data found, and the current row is unchanged. The following conditions set the **SQLSTATE** code to 02000, indicating no data found:

- The active set contains no rows.
- You issue a FETCH NEXT statement when the cursor points to the last row in the active set or points past it.
- You issue a FETCH PRIOR or FETCH PREVIOUS statement when the cursor points to the first row in the active set.
- You issue a FETCH RELATIVE *n* statement when no *n*th row exists in the active set.
- You issue a FETCH ABSOLUTE *n* statement when no *n*th row exists in the active set.

The database server copies the **SQLSTATE** code from the **RETURNED_SQLSTATE** field of the system-diagnostics area. Client-server communication protocols of Dynamic Server, such as SQLI and DRDA[®], support **SQLSTATE** code values. For a list of these codes, and for information about how to get the message text, see “Using the SQLSTATE Error Status Code” on page 2-387. You can use the GET DIAGNOSTICS statement to examine the **RETURNED_SQLSTATE** field directly. The system-diagnostics area can also contain additional error information.

You can also use **SQLCODE** variable of the SQL Communications Area (**sqlca**) to determine the same results.

Fetching from Dynamic Cursors in SPL Routines

Use the FETCH statement in an SPL routine to retrieve the next row of the active set of a specified dynamic cursor into an ordered list of SPL variables that were declared in the same SPL routine.

Syntax

The syntax of the FETCH statement in SPL routines is a subset of the syntax that FETCH supports in Informix ESQL/C routines.

▶▶—FETCH— *cursor_id* —INTO— *output_var* ▶▶

Element	Description	Restrictions	Syntax
<i>cursor_id</i>	Name of a dynamic cursor	Must be open and must have been declared in the same SPL routine	"Identifier " on page 5-21
<i>output_var</i>	An SPL variable to store a fetched value from the row	Must have been declared locally or globally in the calling context, and must be of a data type compatible with the fetched column value	"Identifier " on page 5-21

Just as in ESQL/C routines, the list of output variables must correspond in number, order, and data type with column values that the SQL statement associated with the rows returned by the specified cursor.

All SPL cursors are sequential cursors. Your UDR must include logic to detect the end of the active set of the cursor, because the NOTFOUND condition does not automatically raise an exception in SPL.

The built-in **SQLCODE** function, which can only be called from SPL routines, can return the status code of a FETCH operation.

All other restrictions of ESQL/C on FETCH statements that reference sequential Select or Function cursors also apply to FETCH operations in SPL.

The FETCH statement in SPL routines does not support the following ESQL/C features:

- cursor names specified as host variables
- positional specifications or positional keywords (which require scroll cursors)
- the USING clause with descriptors or with **sqllda** pointers.

In the SPL language, indicator variables are not needed. If the FETCH operation retrieves a NULL value, the SPL variable that receives that fetched value is set to NULL.

The FETCH statement can reference only dynamic cursors that the DECLARE statement defined. The *cursor_id* cannot specify the name of a direct cursor that the FOREACH statement of SPL declared.

Related Information

Related statements: "ALLOCATE DESCRIPTOR " on page 2-2, "CLOSE " on page 2-77, "DEALLOCATE DESCRIPTOR " on page 2-284, "DECLARE " on page 2-286, "DESCRIBE" on page 2-311, "GET DESCRIPTOR " on page 2-381, "OPEN " on page 2-459, "PREPARE " on page 2-468, " SET DEFERRED_PREPARE" on page 2-603, and "SET DESCRIPTOR " on page 2-605

For a task-oriented discussion of the FETCH statement, see the *IBM Informix Guide to SQL: Tutorial*.

For more information about concepts that relate to the FETCH statement, see the *IBM Informix ESQL/C Programmer's Manual*.

FLUSH

Use the FLUSH statement to force rows that a PUT statement buffered to be written to the database.

Use this statement, which is an extension to the ANSI/ISO standard for SQL, with Informix ESQL/C.

Syntax

```
▶▶ FLUSH { cursor_id | cursor_id_var } ▶▶
```

Element	Description	Restrictions	Syntax
<i>cursor_id</i>	Name of a cursor	Must have been declared	"Identifier " on page 5-21
<i>cursor_id_var</i>	Host variable that holds the value of <i>cursor_id</i>	Must be a character data type	Language specific

Usage

The PUT statement adds a row to a buffer, whose content is written to the database when the buffer is full. Use the FLUSH statement to force the insertion when the buffer is not full.

If the program terminates without closing the cursor, the buffer is left unflushed. Rows placed into the buffer since the last flush are lost. Do not expect the end of the program to close the cursor and flush the buffer automatically. The following example shows a FLUSH statement that operates on a cursor called **icurs**:

```
FLUSH icurs
```

Error Checking FLUSH Statements

The SQL Communications Area (**sqlca**) structure contains information on the success of each FLUSH statement and the number of rows that are inserted successfully. The result of each FLUSH statement is described in the fields of the **sqlca**: **sqlca.sqlcode**, **SQLCODE**, and **sqlca.sqlerrd[2]**.

When you use data buffering with an Insert cursor, you do not discover errors until the buffer is flushed. For example, an input value that is incompatible with the data type of the column for which it is intended is discovered only when the buffer is flushed. When an error is discovered, any rows in the buffer that are located after the error are *not* inserted; they are lost from memory.

The **SQLCODE** field is set either to an error code or to zero (0) if no error occurs. The third element of the **SQLERRD** array is set to the number of rows that are successfully inserted into the database:

- If a block of rows is successfully inserted into the database, **SQLCODE** is set to zero (0) and **SQLERRD** to the count of rows.
- If an error occurs while the FLUSH statement is inserting a block of rows, **SQLCODE** shows which error, and **SQLERRD** contains the number of rows that were successfully inserted. (Uninserted rows are discarded from the buffer.)

Tip: When you encounter an **SQLCODE** error, a corresponding **SQLSTATE** error value also exists. Client-server communication protocols of Dynamic Server, such as SQLI and DRDA, support **SQLSTATE** code values. For a list of these codes, and for information about how to get the message text, see “Using the SQLSTATE Error Status Code” on page 2-387.

To count the number of rows actually inserted into the database as well as the number not yet inserted

1. Prepare two integer variables, for example, **total** and **pending**.
2. When the cursor opens, set both variables to 0.
3. Each time a PUT statement executes, increment both **total** and **pending**.
4. Whenever a FLUSH statement executes or the cursor is closed, subtract the third field of the **SQLERRD** array from **pending**.

Examples

The following example assumes that a function named *next_cust* returns either information about a new customer or null data to signal the end of input:

```
EXEC SQL BEGIN WORK;
EXEC SQL OPEN new_custs;

while(SQLCODE == 0)
{
  next_cust();
  if(the_company == NULL)
    break;

  EXEC SQL PUT new_custs;
}

if(SQLCODE == 0) /* if no problem with PUT */
{
  EXEC SQL FLUSH new_custs;
  /* write any rows left */

  if(SQLCODE == 0) /* if no problem with FLUSH */
    EXEC SQL COMMIT WORK; /* commit changes */
}
else
  EXEC SQL ROLLBACK WORK; /* else undo changes */
```

The code in this example calls *next_cust* repeatedly. When it returns non-null data, the **PUT** statement sends the returned data to the row buffer. When the buffer fills,

the rows it contains are automatically sent to the database server. The loop normally ends when *next_cust* has no more data to return.

Related Information

Related statements: "CLOSE " on page 2-77, "DECLARE " on page 2-286, "OPEN " on page 2-459, and "PREPARE " on page 2-468

For a task-oriented discussion of FLUSH, see the *IBM Informix Guide to SQL: Tutorial*.

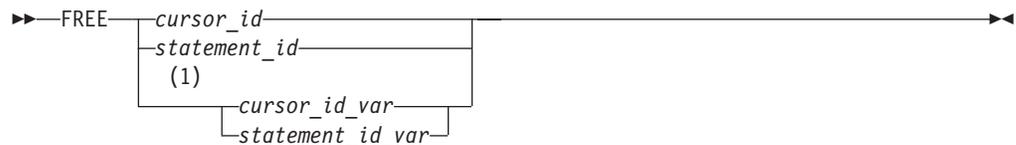
For information about the **sqlca** structure, see the *IBM Informix ESQL/C Programmer's Manual*.

FREE

Use the FREE statement to release resources that are allocated to a prepared statement or to a cursor.

Use this statement, which is an extension to the ANSI/ISO standard for SQL, with Informix ESQL/C or with SPL.

Syntax



Notes:

- 1 ESQL/C only

Element	Description	Restrictions	Syntax
<i>cursor_id</i>	Name of a cursor	Must have been declared	"Identifier " on page 5-21
<i>cursor_id_var</i>	Host variable that holds the value of <i>cursor_id</i>	Must be a character data type	Language specific
<i>statement_id</i>	Identifier of a prepared SQL statement	Must be defined in a previous PREPARE statement	"PREPARE " on page 2-468
<i>statement_id_var</i>	Host variable that stores the name of a prepared object	Must be declared as a character data type.	"PREPARE " on page 2-468

Usage

FREE releases the resources that the database server and (for ESQL/C) the application-development tool allocated for a prepared statement or for a cursor.

If you declared a cursor for a prepared statement, FREE *statement_id* (or *statement_id_var*) releases only the resources in the application development tool; the cursor can still be used. The resources in the database server are released only when you free the cursor.

If you prepared a statement (but did not declare a cursor for it), `FREE statement_id` (or `FREE statement_id_var`) releases the resources in both the application development tool and the database server.

After you free a statement, you cannot execute it or declare a cursor for it until you prepare it again.

The following Informix ESQL/C example shows the sequence of statements that is used to free an implicitly prepared statement:

```
EXEC SQL prepare sel_stmt from 'select * from orders';  
...  
EXEC SQL free sel_stmt;
```

The following Informix ESQL/C example shows the sequence of statements that are used to release the resources of an explicitly prepared statement. The first `FREE` statement in this example frees the cursor. The second `FREE` statement in this example frees the prepared statement.

```
sprintf(demoselect, "%s %s",  
       "select * from customer ",  
       "where customer_num between 100 and 200");  
EXEC SQL prepare sel_stmt from :demoselect;  
EXEC SQL declare sel_curs cursor for sel_stmt;  
EXEC SQL open sel_curs;  
...  
EXEC SQL close sel_curs;  
EXEC SQL free sel_curs;  
EXEC SQL free sel_stmt;
```

If you declared a cursor for a prepared statement, freeing the cursor releases only the resources in the database server. To release the resources for the statement in the application-development tool, use `FREE statement_id` (or `FREE statement_id_var`). If a cursor is not declared for a prepared statement, freeing it releases the resources in both the application-development tool and the database server. For an ESQL/C example of a `FREE` statement that frees a cursor, see the previous example.

After a cursor is freed, it cannot be opened until it is declared again. The cursor should be explicitly closed before it is freed.

When an SPL routine completes execution, the database server automatically releases any resources that had been allocated to the cursor or to prepared statements by `PREPARE` or `DECLARE` statements in the routine, if these have not already been released by the `FREE` statement.

The `FREE` statement in SPL routines cannot reference the `cursor_id` of a direct cursor that the `FOREACH` statement of SPL can declare.

Related Information

Related statements: “`CLOSE` ” on page 2-77, “`DECLARE` ” on page 2-286, “`EXECUTE` ” on page 2-348, “`EXECUTE IMMEDIATE`” on page 2-360, “`OPEN` ” on page 2-459, “`PREPARE` ” on page 2-468, and “ `SET AUTOFREE` ” on page 2-582

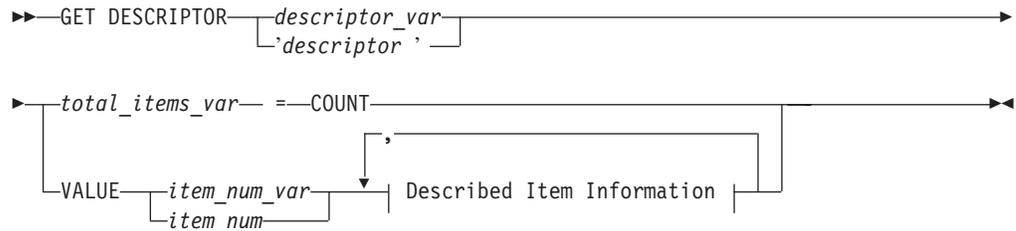
For a task-oriented discussion of the `FREE` statement, see the *IBM Informix Guide to SQL: Tutorial*.

GET DESCRIPTOR

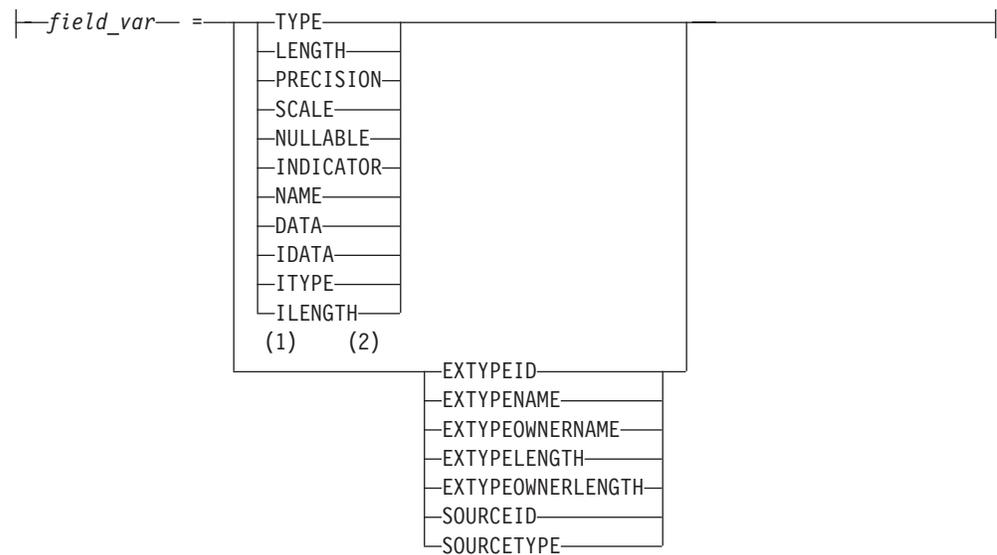
Use the GET DESCRIPTOR statement to read from a system descriptor area.

Use this statement with Informix ESQL/C.

Syntax



Described Item Information:



Notes:

- 1 Informix extension
- 2 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>descriptor</i>	Quoted string that identifies a system-descriptor area (SDA)	System-descriptor area must already have been allocated	"Quoted String" on page 4-189
<i>descriptor_var</i>	Variable that stores <i>descriptor</i> value	Same restrictions as <i>descriptor</i>	Language specific
<i>field_var</i>	Host variable to receive the contents of a field from an SDA	Must be of type that can receive value of a specified SDA field	Language specific
<i>item_num</i>	Unsigned ordinal number of an item described in the SDA	$0 \leq \text{item_num} \leq (\text{number of item descriptors in the SDA})$	"Literal Number" on page 4-185
<i>item_num_var</i>	Host variable storing <i>item_num</i>	Must be an integer data type	Language specific

Element	Description	Restrictions	Syntax
<i>total_items_var</i>	Host variable storing the number of items described in the SDA	Must be an integer data type	Language specific

Usage

Use the GET DESCRIPTOR statement to accomplish any of the following tasks:

- Determine how many items are described in a system-descriptor area.
- Determine the characteristics of each column or expression that is described in the system-descriptor area.
- Copy a value from the system-descriptor area into a host variable after a FETCH statement.

With Dynamic Server, you can use GET DESCRIPTOR after you describe EXECUTE FUNCTION, INSERT, SELECT, or UPDATE statements with the DESCRIBE ... USING SQL DESCRIPTOR statement.

With Extended Parallel Server, you can use GET DESCRIPTOR after you describe EXECUTE PROCEDURE, INSERT, or SELECT statements with the DESCRIBE ... USING SQL DESCRIPTOR statement.

The host variables that you reference in the GET DESCRIPTOR statement must be declared in the BEGIN DECLARE SECTION of a program.

If an error occurs during the assignment of a value to any specified host variable, the contents of the host variable are undefined.

Using the COUNT Keyword

Use the COUNT keyword to determine how many items are described in the system-descriptor area.

The following Informix ESQL/C example shows how to use a GET DESCRIPTOR statement with a host variable to determine how many items are described in the system-descriptor area called **desc1**:

```
main()
{
EXEC SQL BEGIN DECLARE SECTION;
int h_count;
EXEC SQL END DECLARE SECTION;

EXEC SQL allocate descriptor 'desc1' with max 20;

/* This section of program would prepare a SELECT or INSERT
 * statement into the s_id statement id.
 */
EXEC SQL describe s_id using sql descriptor 'desc1';

EXEC SQL get descriptor 'desc1' :h_count = count;
...
}
```

Using the VALUE Clause

Use the VALUE clause to obtain information about a described column or expression or to retrieve values that the database server returns in a system descriptor area.

The *item_num* must be greater than zero (0) but not greater than the number of item descriptors that were specified when the system-descriptor area was allocated with the ALLOCATE DESCRIPTOR statement.

Using the VALUE Clause After a DESCRIBE

After you describe a SELECT, EXECUTE FUNCTION (or EXECUTE PROCEDURE), INSERT, or UPDATE statement, the characteristics of each column or expression in the select list of the SELECT statement, the characteristics of the values returned by the EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement, or the characteristics of each column in an INSERT or UPDATE statement are returned to the system-descriptor area. Each value in the system-descriptor area describes the characteristics of one returned column or expression.

The following Informix ESQL/C example uses GET DESCRIPTOR to obtain data type information from the **demodesc** system-descriptor area:

```
EXEC SQL get descriptor 'demodesc' value :index
        :type = TYPE,
        :len = LENGTH,
        :name = NAME;
printf("Column %d: type = %d, len = %d, name = %s\n",
        index, type, len, name);
```

The value that the database server returns into the TYPE field is a defined integer. To evaluate the data type that is returned, test for a specific integer value. For additional information about integer data type values, see “Setting the TYPE or ITYPE Field” on page 2-608.

In X/Open mode, the X/Open code is returned to the TYPE field. You cannot mix the two modes because errors can result. For example, if a particular data type is not defined under X/Open mode but is defined for IBM Informix products, executing a GET DESCRIPTOR statement can result in an error.

In X/Open mode, a warning message appears if ILENGTH, IDATA, or ITYPE is used. It indicates that these fields are not standard X/Open fields for a system-descriptor area.

If the TYPE of a fetched value is DECIMAL or MONEY, the database server returns the precision and scale information for a column into the PRECISION and SCALE fields after a DESCRIBE statement is executed. If the TYPE is *not* DECIMAL or MONEY, the SCALE and PRECISION fields are undefined.

Using the VALUE Clause After a FETCH

Each time your program fetches a row, it must copy the fetched value into host variables so that the data can be used. To accomplish this task, use a GET DESCRIPTOR statement after each fetch of each value in the select list. If three values exist in the select list, you need to use three GET DESCRIPTOR statements after each fetch (assuming you want to read all three values). The item numbers for each of the three GET DESCRIPTOR statements are 1, 2, and 3.

The following Informix ESQL/C example shows how you can copy data from the DATA field into a host variable (**result**) after a fetch. For this example, it is predetermined that all returned values are the same data type:

```
EXEC SQL get descriptor 'demodesc' :desc_count = count;
..
EXEC SQL fetch democursor using sql descriptor 'demodesc';
for (i = 1; i <= desc_count; i++)
{
```

```

    if (sqlca.sqlcode != 0) break;
    EXEC SQL get descriptor 'demodesc' value :i :result = DATA;
    printf("%s ", result);
    }
printf("\n");

```

Fetching a NULL Value

When you use GET DESCRIPTOR after a fetch, and the fetched value is NULL, the INDICATOR field is set to -1 to indicate the NULL value. The value of DATA is undefined if INDICATOR indicates a NULL value. The host variable into which DATA is copied has an unpredictable value.

Using LENGTH or ILENGTH

If your DATA or IDATA field contains a character string, you must specify a value for LENGTH. If you specify LENGTH=0, LENGTH is automatically set to the maximum length of the string. The DATA or IDATA field might contain a literal character string or a character string that is derived from a character variable of CHAR or VARCHAR data type. This provides a method to determine the length of a string in the DATA or IDATA field dynamically.

If a DESCRIBE statement precedes a GET DESCRIPTOR statement, LENGTH is automatically set to the maximum length of the character field that is specified in your table.

This information is identical for ILENGTH. Use ILENGTH when you create a dynamic program that does not comply with the X/Open standard.

Describing an Opaque-Type Column (IDS)

The DESCRIBE statement sets the following item-descriptor fields when the column to fetch has an opaque type as its data type:

- The EXTTYPEID field stores the extended ID for the opaque type. This integer is the value in the corresponding **extended_id** column of the **sysxdtypes** system catalog table.
- The EXTYPENAME field stores the name of the opaque type. This character value is the value in the **name** column of the row with the matching **extended_id** value in the **sysxdtypes** system catalog table.
- The EXTYPELENGTH field stores the length of the opaque-type name. This integer is the length of the data type name (in bytes).
- The EXTTYPEOWNERNAME field stores the name of the opaque-type owner. This character value is the value in the **owner** column of the row with the matching **extended_id** value in the **sysxdtypes** system catalog table.
- The EXTTYPEOWNERLENGTH field stores the length of the value in the EXTTYPEOWNERNAME field. This integer is the length, in bytes, of the name of the owner of the opaque type.

Use these field names with the GET DESCRIPTOR statement to obtain information about an opaque column.

Describing a Distinct-Type Column (IDS)

The DESCRIBE statement sets the following item-descriptor fields when the column to fetch has a distinct type as its data type:

- The SOURCEID field stores the extended identifier for the source data type.

This integer has the value of the **source** column for the row of the **sysxdtypes** system catalog table whose **extended_id** value matches that of the distinct data type. This field is set only if the source data type is an opaque data type.

- The **SOURCETYPE** field stores the data type constant for the source data type. This value is the data type constant (from the **sqltypes.h** file) for the data type of the source type for the **DISTINCT** data type. The codes for the **SOURCETYPE** field are listed in the description of the **TYPE** field in the **SET DESCRIPTOR** statement. (For more information, see “Setting the **TYPE** or **ITYPE** Field” on page 2-608.) This integer value must correspond to the value in the **type** column for the row of the **sysxdtypes** system catalog table whose **extended_id** value matches that of the **DISTINCT** data type.

Use these field names with the **GET DESCRIPTOR** statement to obtain information about a distinct-type column.

Examples

You use the **GET DESCRIPTOR** statement to accomplish any of the following tasks:

- To determine how many items are described in a system-descriptor area.
- To determine the characteristics of each column or expression that is described in the system-descriptor area.
- To copy a value from the system-descriptor area into a host variable after a **FETCH** statement.

The following **ESQL/C** example shows how to use a **GET DESCRIPTOR** statement with a host variable to determine how many items are described in the system-descriptor area called **desc1**: **GET DESCRIPTOR**

```
main()
{
    EXEC SQL BEGIN DECLARE SECTION;
        int h_count;
    EXEC SQL END DECLARE SECTION;

    EXEC SQL allocate descriptor 'desc1' with max 20;

    /* This section of program would prepare a SELECT or INSERT
    * statement into the s_id statement id.
    */

    EXEC SQL describe s_id using sql descriptor 'desc1';
    EXEC SQL get descriptor 'desc1' :h_count = count;
```

The following **ESQL/C** example uses **GET DESCRIPTOR** to obtain data type information from the **demodesc** system-descriptor area:

```
EXEC SQL get descriptor 'demodesc' value
    :index :type = TYPE,
    :len = LENGTH,
    :name = NAME;
printf("Column %d: type = %d, len = %d, name = %s\n",
    index, type, len, name);
```

The following **ESQL/C** example shows how you can copy data from the **DATA** field into a host variable (**result**) after a **fetch**. For this example, it is predetermined that all returned values are the same data type:

```
EXEC SQL get descriptor 'demodesc' :desc_count = count;
...
EXEC SQL fetch democursor using sql descriptor 'demodesc';
for (i = 1; i <= desc_count; i++)
```


Using the SQLSTATE Error Status Code

When an SQL statement executes, an error status code is automatically generated. This code represents success, failure, warning, or no data found. This error status code is stored in a built-in variable called **SQLSTATE**.

Class and Subclass Codes

The **SQLSTATE** status code is a five-character string that can contain only digits and uppercase letters.

The first two characters of the **SQLSTATE** status code indicate a class. The last three characters of the **SQLSTATE** code indicate a subclass. Figure 2-1 shows the structure of the **SQLSTATE** code. This example uses the value 08001, where 08 is the class code and 001 is the subclass code. The value 08001 represents the error unable to connect with database environment.

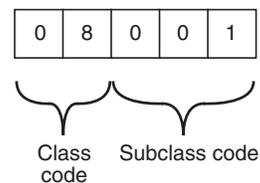


Figure 2-1. Structure of the SQLSTATE Code

The following table is a quick reference for interpreting class code values.

SQLSTATE Class Code Value

Outcome

00	Success
01	Success with warning
02	No data found
> 02	Error or warning

SQLSTATE Support for the ANSI/ISO Standard for SQL

All status codes returned to the **SQLSTATE** variable are ANSI-compliant except in the following cases:

- **SQLSTATE** codes with a class code of 01 and a subclass code that begins with an I are Informix-specific warning messages.
- **SQLSTATE** codes with a class code of IX and any subclass code are Informix-specific error messages.
- **SQLSTATE** codes whose class code begins with a digit in the range 5 to 9 or with an uppercase letter in the range I to Z indicate conditions that are currently undefined by the ANSI/ISO standard for SQL. The only exception is that **SQLSTATE** codes whose class code is IX are Informix-specific error messages.

Client-server communication protocols of Dynamic Server, such as SQLI and DRDA, support these **SQLSTATE** code values.

List of SQLSTATE Codes

This table describes the class codes, subclass codes, and the meaning of all valid warning and error codes associated with the **SQLSTATE** variable.

Class	Subclass	Meaning
00	000	Success.
01	000	Success with warning.
	002	Disconnect error. Transaction rolled back.
	003	NULL value eliminated in set function.
	004	String data, right truncation.
	005	Insufficient item descriptor areas.
	006	Privilege not revoked.
	007	Privilege not granted.
	I01	Database has transactions.
	I03	ANSI-compliant database selected.
	I04	IBM Informix database server selected.
	I05	Float to decimal conversion was used.
	I06	Informix extension to ANSI-compliant syntax.
	I07	UPDATE or DELETE statement does not have a WHERE clause.
	I08	An ANSI keyword was used as a cursor name.
	I09	Cardinalities of the projection list and of the INTO list are not equal.
	I10	Database server running in secondary mode.
	I11	Dataskip is turned on.
02	000	No data found.
07	000	Dynamic SQL error.
	001	USING clause does not match dynamic parameters.
	002	USING clause does not match target specifications.
	003	Cursor specification cannot be executed.
	004	USING clause is required for dynamic parameters.
	005	Prepared statement is not a cursor specification.
	006	Restricted data type attribute violation.
	008	Invalid descriptor count.
	009	Invalid descriptor index.
08	000	Connection exception.
	001	Database server rejected the connection.
	002	Connection name in use.
	003	Connection does not exist.
	004	Client unable to establish connection.
	006	Transaction rolled back.
	007	Transaction state unknown.
	S01	Communication failure.
0A	000	Feature not supported.
	001	Multiple server transactions.
21	000	Cardinality violation.
	S01	Insert value list does not match column list.
	S02	Degree of derived table does not match column list.

Class	Subclass	Meaning
22	000	Data exception.
	001	String data, right truncation.
	002	NULL value, no indicator parameter.
	003	Numeric value out of range.
	005	Error in assignment.
	027	Data exception trim error.
	012	Division by zero (0).
	019	Invalid escape character.
	024	Unterminated string.
	025	Invalid escape sequence.
23	000	Integrity constraint violation.
24	000	Invalid cursor state.
25	000	Invalid transaction state.
2B	000	Dependent privilege descriptors still exist.
2D	000	Invalid transaction termination.
26	000	Invalid SQL statement identifier.
2E	000	Invalid connection name.
28	000	Invalid user-authorization specification.
33	000	Invalid SQL descriptor name.
34	000	Invalid cursor name.
35	000	Invalid exception number.
37	000	Syntax error or access violation in PREPARE or EXECUTE IMMEDIATE.
3C	000	Duplicate cursor name.
40	000	Transaction rollback.
	003	Statement completion unknown.
42	000	Syntax error or access violation.
S0	000	Invalid name.
	001	Base table or view table already exists.
	002	Base table not found.
	011	Index already exists.
	021	Column already exists.
S1	001	Memory allocation failure.
IX	000	Informix reserved error message.

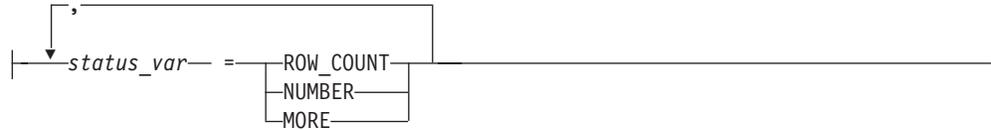
Using SQLSTATE in Applications

You can use a built-in variable, called **SQLSTATE**, which you do not need to declare in your program. **SQLSTATE** contains the status code, essential for error handling, which is generated every time your program executes an SQL statement. **SQLSTATE** is created automatically. You can examine the **SQLSTATE** variable to determine whether an SQL statement was successful. If the **SQLSTATE** variable indicates that the statement failed, you can execute a **GET DIAGNOSTICS** statement to obtain additional error information.

For an example of how to use an **SQLSTATE** variable in a program, see “Using GET DIAGNOSTICS for Error Checking” on page 2-394.

Statement Clause

Statement Clause:



Element	Description	Restrictions	Syntax
<i>status_var</i>	Host variable to receive status information about the most recent SQL statement for the specified status field name	Must match data type of the field	Language specific

When retrieving count and overflow information, GET DIAGNOSTICS can deposit the values of the three statement fields into a corresponding host variable. The host-variable data type must be the same as that of the requested field. The following keywords represent these three fields.

Field Name Keyword	Field Data Type	Field Contents	ESQL/C Host Variable Data Type
MORE	Character	Y or N	char[2]
NUMBER	Integer	1 to 35,000	int
ROW_COUNT	Integer	0 to 999,999,999	int

Using the MORE Keyword

Use the MORE keyword to determine if the most recently executed SQL statement resulted in the following actions by the database server:

- Stored all the exceptions that it detected in the diagnostics area
If so, GET DIAGNOSTICS returns a value of N.
- Detected more exceptions than it stored in the diagnostics area
If so, GET DIAGNOSTICS returns a value of Y. (The value of MORE is always N.)

Using the ROW_COUNT Keyword

The ROW_COUNT keyword returns the number of rows the most recently executed DML statement processed. ROW_COUNT counts these rows:

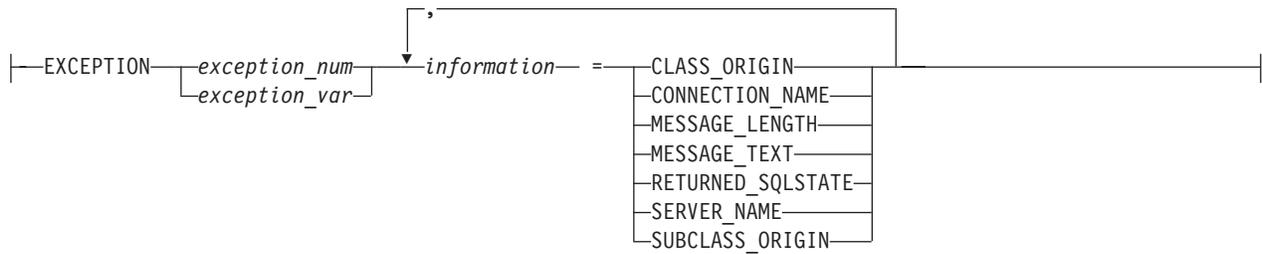
- Inserted into a table
- Updated in a table
- Deleted from a table

Using the NUMBER Keyword

The NUMBER keyword returns the number of exceptions that the most recently executed SQL statement raised. The NUMBER field can hold a value from 1 to 35,000, depending on how many exceptions are counted.

EXCEPTION Clause

Exception Clause:



Element	Description	Restrictions	Syntax
<i>exception_num</i>	Number of exceptions	Integer in range 1 to 35,000	"Literal Number " on page 4-185
<i>exception_var</i>	Variable storing <i>exception_num</i>	Must be SMALLINT or INT	Language specific
<i>information</i>	Host variable to receive the value of a specified exception field	Data type must match that of the specified field	Language specific

The *exception_num* literal indicates one of the exception values from the number of exceptions that the NUMBER field in the Statement clause returns.

When retrieving exception information, GET DIAGNOSTICS writes the values of each of the seven fields into corresponding host variables. These fields are located in the diagnostics area and are derived from an exception raised by the most recent SQL statement.

The host-variable data type must be the same as that of the requested field. The following table describes the seven exception information fields.

Field Name Keyword	Field Data Type	Field Contents	ESQL/C Host Variable Data Type
RETURNED_SQLSTATE	Character	SQLSTATE value	char[6]
CLASS_ORIGIN	Character	String	char[255]
SUBCLASS_ORIGIN	Character	String	char[255]
MESSAGE_TEXT	Character	String	char[255]
MESSAGE_LENGTH	Integer	Numeric value	int
SERVER_NAME	Character	String	char[255]
CONNECTION_NAME	Character	String	char[255]

The application specifies the exception by number, using either an unsigned integer or an integer host variable (an exact numeric with a scale of 0). An exception with a value of 1 corresponds to the **SQLSTATE** value set by the most recent SQL statement other than GET DIAGNOSTICS. The association between other exception numbers and other exceptions raised by that SQL statement is undefined. Thus, no set order exists in which the diagnostic area can be filled with exception values. You always get at least one exception, even if the **SQLSTATE** value indicates success.

If an error occurs within the GET DIAGNOSTICS statement (that is, if an invalid exception number is requested), the Informix internal **SQLCODE** and **SQLSTATE** variables are set to the value of that exception. In addition, the GET DIAGNOSTICS fields are undefined.

Using the RETURNED_SQLSTATE Keyword

The RETURNED_SQLSTATE keyword returns the **SQLSTATE** value that describes the exception.

Using the CLASS_ORIGIN Keyword

Use the CLASS_ORIGIN keyword to retrieve the class portion of the RETURNED_SQLSTATE value. If the ISO standard for SQL defines the class, the value of CLASS_ORIGIN is equal to ISO 9075. Otherwise, the value returned by CLASS_ORIGIN is defined by Informix and cannot be ISO 9075. The terms ANSI SQL and ISO SQL are synonymous.

Using the SUBCLASS_ORIGIN Keyword

The SUBCLASS_ORIGIN keyword returns data on the RETURNED_SQLSTATE subclass. (This value is ISO 9075 if the ISO standard defines the subclass.)

Using the MESSAGE_TEXT Keyword

The MESSAGE_TEXT keyword returns the message text of the exception (for example, an error message).

Using the MESSAGE_LENGTH Keyword

The MESSAGE_LENGTH keyword returns the length in bytes of the current message text string.

Using the SERVER_NAME Keyword

The SERVER_NAME keyword returns the name of the database server associated with a CONNECT or DATABASE statement. GET DIAGNOSTICS updates the SERVER_NAME field when any of the following events occur:

- A CONNECT statement successfully executes.
- A SET CONNECTION statement successfully executes.
- A DISCONNECT statement successfully terminates the current connection.
- A DISCONNECT ALL statement fails.

The SERVER_NAME field is not updated, however, after these events:

- A CONNECT statement fails.
- A DISCONNECT statement fails (but this does not include the DISCONNECT ALL statement).
- A SET CONNECTION statement fails.

The SERVER_NAME field retains the value set in the previous SQL statement. If any of the preceding conditions occur on the first SQL statement that executes, the SERVER_NAME field is blank.

The Contents of the SERVER_NAME Field

The SERVER_NAME field contains different information after you execute the following statements.

Executed Statement

SERVER_NAME Field Contents

CONNECT

Contains the name of the database server to which you connect or fail to connect. Field is blank if you do not have a current connection or if you make a default connection.

SET CONNECTION

Contains the name of the database server to which you switch or fail to switch.

DISCONNECT

Contains the name of the database server from which you disconnect or fail to disconnect. If you disconnect and then you execute a DISCONNECT statement for a connection that is not current, the SERVER_NAME field remains unchanged.

DISCONNECT ALL

Sets the field to blank if the statement executes successfully. If the statement fails, SERVER_NAME contains the names of all the database servers from which you did not disconnect. (This information does not mean that the connection still exists.)

If CONNECT succeeds, SERVER_NAME is set to one of the following values:

- The **INFORMIXSERVER** value (if the connection is to a default database server; that is, if CONNECT specified no database server)
- The name of the database server (if the connection is to a specific database server)

The DATABASE Statement

When you execute a DATABASE statement, the SERVER_NAME field contains the name of the database server on which the database resides.

Using the CONNECTION_NAME Keyword

Use the CONNECTION_NAME keyword to return the name of the connection specified in your CONNECT or SET CONNECTION statement.

When the CONNECTION_NAME Keyword Is Updated

GET DIAGNOSTICS updates the CONNECTION_NAME field when the following situations occur:

- A CONNECT statement successfully executes.
- A SET CONNECTION statement successfully executes.
- A DISCONNECT statement successfully executes at the current connection.
GET DIAGNOSTICS fills the CONNECTION_NAME field with blanks because no current connection exists.
- A DISCONNECT ALL statement fails.

When the CONNECTION_NAME Is Not Updated

The CONNECTION_NAME field is not updated in the following cases:

- A CONNECT statement fails.

- A DISCONNECT statement fails (but this does not include the DISCONNECT ALL statement).
- A SET CONNECTION statement fails.

The CONNECTION_NAME field retains the value set in the previous SQL statement. If any of the preceding conditions occurs on the first SQL statement that executes, the CONNECTION_NAME field is blank.

An implicit connection has no name. After a DATABASE statement successfully creates an implicit connection, the CONNECTION_NAME field is blank.

The Contents of the CONNECTION_NAME Field

The CONNECTION_NAME field contains different information after you execute the following statements.

Executed Statement

CONNECTION_NAME Field Contents

CONNECT

Contains connection name specified in the CONNECT statement, to which you connect or fail to connect. The field is blank for no current connection or a default connection.

SET CONNECTION

Contains the connection name specified in the CONNECT statement, to which you switch or fail to switch.

DISCONNECT

Contains the connection name specified in the CONNECT statement, from which you disconnect or fail to disconnect. If you disconnect, and then execute a DISCONNECT statement for a connection that is not current, the CONNECTION_NAME field remains unchanged.

DISCONNECT ALL

Contains no information if the statement executes successfully. If the statement does not execute successfully, the CONNECTION_NAME field contains the names of all the connections specified in your CONNECT statement from which you did not disconnect. This information does not mean, however, that the connection still exists.

If CONNECT is successful, CONNECTION_NAME takes one of these values:

- The name of the database environment as specified in the CONNECT statement if the CONNECT statement does not include the AS clause
- The name of the connection (the identifier that was declared after the AS keyword) if the CONNECT statement includes the AS clause

Using GET DIAGNOSTICS for Error Checking

GET DIAGNOSTICS returns values from various fields in the diagnostics area. For each field in the diagnostics area that you wish to access, you must supply a host variable of a compatible data type.

The following example illustrates how to use the GET DIAGNOSTICS statement to display error information. The example shows an Informix ESQL/C error display routine called **disp_sqlstate_err()**:

```
void disp_sqlstate_err()
{
  int j;
```

```

EXEC SQL BEGIN DECLARE SECTION;
    int exception_count;
    char overflow[2];
    int exception_num=1;
    char class_id[255];
    char subclass_id[255];
    char message[255];
    int messlen;
    char sqlstate_code[6];
    int i;
EXEC SQL END DECLARE SECTION;
printf("-----");
printf("-----\n");
printf("SQLSTATE:
printf("SQLCODE: %d\n", SQLCODE);
printf("\n");
EXEC SQL get diagnostics :exception_count = NUMBER,
    :overflow = MORE;
printf("EXCEPTIONS: Number=%d\t", exception_count);
printf("More? %s\n", overflow);
for (i = 1; i <= exception_count; i++)
{
    EXEC SQL get diagnostics exception :i
        :sqlstate_code = RETURNED_SQLSTATE,
        :class_id = CLASS_ORIGIN, :subclass_id = SUBCLASS_ORIGIN,
        :message = MESSAGE_TEXT, :messlen = MESSAGE_LENGTH;
    printf("- - - - -\n");
    printf("EXCEPTION %d: SQLSTATE=%s\n", i, sqlstate_code);
    message[messlen-1] = '\0';
    printf("MESSAGE TEXT: %s\n", message);
    j = stleng(class_id);
    while((class_id[j] == '\0') ||
        (class_id[j] == ' '))
        j--;
    class_id[j+1] = '\0';
    printf("CLASS ORIGIN:
    j = stleng(subclass_id);
    while((subclass_id[j] == '\0') ||
        (subclass_id[j] == ' '))
        j--;
    subclass_id[j+1] = '\0';
    printf("SUBCLASS ORIGIN:
}
printf("-----");
printf("-----\n");
}

```

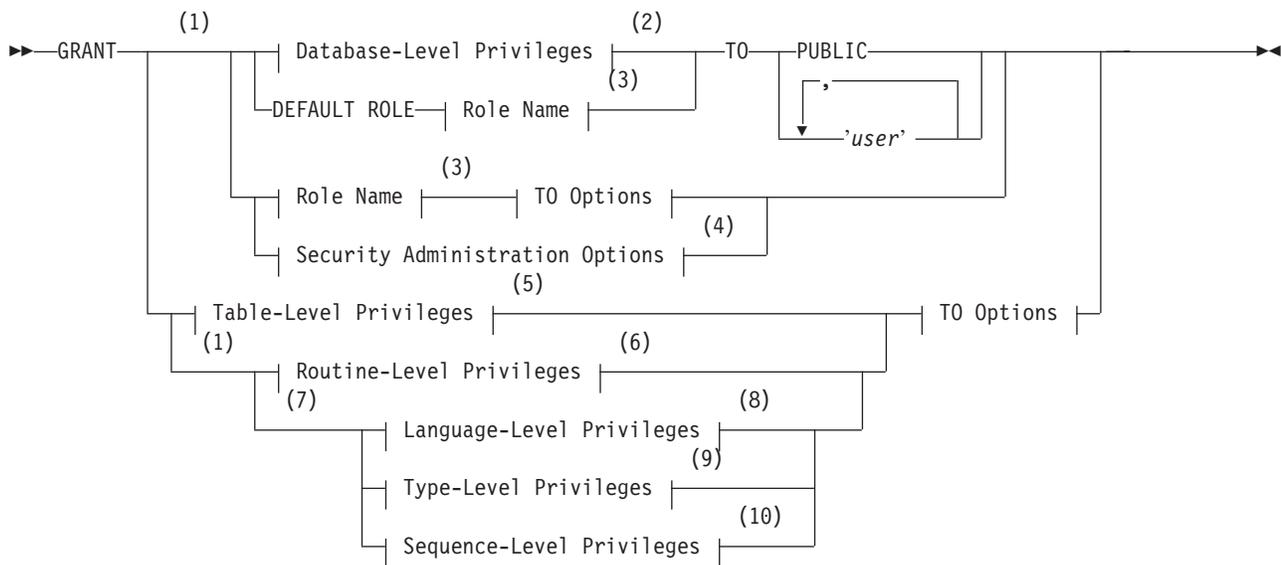
Related Information

For a task-oriented discussion of error handling and the **SQLSTATE** variable, see the *IBM Informix Guide to SQL: Tutorial*. For a discussion of concepts related to the GET DIAGNOSTICS statement and the **SQLSTATE** variable, see the *IBM Informix ESQL/C Programmer's Manual*.

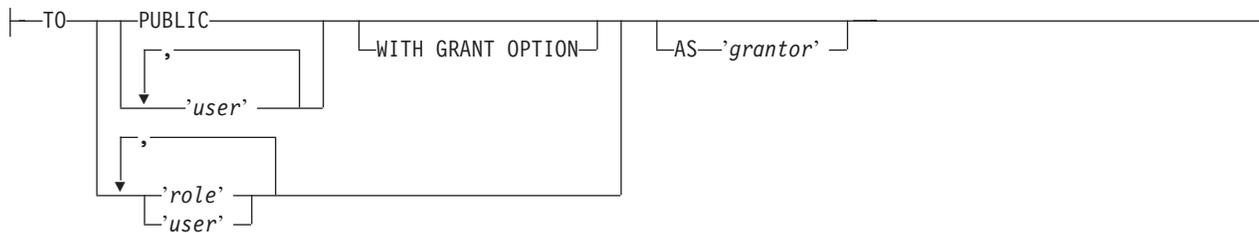
GRANT

Use the GRANT statement to assign access privileges and roles to users and to other roles. Users who hold the DBSECADM role can use this statement to assign user security labels and exemptions from label-based access control (LBAC) security rules.

Syntax



TO Options:



Notes:

- 1 Informix extension
- 2 See "Database-Level Privileges" on page 2-397
- 3 See "Role Name" on page 2-408
- 4 See "Security Administration Options" on page 2-413
- 5 See "Table-Level Privileges" on page 2-399
- 6 See "Routine-Level Privileges" on page 2-405
- 7 Dynamic Server only
- 8 See "Langage-Level Privileges (IDS)" on page 2-406
- 9 See "Type-Level Privileges (IDS)" on page 2-403
- 10 See "Sequence-Level Privileges (IDS)" on page 2-407

Element	Description	Restrictions	Syntax
<i>grantor</i>	Authorization identifier of a user who can use REVOKE to cancel the effects of this GRANT statement. If AS clause is omitted, default is login name of user issuing this statement	Must be valid <i>user</i> name (not a <i>role</i> name)	" Owner Name " on page 5-45
<i>role</i>	Name of an existing role to which you grant one or more access privileges, or to which you assign another role	Must exist in the database	" Owner Name " on page 5-45

Element	Description	Restrictions	Syntax
<i>user</i>	Authorization identifier of a user to whom you grant one or more access privileges, or to whom you assign a role	Same as for <i>grantor</i>	“ Owner Name ” on page 5-45

Usage

The GRANT statement extends access privileges to other users that would normally accrue only to the DBA or to the creator of an object. Subsequent GRANT statements do not affect privileges that have already been granted to a user.

You can use the GRANT statement for operations like the following:

- Authorize others to use or administer a database that you create
- Allow others to view, alter, or drop a table, synonym, view or (for Dynamic Server only) a sequence object that you create
- Allow others to use a data type or the SPL language, or (for Dynamic Server only) to execute a user-defined routine (UDR) that you create
- Assign a role and its privileges to users, or to PUBLIC, or to another role
- Assign a default role to one or more users or to PUBLIC
- If you hold the DBSECADM role, assign LBAC security labels or exemptions from rules of LBAC security policies to users,

You can grant privileges to a previously created role or to a built-in role. You can grant a role to PUBLIC, to individual users, or to another role.

If you enclose *grantor*, *role*, or *user* in quotation marks, the name is case sensitive and is stored exactly as you typed it. In an ANSI-compliant database, if you do not use quotation marks as delimiters, the name is stored in uppercase letters.

Privileges that you grant remain in effect until you cancel them with a REVOKE statement. Only the grantor of a privilege can revoke that privilege. The grantor is the person who issues the GRANT statement, unless the AS *grantor* clause transfers the right to revoke those privileges to another user.

Only the owner of an object or a user to whom privileges were explicitly granted with the WITH GRANT OPTION keywords can grant privileges on an object. Having DBA privileges is not sufficient. As DBA, however, you can grant a privilege on behalf of another user by using the AS *grantor* clause. For privileges on database objects whose owner is not a user recognized by the operating system (for example, user **informix**), the AS *grantor* clause is useful.

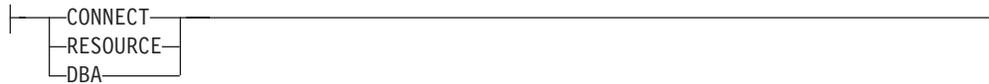
The keyword PUBLIC extends the specified privilege or role to all users. If you intend to restrict privileges that PUBLIC already holds to only a subset of users, you must first revoke those privileges from PUBLIC.

To grant privileges on one or more fragments of a table that has been fragmented by expression, see “ GRANT FRAGMENT ” on page 2-423.

Database-Level Privileges

Database-level access privileges affect access to a database. Only individual users, rather than roles, can hold database privileges.

Database-Level Privileges:



When you create a database with the CREATE DATABASE statement, you are the owner and automatically receive all database-level privileges.

The database remains inaccessible to any other users until you, as DBA, grant database privileges to them.

As database owner, you also receive table-level privileges on all tables in the database automatically. For more information about table-level privileges, see “Table-Level Privileges” on page 2-399.

Recommendation: Only user **informix** can modify system catalog tables directly. Except as noted specifically in your database server documentation, however, do not use DML statements to insert, delete, or update rows of system catalog tables directly, because modifying data in these tables can destroy the integrity of the database.

When database-level privileges conflict with table-level privileges, the more restrictive privileges take precedence.

Database access levels are, from lowest to highest, Connect, Resource, and DBA. Use the corresponding keyword to grant a level of access privilege.

Privilege	Effect
CONNECT	Lets you query and modify data You can modify the database schema if you own the database object that you intend to modify. Any user with the Connect privilege can perform the following operations: <ul style="list-style-type: none">• Connect to the database with the CONNECT statement or another connection statement• Execute SELECT, INSERT, UPDATE, and DELETE statements, provided the user has the necessary table-level privileges• Create views, provided the user has the Select privilege on the underlying tables• Create synonyms• Create temporary tables and create indexes on the temporary tables• Alter or drop a table or an index, provided the user owns the table or index (or has Alter, Index, or References privileges on the table)• Grant privileges on a table or view, provided the user owns the table (or was given privileges on the table with the WITH GRANT OPTION keywords)
RESOURCE	Lets you extend the structure of the database In addition to the capabilities of the Connect privilege, the holder of the Resource privilege can perform the following functions: <ul style="list-style-type: none">• Create new tables• Create new indexes• Create new UDRs• Create new data types

Privilege	Effect
DBA	<p>Has all the capabilities of the Resource privilege and can perform the following additional operations:</p> <ul style="list-style-type: none"> • Grant any database-level privilege, including the DBA privilege, to another user • Grant any table-level privilege to another user or to a role • Grant a role to a user or to another role • Revoke a privilege whose grantor you specify as the <i>revoker</i> in the AS clause of the REVOKE statement • Restrict the Execute privilege to DBAs when registering a UDR • Execute the SET SESSION AUTHORIZATION statement • Create any database object • Create tables, views, and indexes, designating another user as owner of these objects • Alter, drop, or rename database objects, regardless of who owns them • Execute the DROP DISTRIBUTIONS option of the UPDATE STATISTICS statement • Execute DROP DATABASE and RENAME DATABASE statements

User **informix** has the privilege required to alter the tables of the system catalog, including the **systables** table.

The following example uses the PUBLIC keyword to grant the Connect privilege on the currently active database to all users:

```
GRANT CONNECT TO PUBLIC;
```

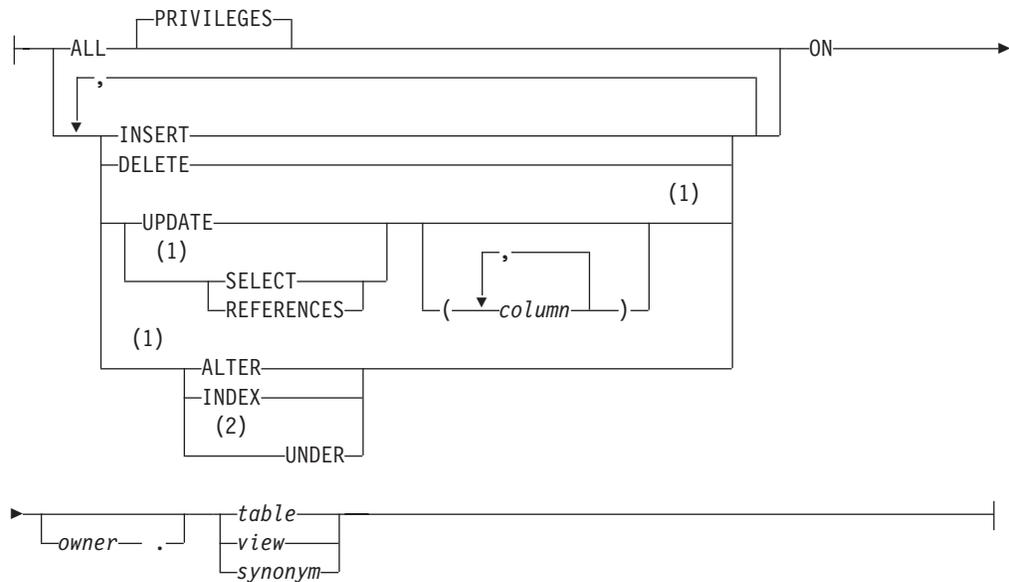
You cannot grant database-level privileges to a role. Only individual users or PUBLIC can hold database-level privileges.

Table-Level Privileges

When you create a table with the CREATE TABLE statement, you are the table owner and automatically receive all table-level privileges. You cannot transfer ownership to another user, but you can grant table-level privileges to another user or to a role. (See, however, "RENAME TABLE" on page 2-492, which can change both the name and the ownership of a table.)

A user with the database-level DBA privilege automatically receives all table-level privileges on every table in that database.

Table-Level Privileges:



Notes:

- 1 Informix extension
- 2 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>column</i>	Column on which the References, Select, or Update privilege is granted. Default scope is all columns of <i>table, view, or synonym</i> .	Must be a column of the <i>table, view, or synonym</i>	"Identifier " on page 5-21
<i>owner</i>	Name of the user who owns the <i>table, view, or synonym</i>	Must be a valid authorization identifier	" Owner Name " on page 5-45
<i>synonym, table, view</i>	Synonym, table, or view on which privileges are granted.	Must exist in the current database	"Identifier " on page 5-21

The GRANT statement can list one or more of the following keywords to specify the table privileges that you grant to the same users or roles.

Privilege	Effect
INSERT	Lets you insert rows
DELETE	Lets you delete rows
SELECT	Lets you access any column in SELECT statements. You can restrict the Select privilege to one or more columns by listing the columns.
UPDATE	Lets you access any column in UPDATE statements. You can restrict the Update privilege to one or more columns by listing the columns.
REFERENCES	Lets you define referential constraints on columns. You must have the Resource privilege to take advantage of the References privilege. (You can add, however, a referential constraint during an ALTER TABLE statement without holding the Resource privilege on the database.) You need only the References privilege to indicate cascading deletes. You do not need the Delete privilege to place cascading deletes on a table. You can restrict the References privilege to one or more columns by listing the columns.

Privilege	Effect
INDEX	Lets you create permanent indexes. You must have the Resource privilege to use the Index privilege. (Any user with the Connect privilege can create an index on temporary tables.)
ALTER	Lets you add or delete columns, modify column data types, add or delete constraints, change the locking mode of the table from PAGE to ROW, or add or drop a corresponding ROW data type for your table. It also lets you enable or disable indexes, constraints and triggers, as described in "SET Database Object Mode" on page 2-591. You must have the Resource privilege to use the Alter privilege. In addition, you also need the Usage privilege for any user-defined data type affected by the ALTER TABLE statement.
UNDER	Lets you create subtables under a typed table (for IDS only)
ALL	Provides all privileges listed above. The PRIVILEGES keyword is optional.

You can narrow the scope of a Select, Update, or References privilege by specifying the columns to which the privilege applies.

Specify the keyword PUBLIC as *user* if you intend the GRANT statement to apply to all users.

Some simple examples that follow illustrate how to give table-level privileges with the GRANT statement.

The following statement grants the privilege to delete and select values in any column in the table **customer** to users **mary** and **john**. It also grants the Update privilege, but only for columns **customer_num**, **fname**, and **lname**:

```
GRANT DELETE, SELECT, UPDATE (customer_num, fname, lname)
  ON customer TO mary, john;
```

To grant the same privileges as those above to all authorized users, use the keyword PUBLIC as the following example shows:

```
GRANT DELETE, SELECT, UPDATE (customer_num, fname, lname)
  ON customer TO PUBLIC;
```

For a Dynamic Server database, suppose a user named **mary** has created a typed table named **tab1**. By default, only user **mary** can create subtables under the **tab1** table. If **mary** wants to grant the ability to create subtables under the **tab1** table to a user named **john**, **mary** must enter the following GRANT statement:

```
GRANT UNDER ON tab1 TO john;
```

After receiving the Under privilege on table **tab1**, user **john** can create one or more subtables under **tab1**.

Effect of the ALL Keyword

The ALL keyword grants all possible table-level privileges to the specified user. If any or all of the table-level privileges do not exist for the grantor, the GRANT statement with the ALL keyword succeeds (in the sense of **SQLCODE** being set to zero, even if the possible privileges are an empty set for the grantor on the table). In this case, however, the following **SQLSTATE** warning is returned:

```
01007 - Privilege not granted.
```

For example, assume that user **ted** has the Select and Insert privileges on the **customer** table with the authority to grant those privileges to other users.

User **ted** wants to grant all table-level privileges to user **tania**. So user **ted** issues the following GRANT statement:

```
GRANT ALL ON customer TO tania;
```

This statement executes successfully but returns **SQLSTATE** code 01007 for the following reasons:

- The statement succeeds in granting the Select and Insert privileges to user **tania** because user **ted** has those privileges and the right to grant those privileges to other users.
- The other privileges implied by the ALL keyword were not grantable by user **ted** and, therefore, were not granted to user **tania**.

With Dynamic Server, if you grant all table-level privileges with the ALL keyword, the privileges includes the Under privilege only if the table is a typed table. The grant of ALL privileges does not include the Under privilege if the table is not based on a ROW type.

If the table owner grants ALL privileges on a traditional relational table and later changes that table to a typed table, the table owner must explicitly grant the Under privilege to allow other users to create subtables under it.

Table Reference

You grant table-level privileges directly by specifying the name or an existing synonym of a table or of a view, which you can qualify with the *owner* name.

Table Reference:



Element	Description	Restrictions	Syntax
<i>owner</i>	Name of the user who owns the <i>table</i> , <i>view</i> , or <i>synonym</i>	Must be a valid authorization identifier	" Owner Name " on page 5-45
<i>synonym, table, view</i>	Synonym, table, or view on which privileges are granted	The <i>table</i> , <i>view</i> , or <i>synonym</i> must exist in the database	"Identifier " on page 5-21

The object on which you grant privileges must reside in the current database.

In an ANSI-compliant database, if *owner* is not enclosed between quotation marks, the database stores the owner name in lowercase letters.

Privileges on Tables and Synonyms

In an ANSI-compliant database, if you create a table, only you, its owner, have any table-level privileges until you explicitly grant them to others.

When you create a table in a database that is *not* ANSI compliant, however, PUBLIC receives Select, Insert, Delete, Under, and Update privileges for that table

and its synonyms. (The `NODEFDAC` environment variable, when set to `yes`, prevents `PUBLIC` from automatically receiving these table-level privileges.)

To allow access only to some users, or only on some columns in a database that is not ANSI compliant, you must explicitly revoke the privileges that `PUBLIC` receives by default, and then grant only the privileges that you intend. For example, this series of statements grants privileges on the entire `customer` table to users `john` and `mary`, but restricts `PUBLIC` access to the `Select` privilege on only four of the columns in that table:

```
REVOKE ALL ON customer FROM PUBLIC;  
GRANT ALL ON customer TO john, mary;  
GRANT SELECT (fname, lname, company, city) ON customer TO PUBLIC;
```

Privileges on a View

You must have at least the `Select` privilege on a table or columns to create a view on that table. For views that reference only tables in the current database, if the owner of a view loses the `Select` privilege on any base table underlying the view, the view is dropped.

You have the same privileges for the view that you have for the table or tables contributing data to the view. For example, if you create a view from a table to which you have only `Select` privileges, you can select data from your view but you cannot delete or update data. For information on how to create a view, see “`CREATE VIEW`” on page 2-274.

When you create a view, `PUBLIC` does not automatically receive any privileges for a view that you create. Only you have access to table data through that view. Even users who have privileges on the base table of the view do not automatically receive privileges for the view.

You can grant (or revoke) privileges on a view only if you are the owner of the underlying base tables, or if you received these privileges on the base table with the right to grant them (specified by the `WITH GRANT OPTION` keywords). You must explicitly grant those privileges within your authority, because `PUBLIC` does not automatically receive any privileges on a view when it is created.

The creator of a view can explicitly grant `Select`, `Insert`, `Delete`, and `Update` privileges for the view to other users or to a role. You cannot grant `Index`, `Alter`, `Under`, or `References` privileges on a view (nor can you specify the `ALL` keyword for views, because `ALL` confers `Index`, `References`, and `Alter` privileges).

Type-Level Privileges (IDS)

You can specify two privileges on data types that are not built-in data types:

- The `Usage` privilege on a user-defined data type
- The `Under` privilege on a named `ROW` type

Type-Level Privileges:

```
|-----|  
|  USAGE ON TYPE—type_name-----|  
|  UNDER ON TYPE—row_type_name-----|
```

Element	Description	Restrictions	Syntax
<i>row_type_name</i>	Named ROW type on which the Under privilege is granted	Named ROW data type must exist	"Identifier " on page 5-21; "Data Type " on page 4-21
<i>type_name</i>	User-defined type on which the Usage privilege is granted	User-defined data type must exist.	"Identifier " on page 5-21; "Data Type " on page 4-21

To see what privileges exist on user-defined data types, check the **sysxdtypes** system catalog table for the *owner* of each UDT, and the **sysxdttypeauth** system catalog table for any other users or roles that hold privileges on UDTs. See the *IBM Informix Guide to SQL: Reference* for information on the system catalog tables.

For all the *built-in data types*, however, these access privileges are automatically available to PUBLIC and cannot be revoked.

USAGE Privilege

You own any user-defined data type (UDT) that you create. As owner, you automatically receive the Usage privilege on that data type and can grant the Usage privilege to others so that they can reference the type name or data of that type in SQL statements. DBAs can also grant the Usage privilege for UDTs.

The following example grants user **mark** access privileges to use the **widget** user-defined type:

```
GRANT USAGE ON TYPE widget TO mark;
```

If you grant Usage privilege to a user (or to a role) that has Alter privileges, that grantee can add a column to a table that contains values of your UDT.

Without privileges from the GRANT statement, any user can issue SQL statements that reference built-in data types. In contrast, a user must receive an explicit Usage privilege from a GRANT statement to use a distinct data type, even if the distinct type is based on a built-in type.

For more information about user-defined types, see "CREATE OPAQUE TYPE " on page 2-145, "CREATE DISTINCT TYPE " on page 2-100, the discussion of data types in the *IBM Informix Guide to SQL: Reference* and the *IBM Informix Database Design and Implementation Guide*.

UNDER Privilege

You own any named ROW type that you create. If you want other users to be able to create subtypes under this named ROW type, you must grant to these users the Under privilege on your named ROW type.

For example, suppose that you create a ROW type named **rtype1**:

```
CREATE ROW TYPE rtype1 (cola INT, colb INT);
```

If you want another user named **kathy** to be able to create a subtype under this named ROW type, you must grant the Under privilege on this named ROW type to user **kathy**:

```
GRANT UNDER ON ROW TYPE rtype1 TO kathy;
```

Now user **kathy** can create another ROW type under the **rtype1** ROW type, even though **kathy** is not the owner of the **rtype1** ROW type:

```
CREATE ROW TYPE rtype2 (colc INT, cold INT) UNDER rtype1;
```

For more about named ROW types, see “CREATE ROW TYPE ” on page 2-168, and the discussion of data types in the *IBM Informix Guide to SQL: Reference* and the *IBM Informix Database Design and Implementation Guide*.

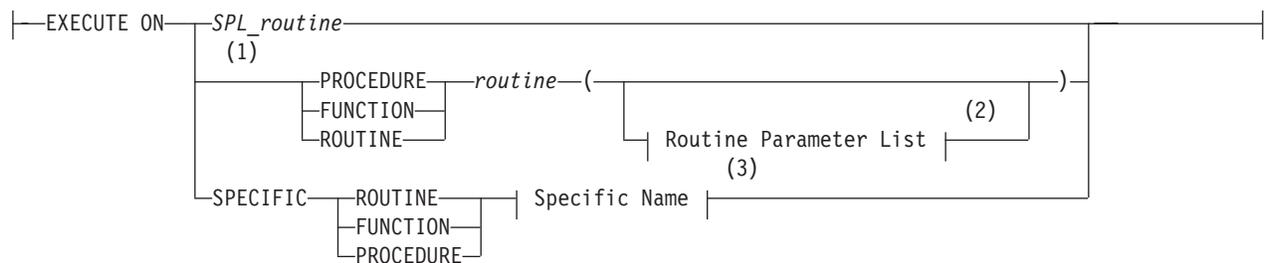
Routine-Level Privileges

When you create a user-defined routine (UDR), you become owner of the UDR and you automatically receive the Execute privilege on that UDR.

The Execute privilege allows you to invoke the UDR with an EXECUTE FUNCTION or EXECUTE PROCEDURE statement, whichever is appropriate, or with a CALL statement in an SPL routine. The Execute privilege also allows you to use a user-defined function in an expression, as in this example:

```
SELECT * FROM table WHERE in_stock(partnum) < 20;
```

Routine-Level Privileges:



Notes:

- 1 Dynamic Server only
- 2 See “Routine Parameter List” on page 5-67
- 3 See “Specific Name ” on page 5-73

Element	Description	Restrictions	Syntax
<i>routine</i>	A user-defined routine	Must exist	“Identifier ” on page 5-21
<i>SPL_routine</i>	An SPL routine	Must be unique in the database	“Identifier ” on page 5-21

The following statement grants Execute privilege on the **delete_order** routine to user **finn**:

```
GRANT EXECUTE on delete_order TO finn;
```

Whether you must grant the Execute privilege explicitly depends on the following conditions:

- If you have DBA-level privileges, you can use the DBA keyword of CREATE FUNCTION or CREATE PROCEDURE to restrict the default Execute privilege to users with the DBA privilege. You must explicitly grant the Execute privilege on that UDR to users who do not have the DBA privilege.
- If you have the Resource database-level privilege but not the DBA privilege, you cannot use the DBA keyword when you create a UDR:

- When you create a UDR in a database that is *not* ANSI compliant, PUBLIC can execute that UDR. You do not need to issue a GRANT statement for other users to receive the Execute privilege.
- Setting the **NODEFDAC** environment variable to yes prevents PUBLIC from executing the UDR until you explicitly grant the Execute privilege.
- In an ANSI-compliant database, the creator of a UDR must explicitly grant the Execute privilege on the UDR for other users to be able to execute it.

In Dynamic Server, if two or more UDRs have the same name, use a keyword from this list to specify which of those UDRs a user list can execute.

Keyword

UDR that the User Can Execute

SPECIFIC

The UDR identified by *specific name*

FUNCTION

Any function with the specified *routine name* (and parameter types that match *routine parameter list*, if specified)

PROCEDURE

Any procedure with the specified *routine name* (and parameter types that match *routine parameter list*, if specified)

ROUTINE

Functions or procedures with the specified *routine name* (and parameter types that match *routine parameter list*, if specified)

If both a user-defined function and a user-defined procedure of Dynamic Server have the same name and the same list of parameter data types, you can grant the Execute privilege to both with the keyword ROUTINE.

To limit the Execute privilege to one routine among several that have the same identifier, use the FUNCTION, PROCEDURE, or SPECIFIC keyword.

To limit the Execute privilege to a UDR that accepts certain data types as arguments, include the routine parameter list or use the SPECIFIC keyword to introduce the specific name of a UDR.

In Dynamic Server, if an external function has a negator function, you must grant the Execute privilege on both the external function and its negator function before users can execute the external function.

A user must have the Usage privilege on a language to register a user-defined routine of Dynamic Server that is written in that language.

Language-Level Privileges (IDS)

Dynamic Server also supports *language-level privileges*, which specify the programming languages of UDRs that users who have been granted Usage privileges for a given language can register in the database. This is the syntax of the USAGE ON LANGUAGE clause for specifying language-level privileges:

Language-Level Privileges:

OPTION keywords when you grant ALTER, SELECT, or ALL to a user or to PUBLIC (but not to a role) as privileges on a sequence object.

Alter Privilege

You can grant the Alter privilege on a sequence to another user or to a role. The Alter privilege enables a specified user or role to modify the definition of a sequence with the ALTER SEQUENCE statement or to rename the sequence with the RENAME SEQUENCE statement.

The following statement grants the Alter privilege to user **mark** on the **cust_seq** sequence object:

```
GRANT ALTER ON cust_seq TO mark;
```

Select Privilege

You can grant the Select privilege on a sequence to another user or to a role. The Select privilege enables a specified user or role to use *sequence.CURRVAL* and *sequence.NEXTVAL* expressions in SQL statements to read and to increment (respectively) the value of a sequence.

The following statement grants the Select privilege to user **mark** on the **cust_seq** sequence object:

```
GRANT SELECT ON cust_seq TO mark;
```

ALL Keyword

You can specify the ALL keyword to grant both Alter and Select privileges on a sequence object to another user or to a role, or to a list of users or roles.

The User List

You can grant privileges to an individual user or to a list of users. You can also specify the PUBLIC keyword to grant privileges to all users.

User List:



Element	Description	Restrictions	Syntax
<i>user</i>	Login name of a user to whom you are granting privilege or granting a role	Must be a valid authorization identifier	“ Owner Name ” on page 5-45

The following example grants the Insert table-level privilege on **table1** to the user **mary** in a database that is not ANSI-compliant:

```
GRANT INSERT ON table1 TO mary;
```

In an ANSI-compliant database, if you do not include quotation marks as delimiters around *user*, the name of the user is stored in uppercase letters.

Role Name

You can use the GRANT statement to associate a list of one or more users (or all users, using the PUBLIC keyword) with a *role* name that can describe what they

do. After you declare and grant a role, access privileges that you grant to that role are thereby granted to all users who are currently associated with that role.

Role Name:



Element	Description	Restrictions	Syntax
<i>role</i>	Role that is granted, or to which a privilege or another role is granted	Must exist. If enclosed between quotation marks, <i>role</i> is case sensitive.	" Owner Name " on page 5-45

You can also grant an existing role to another role. This action gives whatever privileges the granted role possesses to all users who have the receiving role.

Granting a Role to a User or to Another Role

You must register a role in the database before the role can be used in a GRANT statement. For more information, see "CREATE ROLE " on page 2-164.

A DBA has the authority to grant a new role to another user. If a user receives a role WITH GRANT OPTION, that user can grant the role to other users or to another role. Users keep a role that was granted to them until the REVOKE statement breaks the association between their login names and the role name.

Important: The CREATE ROLE and GRANT statements do not activate the role. A non-default role has no effect until SET ROLE enables it. The grantor or the grantee of a role can issue the SET ROLE statement.

The following example shows the actions required to grant and activate the role **payables** to a group of employees who perform account payable functions. First the DBA creates role **payables**, then grants it to **maryf**.

```
CREATE ROLE payables;
GRANT payables TO maryf WITH GRANT OPTION;
```

The DBA or **maryf** can activate the role with the following statement:

```
SET ROLE payables;
```

User **maryf** has the WITH GRANT OPTION authorization to grant **payables** to other employees who pay accounts.

```
GRANT payables TO charly, gene, marvin, raoul;
```

If you grant privileges for one role to another role, the recipient role has the combined set of privileges that have been granted to both roles. The following example grants the role **petty_cash** to the role **payables**:

```
CREATE ROLE petty_cash;
SET ROLE petty_cash;
GRANT petty_cash TO payables;
```

After all of these statements execute successfully, if user **raoul** uses the SET ROLE statement to make **payables** his current role, then (aside from the effects of any REVOKE operations) he holds the following combined set of access privileges:

- The privileges granted to the **payables** role
- The privileges granted to the **petty_cash** role

- The privileges granted individually to **raoul**
- The privileges granted to PUBLIC

If you attempt to grant a role to yourself, either directly or indirectly, the database server generates an error. (For an important exception to this rule, however, see the description of the “DBSECADM Clause” on page 2-414.)

The database server also generates an error if you include the WITH GRANT OPTION keywords in a GRANT statement that assigns a role to another role.

Granting a Privilege to a Role

You can grant table-level and routine-level privileges to a role if you have the authority to grant these same privileges to login names or to PUBLIC. A role cannot hold database-level privileges.

In Dynamic Server, you can also grant type-level privileges to a role.

The syntax is more restricted for granting privileges to a role than to a user:

- You can specify the AS *grantor* clause only in Dynamic Server.
In this way, whoever has the role can revoke these same privileges. For more information, see “AS grantor Clause” on page 2-412.
- You cannot include the WITH GRANT OPTION clause.
A role cannot, in turn, grant the same access privileges to another user.

This example grants Insert privilege on the **supplier** table to the role **payables**:

```
GRANT INSERT ON supplier TO payables;
```

Anyone who has been granted the **payables** role, and who successfully activates it by issuing the SET ROLE statement, can now insert rows into **supplier**.

Granting a Default Role

The DBA or the owner of the database (by default, user **informix**) can define a *default role* for one or more users or for PUBLIC with the GRANT DEFAULT ROLE statement. A default role is activated when the user connects to the database. The SET ROLE statement is not required to activate a default role.

Default roles are useful if users access databases through client applications that cannot modify access privileges nor set roles.

A default role can specify a set of access privileges to all the users who are assigned that role, as in the following example:

```
CREATE ROLE accounting;
GRANT ALTER, INSERT, SELECT ON stock TO accounting;
GRANT DEFAULT ROLE accounting TO mary, asok, vlad;
```

The last statement provides users **mary**, **asok**, and **vlad** with **accounting** as their default role. If any of these users connects to a database, that user activates whatever privileges the **accounting** role holds, in addition to any privileges that the user already possesses as an individual or as PUBLIC.

The role must already exist and the user must have the access privileges to set the role. If the role has not previously been granted to a user, it is granted as part of setting the default role.

If no default role is defined for a user nor for PUBLIC, then no role is set, and the existing privileges of the user are in effect.

The following example shows how the default role can be assigned to all users:

```
DATABASE hrdb;
CREATE ROLE emprole;
GRANT CONNECT TO PUBLIC;
GRANT SELECT ON emptab TO emprole;
GRANT emprole TO PUBLIC;
GRANT DEFAULT ROLE emprole TO PUBLIC;
```

Note: Using GRANT DEFAULT ROLE is an alternative to issuing the SET ROLE statement in the **sysdbopen()** procedure. Default roles defined using the **sysdbopen()** procedure, however, have precedence over any other role when a user establishes a connection.

Changing the default role for a user or for PUBLIC only affects new database connections. Existing connections continue to run under currently assigned roles. If one default role was granted to *user*, and another default role was granted to PUBLIC, the default role granted to *user* takes precedence at connection time.

A default role cannot be assigned to another role. Because roles are not defined across databases, the default role must be assigned for each database. No options besides the *user-list* are valid after the TO keyword in the GRANT DEFAULT ROLE statement. The database server issues an error if you attempt to include the *AS grantor* clause or the WITH GRANT OPTION clause.

Granting the EXTEND Role (IDS)

The Database Server Administrator (DBSA), by default user **informix**, can grant the built-in EXTEND roles to one or more users or to PUBLIC with the GRANT EXTEND TO *user-list* statement. If the IFX_EXTEND_ROLE configuration parameter is set to ON, only users who hold the EXTEND role can create or drop UDRs that are written in C or Java, external languages that can support shared libraries. The SET ROLE statement is not required for the EXTEND role to have this effect; it is sufficient for a user to hold the EXTEND role without using SET ROLE to enable it. The following example grants this role to user **max**:

```
GRANT EXTEND TO 'max';
```

This statement enables user **max** to create or drop UDRs, without requiring **max** to issued the SET ROLE EXTEND statement. (Here the quotation marks preserve the lowercase letters in the authorization identifier **max**.)

In databases for which this security feature is not needed, the DBSA can disable this restriction on who can create or drop external UDRs by setting the IFX_EXTEND_ROLE configuration parameter to OFF in the ONCONFIG file. When IFX_EXTEND_ROLE is set to OFF, any user who holds the Resource privilege can create or drop external UDRs. See “Database-Level Privileges” on page 2-397 for information about the Resource privileges.

WITH GRANT OPTION Keywords

The WITH GRANT OPTION keywords convey the privilege or role to *user* with the right to grant the same privileges or role to other users. You create a chain of privileges that begins with you and extends to *user* as well as to whomever *user* subsequently conveys the right to grant privileges. If you include WITH GRANT OPTION, you can no longer control the dissemination of privileges.

The following example grants the Alter and Select privileges to user **mark** on the **cust_seq** sequence object, with the ability to grant those privileges to others:

```
GRANT ALL ON cust_seq TO mark WITH GRANT OPTION;
```

If you revoke from *user* the privilege that you granted using the WITH GRANT OPTION keyword, you sever the chain of privileges. That is, when you revoke privileges from *user*, you automatically revoke the privileges of all users who received privileges from *user* or from the chain that *user* created (unless *user*, or the users who received privileges from *user*, were granted the same set of privileges by someone else).

The following examples illustrate this situation. You, as the owner of the table **items**, issue the following statements to grant access to user **mary**:

```
REVOKE ALL ON items FROM PUBLIC;  
GRANT SELECT, UPDATE ON items TO mary WITH GRANT OPTION;
```

User **mary** uses her privilege to grant users **cathy** and **paul** access to the table:

```
GRANT SELECT, UPDATE ON items TO cathy;  
GRANT SELECT ON items TO paul;
```

Later you revoke the access privileges for user **mary** on the **items** table:

```
REVOKE SELECT, UPDATE ON items FROM mary;
```

This single statement effectively revokes all privileges on the **items** table from users **mary**, **cathy**, and **paul**. If you want to create a chain of privileges with another user as the source of the privilege, use the AS *grantor* clause.

In Dynamic Server, the WITH GRANT OPTION keywords are valid only for users. They are not valid when a role is the grantee of a privilege or of another role.

The Database Server Administrator cannot include the WITH GRANT OPTION keywords in the GRANT EXTEND or GRANT DBSECADM statements. The DBSA cannot delegate to another user the authorization to grant the built-in EXTEND or DBSECADM roles. If more than one user needs either of these authorizations, they should be included in the **DBSA** group when the database server is installed.

In addition to the GRANT DBSECADM statement, none of the other security administration options of the GRANT statement support the WITH GRANT OPTION keywords. For more information about these statements and their syntax, see "Security Administration Options" on page 2-413.

AS grantor Clause

When you grant privileges, by default, you are the one who can revoke those privileges. The AS *grantor* clause lets you establish another user as the source of the privileges you are granting. When you use this clause, the login provided in the AS *grantor* clause replaces your login in the appropriate system catalog table. You can use this clause only if you have the DBA privilege on the database.

After you use this clause, only the specified *grantor* can REVOKE the effects of the current GRANT. Even a DBA cannot revoke a privilege unless that DBA is listed in the system catalog table as the source who granted the privilege.

The following example illustrates this situation. You are the DBA and you grant all privileges on the **items** table to user **tom** with the right to grant all privileges:

```
REVOKE ALL ON items FROM PUBLIC;
GRANT ALL ON items TO tom WITH GRANT OPTION;
```

The following example illustrates a different situation. You also grant Select and Update privileges to user **jim**, but you specify that the grant is made as user **tom**. (The records of the database server show that user **tom** is the grantor of the grant in the **sysstbauth** system catalog table, rather than you.)

```
GRANT SELECT, UPDATE ON items TO jim AS tom;
```

Later, you decide to revoke privileges on the **items** table from user **tom**, so you issue the following statement:

```
REVOKE ALL ON items FROM tom;
```

If instead, however, you try to revoke privileges from user **jim** with a similar statement, the database server returns an error, as the next example shows:

```
REVOKE SELECT, UPDATE ON items FROM jim;
```

580: Cannot revoke permission.

You receive an error because the database server record shows the original grantor as user **tom**, and you cannot revoke the privilege. Although you are the DBA, you cannot revoke a privilege that another user granted.

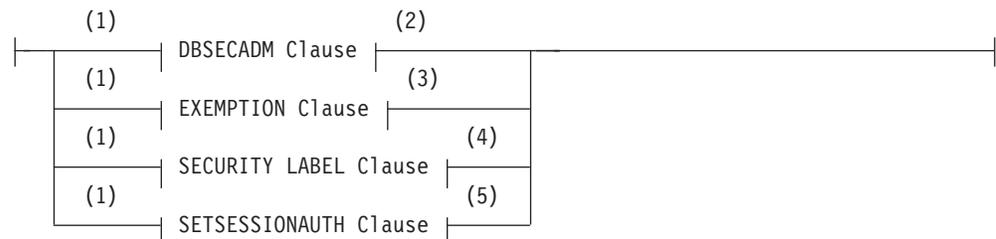
The AS *grantor* clause is not valid in the GRANT DEFAULT ROLE statement.

Security Administration Options

In conjunction with the REVOKE statement, the GRANT statement supports the discretionary access control (DAC) data security feature of Dynamic Server by specifying which users or roles hold privileges that are required to access the database or objects within the database.

The Security Administration Options of the GRANT statement, like their counterparts for the REVOKE statement, support an additional set of data security features, called label-based access control (LBAC). These features enable Dynamic Server to allow or withhold access to protected data on the basis of a comparing a row security label or column security label that is contained in the data object to the user security label and other credentials that have been granted to the user who is seeking access.

Security Administration Options:



Notes:

- 1 Dynamic Server only
- 2 See "DBSECADM Clause" on page 2-414

- 3 See "EXEMPTION Clause" on page 2-416
- 4 See "SECURITY LABEL Clause" on page 2-418
- 5 See "SETSESSIONAUTH Clause" on page 2-421

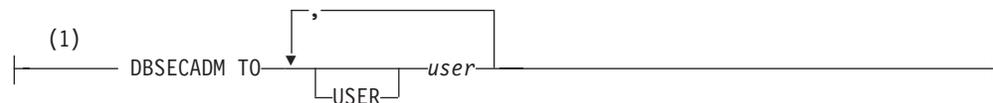
Use of these GRANT statement security administration options is restricted:

- Only the Database Server Administrator (DBSA), by default user **informix**, or (on UNIX) a member of the **DBSA** group, or (on Windows) a member of the **Informix-Admin** group, can use the GRANT DBSECADM statement to grant the DBSECADM role.
- Only a user who holds the DBSECADM role can issue the GRANT EXEMPTION, GRANT SECURITY LABEL, or GRANT SETSESSIONAUTH statements, or the corresponding REVOKE statements.

DBSECADM Clause

The GRANT DBSECADM statement enables the user to whom the DBSECADM role is granted to issue DDL statements that can create, alter, rename, or drop security objects, including security policies, security labels, and security components.

DBSECADM Clause:



Notes:

- 1 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>user</i>	User to whom the role is to be granted	Must be the authorization identifier of a user	" Owner Name " on page 5-45

The DBSECADM role is a built-in role that only the DBSA can grant. Unlike user-defined roles, whose scope is the database in which the role is created, the scope of the DBSECADM role is all of the databases of the Dynamic Server instance. It is not necessary for DBSA to reissue the GRANT DBSECADM statement in other databases of the same server. Like all built-in roles of Dynamic Server, the DBSECADM role is enabled when it is granted, without requiring activation by the SET ROLE statement, and it remains in effect until it is revoked.

Only a user who holds the DBSECADM role can issue the following SQL statements that create or modify security objects:

- ALTER SECURITY LABEL COMPONENT
- CREATE SECURITY LABEL
- CREATE SECURITY LABEL COMPONENT
- CREATE SECURITY POLICY
- DROP SECURITY LABEL
- DROP SECURITY LABEL COMPONENT
- DROP SECURITY POLICY

- RENAME SECURITY LABEL
- RENAME SECURITY LABEL COMPONENT
- RENAME SECURITY POLICY

Only a user who holds the DBSECADM role can use the following SQL statements to reference tables that are protected by a security policy:

- ALTER TABLE ... ADD SECURITY POLICY
- ALTER TABLE ... ADD ... IDSSECURITYLABEL [DEFAULT *label*]
- ALTER TABLE ... ADD ... [COLUMN] SECURED WITH
- ALTER TABLE ... DROP SECURITY POLICY
- ALTER TABLE ... MODIFY ... [COLUMN] SECURED WITH
- ALTER TABLE ... MODIFY ... DROP COLUMN SECURITY
- CREATE TABLE ... COLUMN SECURED WITH
- CREATE TABLE ... IDSSECURITYLABEL [DEFAULT *label*]
- CREATE TABLE ... SECURITY POLICY

The following GRANT and REVOKE statements also cannot be issued by a user who does not hold the DBSECADM role:

- GRANT EXEMPTION
- GRANT SECURITY LABEL
- GRANT SETSESSIONAUTH
- REVOKE EXEMPTION
- REVOKE SECURITY LABEL
- REVOKE SETSESSIONAUTH

The USER keyword that can follow the TO keyword is optional, and has no effect, but any authorization identifier that the DBSA specifies in the GRANT DBSECADM statement must be the identifier of an individual user, rather than the identifier of a role, or the PUBLIC group.

The *user* can be the DBSA who issues this GRANT DBSECADM statement. This is an important exception to the general restriction that the TO clause of the GRANT statement (like the FROM clause in the REVOKE statements) cannot explicitly reference the authorization identifier of the user who issues the statement. Unlike other roles, access privileges, user security labels, and exemptions on rules that the GRANT statement can specify, you can grant the DBSECADM role to yourself, if you are user **informix**, or a member of the DBSA group, or (on Windows[®]) if you are a member of the **Informix-Admin** group.

In the following example, the DBSA grants the DBSECADM role to user **niccolo**:

```
GRANT DBSECADM TO niccolo;
```

If this statement executes successfully, user **niccolo** can perform the LBAC operations listed above, provided that **niccolo** also holds sufficient discretionary access privileges on the database and on the database objects that those SQL statements reference.

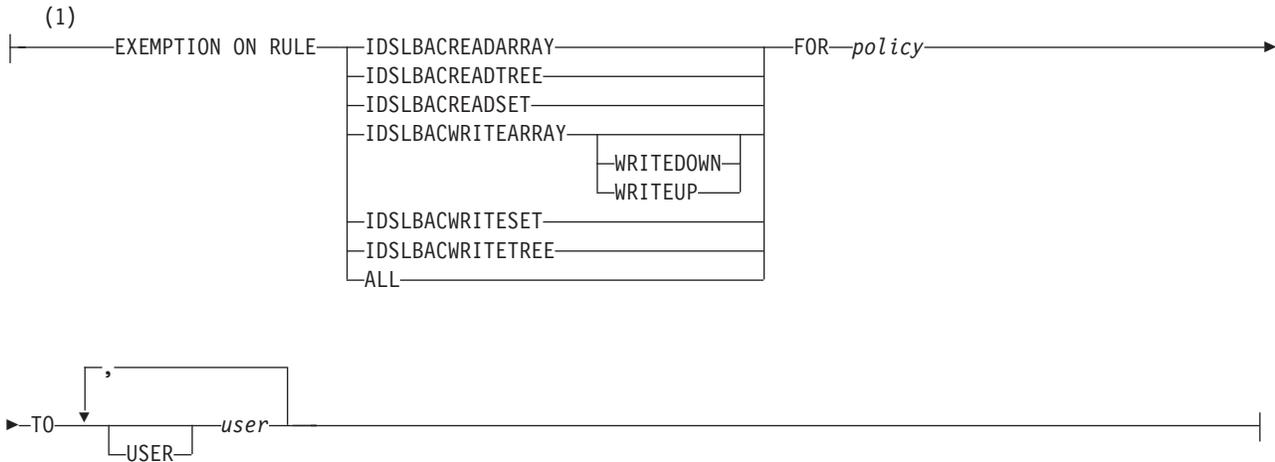
After a user is granted the DBSECADM role, only the DBSA can revoke it.

For a discussion of LBAC security objects, see your IBM Informix Dynamic Server Security Guide.

EXEMPTION Clause

The GRANT EXEMPTION statement modifies the security credentials of the specified user (or list of users) by disabling one or all of the rules of a specified security policy.

EXEMPTION Clause:



Notes:

1 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>policy</i>	The security policy from which the exemption is granted	Must exist in the database	"Identifier " on page 5-21
<i>user</i>	User to whom the exemption is to be granted	Must be the authorization identifier of a user	" Owner Name " on page 5-45

Only a user who holds the DBSECADM role can issue the GRANT EXEMPTION statement.

Rules on Which Exemptions Are Granted: The keyword that follows the ON keyword specifies the predefined LBAC access rule of the security policy (whose identifier follows the FOR keyword) for which an exemption is granted. The access rule for which exemption is granted does not apply when a table that is protected by the specified security policy is accessed by a user to whom the exemption is granted. For descriptions of the predefined rules for read access and for write access that are associated with a security policy, see the section "Rules Associated with a Security Policy" on page 2-182.

The following keywords of the GRANT EXEMPTION statement identify specific **IDSLBACRULES** rules from which this statement can exempt users:

- IDSLBACREADARRAY exempts the user from the **IDSLBACREADARRAY** rule for the specified security policy. That rule requires that each array component of the user security label must be greater than or equal to the corresponding array component of the data row security label.
- IDSLBACREADSET exempts the user from the **IDSLBACREADSET** rule for the specified security policy. That rule requires that each set component of the user security label must include the set component of the data row security label

- **IDSLBACREADTREE** exempts the user from the **IDSLBACREADTREE** rule for the specified security policy. That rule requires that each tree component of the user security label must include at least one of the elements in the tree component of the data row security label, or else the ancestor of one such element.
- **IDSLBACWRITEARRAY WRITEDOWN** exempts the user from one aspect of the **IDSLBACWRITEARRAY** rule for the specified security policy. That rule requires that each array component of the user security label must be equal to the array component of the data row security label. The user who holds this exemption can write to a row whose array component level is below the level in the label of the user. The user cannot, however, write to a row in whose label the array component level is above the level in the label of the user.
- **IDSLBACWRITEARRAY WRITEUP** exempts the user from one aspect of the **IDSLBACWRITEARRAY** rule for the specified security policy. The user who holds this exemption can write to a row whose array component level is above the level in the label of the user. The user cannot, however, write to a row in whose label the array component level is below the level in the label of the user.
- **IDSLBACWRITEARRAY** (with no **WRITEDOWN** or **WRITEUP** keyword) exempts the user from the **IDSLBACWRITEARRAY** rule for the specified security policy. The user who holds this exemption can write to a row without regard to the corresponding array component level of the row label.
- **IDSLBACWRITESET** exempts the user from the **IDSLBACWRITESET** rule for the specified security policy. That rule requires that each set component of the user security label must include the set component of the data row security label
- **IDSLBACWRITETREE** exempts the user from the **IDSLBACWRITETREE** rule for the specified security policy. That rule requires that each tree component of the user security label must include at least one of the elements in the tree component of the data row security label, or an ancestor of one such element.
- **ALL** exempts the user from all **IDSLBACRULES** rules for the specified security policy. This form of exemption is required to load data into a protected table.

In the following example, **DBSECADM** grants an exemption from all of the rules of the **MegaCorp** security policy to users **manoj** and **sam**:

```
GRANT EXEMPTION ON RULE ALL FOR MegaCorp TO manoj, sam;
```

Security Policies and Grantees of Exemptions: An exemption applies only to the rules of the security policy whose name follows the **FOR** keyword. A protected table can have multiple security labels, but no more than one security policy.

The **GRANT EXEMPTION** statement fails with an error if the specified policy does not exist in the database.

The **USER** keyword that can follow the **TO** keyword is optional, and has no effect, but any authorization identifier specified in the **GRANT EXEMPTION** statement must be the identifier of an individual user, rather than the identifier of a role. This *user* cannot be the **DBSECADM** who issues the same **GRANT EXEMPTION** statement.

In the following example, **DBSECADM** grants an exemption to user **lynette** from rule **IDSLBACREADARRAY** of the **MegaCorp** security policy:

```
GRANT EXEMPTION ON RULE IDSLBACREADARRAY FOR MegaCorp TO lynette;
```

This exemption bypasses the read access rules for all array components of security labels of the specified policy.

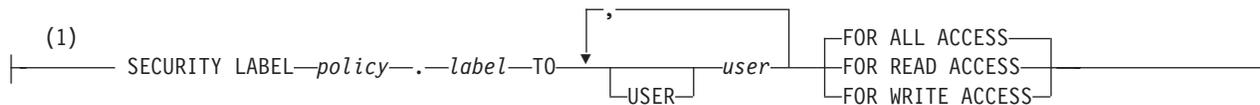
When the GRANT EXEMPTION statement successfully grants an exemption to a user, the database server updates the **syssecpolicyexemptions** table of the system catalog to register the new exemption (or multiple exemptions, if several users are listed after the TO keyword).

For a discussion of LBAC security objects, see your IBM Informix Dynamic Server Security Guide.

SECURITY LABEL Clause

The GRANT SECURITY LABEL statement grants a security label to a user or to a list of users.

SECURITY LABEL Clause:



Notes:

1 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>label</i>	Name of an existing security label	Must exist as a label for the specified security <i>policy</i>	"Identifier " on page 5-21
<i>policy</i>	The security policy of this <i>label</i>	Must already exist in the database	"Identifier " on page 5-21
<i>user</i>	User to whom the label is to be granted	Must be the authorization identifier of a user	" Owner Name " on page 5-45

Only a user who holds the DBSECADM role can issue the GRANT SECURITY LABEL statement.

A security label is a database object that is always associated with a security policy. That policy defines the set of valid security components that make up the security label. The label stores a set of one or more values for each component of the security policy.

The DBSECADM can associate a security label with the following entities:

- A column of a database table, which a *column security label* can protect
- A row of a database table, which a *row security label* can protect
- A user, whose *user security label* (and any exemptions from rules of the security policy that have been granted to the user) are called the *security credentials* of the user.

When a user who holds a security label for a specific security policy attempts to access a row that is protected by a row security label of the same security policy, the database server compares the sets of values in the user security label and in the row security label to determine whether the user should be allowed access to the data. Similarly, LBAC takes into account the user security label and the column security label in determining whether the credentials of the user should allow access to the protected column.

The GRANT SECURITY LABEL statement is the mechanism by which DBSECADM associates a user with a security label. (Data values in a protected table are associated with a row security label or with a column security label by options to the CREATE TABLE or ALTER TABLE statements that only DBSECADM can execute, rather than by the GRANT SECURITY LABEL statement.)

The USER keyword that can follow the TO keyword is optional, and has no effect, but any authorization identifier specified in the GRANT SECURITY LABEL statement must be the identifier of an individual user, rather than of a role.

Access Specifications: The list of users to whom the security label is granted can optionally be followed by keywords that specify the type of access to data that the security policy of the label protects

- FOR WRITE ACCESS

These keywords restrict the label to the write access rules of **IDSLBACRULES**, namely **IDLSBACWRITEARRAY**, **IDLSBACWRITESET**, and **IDLSBACWRITETREE**. These rules affect INSERT, DELETE, and UPDATE operations on protected data.

- FOR READ ACCESS

These keywords restrict the label to the read access rules of **IDSLBACRULES**, namely **IDLSBACWREADARRAY**, **IDLSBACREADSET**, and **IDLSBACREADTREE**. These rules affect SELECT, DELETE, and UPDATE operations on protected data.

- FOR ALL ACCESS

These keywords apply the label to all of the read and write access rules that are listed above. If the GRANT SECURITY LABEL statement includes no FOR ... ACCESS specification, this option takes effect as the default.

For more information about these **IDSLBACRULES** rules for label-based read and write access, see “Rules Associated with a Security Policy” on page 2-182. For information about exemptions to these rules that can be granted for a specific security policy, see “Rules on Which Exemptions Are Granted” on page 2-416.

If a user is granted a different security label for read access than for write access, then the values given for the security label components must follow these rules:

- For security label components of type ARRAY, the value must be the same in both security labels.
- For security label components of type SET, the values given in the security label used for WRITE access must be a subset of the values given in the security label used for READ access. If all of the values are the same, this is interpreted as being a subset, and is allowed.
- For security label components of type TREE, every element in the tree component of the security label for write access must be either an element or a descendent of an element in the tree component of the security label for read access.

In summary, when DBSECADM attempts to grant a security label for read access to a user who already holds a security label for write access, or vice versa, the read label cannot be more restrictive than the write label. Otherwise, the GRANT SECURITY LABEL statement fails with an error.

A user can be granted no more than two labels for the same security policy. If two labels are granted for the same policy, one label must be for read access and the other for write access. If DBSECADM attempts to grant a security label for read access to a user who already holds a security label for read access that is based on

the same security policy, the GRANT SECURITY LABEL statement fails with an error. A similar failure result if both labels are for write access and are on the same security policy.

In both of these cases, the first security label must be revoked explicitly by the REVOKE SECURITY LABEL statement before a second label can be granted for the same access mode and the same security policy. The only exception to this rule is if both labels specify the same values for component elements, because in this case both security labels are functionally identical, and no error is issued.

Rules for User Security Labels: The following rules affect security labels that are granted to users by the GRANT SECURITY LABEL statement:

- The *user* cannot be the DBSECADM who issues this GRANT SECURITY LABEL statement.
- A user without a security label has a NULL or zero label. A user with no security label cannot access data in a protected table, unless the user holds the necessary exemptions on the policy.
- By default, the IDSSECURITYLABEL column of a protected table cannot have NULL values. A user with no security label cannot insert data into a table with row protection, even if the user has been granted the necessary exemptions on the security policy, unless the row label is explicitly specified in the INSERT statement. For details of how to specify a security label explicitly in the INSERT statement, see “ Security Label Support Functions (IDS)” on page 4-119.
- User security labels have no effect on the following types of database tables, because these tables cannot be protected by a security policy:
 - Virtual Table Interface tables,
 - tables with Virtual Index Interface indexes,
 - tables in a typed-table hierarchy,
 - temporary tables.

Examples of Granting User Security Labels:

The following three statements create three security label components called **level**, **compartments**, and **groups** respectively:

```
CREATE SECURITY LABEL COMPONENT
  level ARRAY ['TS','S','C','U'];

CREATE SECURITY LABEL COMPONENT
  compartments SET {'A','B','C','D'};

CREATE SECURITY LABEL COMPONENT
  groups TREE ('G1' ROOT,
              'G2' UNDER ROOT,
              'G3' UNDER ROOT);
```

The following statement creates a security policy called **secPolicy** based on the three components above:

```
CREATE SECURITY POLICY secPolicy COMPONENTS
  level, compartments, groups;
```

The following statement creates a security label called **secLabel1**:

```
CREATE SECURITY LABEL secPolicy.secLabel1
  COMPONENT level 'S',
  COMPONENT compartments 'A', 'B',
  COMPONENT groups 'G2';
```

The following statement creates a security label called **secLabel2**:

```
CREATE SECURITY LABEL secPolicy.secLabel2
  COMPONENT level 'S',
  COMPONENT compartments 'B',
  COMPONENT groups 'G2';
```

The following statement creates a security label called **secLabel3**:

```
CREATE SECURITY LABEL secPolicy.secLabel3
  COMPONENT level 'S',
  COMPONENT compartments 'A',
  COMPONENT groups 'G3';
```

The following statement creates a security label called **secLabel4**:

```
CREATE SECURITY LABEL secPolicy.secLabel4
  COMPONENT level 'TS',
  COMPONENT compartments 'A',
  COMPONENT groups 'G1';
```

The following statement grants a security label for read access to user **sam**:

```
GRANT SECURITY LABEL secPolicy.secLabel1
  TO sam FOR READ ACCESS;
```

The following statement grants a security label for write access to user **sam**. This statement succeeds because it satisfies the rules given above.

```
GRANT SECURITY LABEL secPolicy.secLabel2
  TO sam FOR WRITE ACCESS;
```

The following statement grants a security label for read access to user **lynette**:

```
GRANT SECURITY LABEL secPolicy.secLabel1
  TO lynette FOR READ ACCESS;
```

The following statement attempts to grant a security label for write access to user **sam**. This statement fails because it violates the rule with respect to the tree component.

```
GRANT SECURITY LABEL secPolicy.secLabel3
  TO sam FOR WRITE ACCESS;
```

The following statement attempts to grant a security label for write access to user **sam**. This statement fails because it violates the rule with respect to the array component.

```
GRANT SECURITY LABEL secPolicy.secLabel4
  TO sam FOR WRITE ACCESS;
```

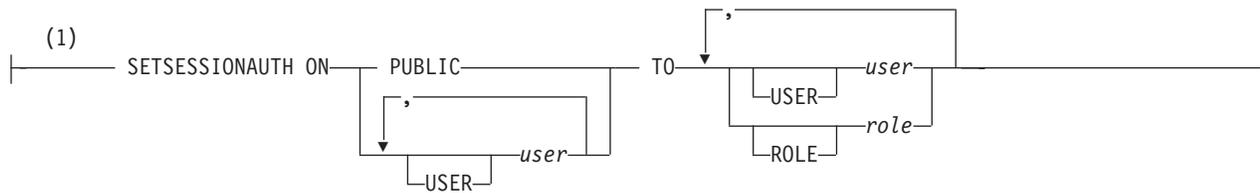
When the GRANT SECURITY LABEL statement successfully grants a security label to a user, the database server updates the **sysseclabelauth** table of the system catalog to register the new holder of the security label.

For a discussion of LBAC security objects, see your IBM Informix Dynamic Server Security Guide.

SETSESSIONAUTH Clause

The GRANT SETSESSIONAUTH statement grants the SETSESSIONAUTH privilege to one or more users or roles. This privilege allows the holder to use the SET SESSION AUTHORIZATION statement to set the session authorization to PUBLIC or to any one of a list of specified users.

SETSESSIONAUTH Clause:



Notes:

1 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>role</i>	Role to which the privilege is to be granted	Must be the authorization identifier of a role	“ Owner Name ” on page 5-45
<i>user</i>	After the TO keyword, a user to whom the privilege is to be granted. After the ON keyword, a user whose identity the grantee can specify in the SET AUTHORIZATION statement.	Must be the authorization identifier of a user	“ Owner Name ” on page 5-45

Only a user who holds the DBSECADM role can grant the SETSESSIONAUTH privilege. Both the SETSESSIONAUTH privilege and the DBA privilege are required to execute the SET AUTHORIZATION statement.

The user or PUBLIC specification that follows the ON keyword specifies whose identity the grantee of the SETSESSIONAUTH privilege can take while using SET SESSION AUTHORIZATION statement. This can be a user or PUBLIC but not a role. If PUBLIC is specified, then the grantee of the privilege can assume the identity of any database user.

The USER and ROLE keywords that can follow the TO keyword are optional. Neither the *user* nor the *role* can be the holder of the DBSECADM role who issues the GRANT SETSESSIONAUTH statement. The TO clause cannot specify PUBLIC as the grantee.

The following example grants to user **sam** the ability to set the session authorization to users **lynette** and **manoj**:

```
GRANT SETSESSIONAUTH ON lynette, manoj TO sam;
```

The next example grants to user **lynette** the ability to set the session authorization to PUBLIC:

```
GRANT SETSESSIONAUTH ON PUBLIC TO lynette;
```

Only a user who holds the DBSECADM role can revoke the SETSESSIONAUTH privilege. For a discussion of LBAC security objects, see your IBM Informix Dynamic Server Security Guide.

Related Information

Related statements: “ GRANT FRAGMENT ” on page 2-423, “ REVOKE ” on page 2-494, and “ REVOKE FRAGMENT ” on page 2-517

For information about roles, see the following statements: "CREATE ROLE " on page 2-164, "DROP ROLE " on page 2-334, and "SET ROLE " on page 2-647.

In the *IBM Informix Database Design and Implementation Guide*, see the discussion of privileges.

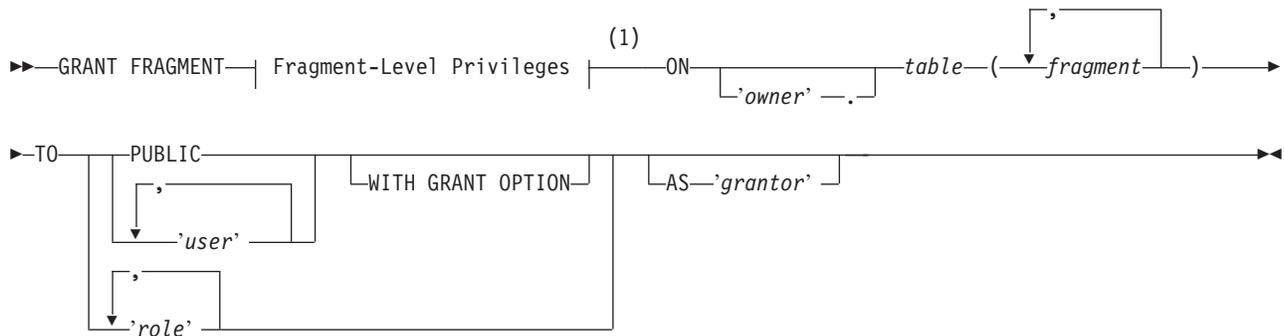
For a discussion of how to embed GRANT and REVOKE statements in programs, see the *IBM Informix Guide to SQL: Tutorial*.

GRANT FRAGMENT

Use the GRANT FRAGMENT statement to grant to one or more users or roles any or all of the Insert, Update, and Delete access privileges on individual fragments of a table that has been fragmented by expression.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 See "Fragment-Level Privileges" on page 2-424

Element	Description	Restrictions	Syntax
<i>fragment</i>	Partition or dbspace that stores a fragment of table	Must exist; cannot be delimited by quotation marks	"Identifier " on page 5-21
<i>grantor</i>	User who can revoke the privileges	Same as for <i>user</i>	" Owner Name " on page 5-45
<i>owner</i>	User who owns <i>table</i>	Must be owner of <i>table</i>	" Owner Name " on page 5-45
<i>role</i>	Role to receive privileges	Must exist in sysusers	" Owner Name " on page 5-45
<i>table</i>	Fragmented table on which fragment privileges are granted	Must exist and must be fragmented by expression	"Identifier " on page 5-21
<i>user</i>	User to whom privileges are to be granted	Must be a valid authorization identifier	" Owner Name " on page 5-45

Usage

The GRANT FRAGMENT statement is a special case of the GRANT statement for assigning privileges on table fragments in the local database. GRANT FRAGMENT is valid only for tables that are fragmented according to an expression-based distribution scheme. For an explanation of this type of fragmentation strategy, see “Expression Distribution Scheme” on page 2-15.

Fragment-Level Privileges

The keyword or keywords that follow the FRAGMENT keyword specify *fragment-level privileges*, which are a logical subset of table-level privileges:

Fragment-Level Privileges:



These keywords correspond to the following fragment-level privileges:

Keyword

Effect on Grantee

ALL Receives Insert, Delete, and Update privileges on the fragment

INSERT

Can insert rows into the fragment

DELETE

Can delete rows from the fragment

UPDATE

Can update rows in the fragment and in any columns.

Definition of Fragment-Level Authorization

In an ANSI-compliant database, the owner implicitly receives all table-level privileges on a newly created table, but no other users receive privileges.

A user who has table privileges on a fragmented table has the privileges implicitly on all fragments of the table. These privileges are not recorded in the **sysfragauth** system catalog table.

When a fragmented table is created in a database that is not ANSI compliant, the table owner implicitly receives all table-level privileges on the table, and other users (that is, PUBLIC) receive all fragment-level privileges by default. The privileges granted to PUBLIC are explicitly recorded in the **sysstabauth** system catalog table.

If you use the REVOKE statement to withdraw existing table-level privileges, however, you can then use the GRANT FRAGMENT statement to restore specified table-level privileges to users, roles, or PUBLIC on some subset of the fragments.

Whether or not the database is ANSI compliant, you can use the GRANT FRAGMENT statement to grant explicit Insert, Update, and Delete privileges on

one or more fragments of a table that is fragmented by expression. The privileges that the GRANT FRAGMENT statement grants are explicitly recorded in the `sysfragauth` system catalog table.

The Insert, Update, and Delete privileges that are conferred on table fragments by the GRANT FRAGMENT statement are collectively known as *fragment-level privileges* or *fragment-level authority*.

Effect of Fragment-Level Authorization in Statement Validation

Fragment-level privilege enables users to execute INSERT, DELETE, and UPDATE data manipulation language (DML) statements on table fragments, even if the grantees lack Insert, Update, and Delete privileges on the table as a whole. Users who lack the table privileges can insert, delete, and update rows in authorized fragments because of the algorithm by which the database server validates DML statements. This algorithm consists of the following checks:

1. When a user executes an INSERT, DELETE, or UPDATE statement, the database server first checks whether the user has the table privileges necessary for the operation attempted. If the table privileges exist, the statement continues processing.
2. If the table privileges do not exist, the database server checks whether the table is fragmented by expression. If the table is not fragmented by expression, the database server returns an error to the user. This error indicates that the user does not have the privilege to execute the statement.
3. If the table is fragmented by expression, the database server checks whether the user holds the fragment privileges necessary for the attempted operation. If the user holds the required fragment privileges, the database server continues to process the statement. If the fragment privileges do not exist, the database server returns an error to the user. This error indicates that the user does not have the privilege to execute the statement.

Duration of Fragment-Level Privileges

The duration of fragment-level privileges is tied to the duration of the fragmentation strategy for the table as a whole.

If you drop a fragmentation strategy by means of a DROP TABLE statement or by the INIT, DROP, or DETACH clauses of an ALTER FRAGMENT statement, you also drop any privileges that exist for the affected fragments. Similarly, if you drop a fragment or partition of a table, you also drop any privileges that exist for the fragment.

Tables that are created as a result of a DETACH or INIT clause of an ALTER FRAGMENT statement do not keep the privileges that the former fragment or fragments had when they were part of the fragmented table. Instead, such tables assume the default table privileges.

If a table on which fragment privileges are defined is changed to a table with a round-robin strategy or some other expression strategy, the fragment privileges are also dropped, and the table assumes the default table privileges.

Specifying Fragments

You can specify one fragment or a comma-separated list of fragments, with the name (or list of names) enclosed between parentheses that immediately follow the ON *table* specification. You cannot use quotation marks to delimit fragment names. The database server issues an error if you include no fragment, or if no fragment of the specified table matches a fragment that you list.

Each *fragment* must be referenced by the name of the partition where it is stored. If you did not declare an explicit identifier when you created the partition, its name defaults to the name of the dbspace that contains the partition.

After a dbspace is renamed successfully by the onspaces utility, only the new name is valid. Dynamic Server automatically updates existing fragmentation strategies in the system catalog to substitute the new dbspace name, but you must specify the new name in GRANT FRAGMENT statement to reference a fragment whose default name is the name of a renamed dbspace.

The TO Clause

The list of one or more users or roles that follows the TO keyword identifies the grantees. You can specify the PUBLIC keyword to grant the specified fragment-level privileges to all users.

You cannot use GRANT FRAGMENT to grant fragment-level privileges to yourself, either directly or through roles.

If you enclose *user* or *role* in quotation marks, the name is case sensitive and is stored exactly as you typed it. In an ANSI-compliant database, if you do not use quotation marks around *user* or around *role*, the name is stored in uppercase letters.

The following statement grants the Insert, Update, and Delete privileges on the fragment of the **customer** table in **part1** to user **larry**:

```
GRANT FRAGMENT ALL ON customer (part1) TO larry;
```

The following statement grants the Insert, Update, and Delete privileges on the fragments of the **customer** table in **part1** and **part2** to user **millie**:

```
GRANT FRAGMENT ALL ON customer (part1, part2) TO millie;
```

To grant privileges on all fragments of a table to the same user or users, you can use the GRANT statement instead of the GRANT FRAGMENT statement. You can also use the GRANT FRAGMENT statement for this purpose.

Assume that the **customer** table is fragmented by expression into three fragments, and these fragments reside in the dbspaces named **part1**, **part2**, and **part3**. You can use either of the following statements to grant the Insert privilege on all fragments of the table to user **helen**:

```
GRANT FRAGMENT INSERT ON customer (part1, part2, part3) TO helen;
```

```
GRANT INSERT ON customer TO helen;
```

Granting Privileges to One User or a List of Users

You can grant fragment-level privileges to a single user or to a list of users.

The following statement grants the Insert, Update, and Delete privileges on the fragment of the **customer** table in **part3** to user **oswald**:

```
GRANT FRAGMENT ALL ON customer (part3) TO oswald;
```

The following statement grants the Insert, Update, and Delete privileges on the fragment of the **customer** table in **part3** to users **jerome** and **hilda**:

```
GRANT FRAGMENT ALL ON customer (part3) TO jerome, hilda;
```

Granting One Privilege or a List of Privileges

When you specify fragment-level privileges in a GRANT FRAGMENT statement, you can specify one privilege, a list of privileges, or all privileges.

The following statement grants the Update privilege on the fragment of the **customer** table in **part1** to user **ed**:

```
GRANT FRAGMENT UPDATE ON customer (part1) TO ed;
```

The following statement grants the Update and Insert privileges on the fragment of the **customer** table in **part1** to user **susan**:

```
GRANT FRAGMENT UPDATE, INSERT ON customer (part1) TO susan;
```

The following statement grants the Insert, Update, and Delete privileges on the fragment of the **customer** table in **part1** to user **harry**:

```
GRANT FRAGMENT ALL ON customer (part1) TO harry;
```

WITH GRANT OPTION Clause

As in other GRANT statements, the WITH GRANT OPTION keywords specify that the grantee can grant the same fragment-level privileges to other users. WITH GRANT OPTION is not valid if the TO clause specifies a *role* as grantee. For additional information, see “WITH GRANT OPTION Keywords” on page 2-411.

The following statement grants the Update privilege on the fragment of the **customer** table in **part3** to user **george** and also gives george the right to grant the Update privilege on the same fragment to other users:

```
GRANT FRAGMENT UPDATE ON customer (part3) TO george WITH GRANT OPTION;
```

AS grantor Clause

The AS *grantor* clause of the GRANT FRAGMENT statement can specify the grantor of the privilege. You can use this clause only if you have the DBA privilege on the database. When you include the AS *grantor* clause, the database server lists the user or role who is specified as *grantor* as the grantor of the privilege in the **grantor** column of the **sysfragauth** system catalog table.

In the next example, the DBA grants the Delete privilege on the fragment of the **customer** table in **part3** to user **martha**, and uses the AS *grantor* clause to specify that user **jack** is listed in **sysfragauth** as the grantor of the privilege:

```
GRANT FRAGMENT DELETE ON customer (part3) TO martha AS jack;
```

One effect of the AS *grantor* clause in the previous example is that user **jack** can execute the REVOKE FRAGMENT statement to cancel the Delete fragment-level privilege that **martha** holds, if this GRANT FRAGMENT statement were the only source of the fragment authority of **martha** on the **customer** rows in **part3**.

Omitting the AS grantor Clause

When GRANT FRAGMENT does not include the AS *grantor* clause, the user who issues the statement is the default grantor of the specified fragment privileges.

In the next example, the user grants the Update privilege on the fragment of the **customer** table in **part3** to user **fred**. Because this statement does not specify the

AS *grantor* clause, the user who issues the statement is listed by default as the grantor of the privilege in the **sysfragauth** system catalog table.

```
GRANT FRAGMENT UPDATE ON customer (part3) TO fred;
```

If you omit the AS *grantor* clause of GRANT FRAGMENT, or if you specify your own login name as the *grantor*, you can later use the REVOKE FRAGMENT statement to revoke the privilege that you granted to the specified user. For example, if you grant the Delete privilege on the fragment of the **customer** table in **part3** to user **martha** but specify user **jack** as the grantor of the privilege, user **jack** can revoke that privilege from user **martha**, but you cannot revoke that privilege from user **martha**.

The DBA, or the owner of the fragment, can use the AS clause of the REVOKE FRAGMENT statement to revoke privileges on the fragment.

Related Information

Related statements: "GRANT " on page 2-395 and "REVOKE FRAGMENT " on page 2-517

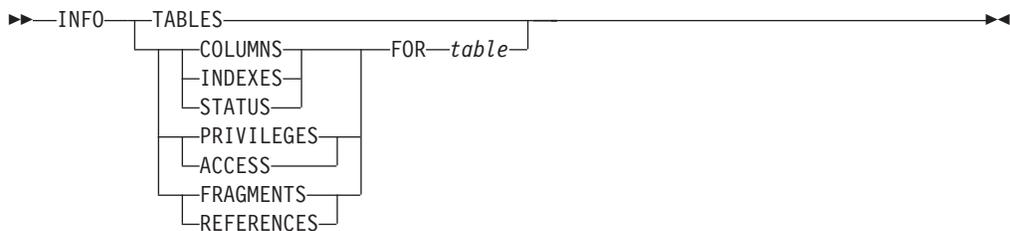
For a discussion of fragment-level and table-level privileges, see the *IBM Informix Database Design and Implementation Guide*.

INFO

Use the INFO statement to display information about databases tables.

This statement is an extension to the ANSI/ISO standard for SQL. You can use this statement only with DB–Access.

Syntax



Element	Description	Restrictions	Syntax
<i>table</i>	Table about which you seek information	Must exist	"Database Object Name " on page 5-16

Usage

You can use the INFO statement with DB–Access to display information about database tables.

Keywords of the INFO statement can display the following information.

INFO Keyword Information Displayed

TABLES

Table names in the current database

COLUMNS

Column information for a specified table

INDEXES

Index information for a specified table

FRAGMENTS

Fragmentation strategy for a specified table

ACCESS or PRIVILEGES

Access privileges for a specified table. (The ACCESS and PRIVILEGES keywords are synonyms in this context.)

REFERENCES

References privileges for columns of a specified table

STATUS

Status information for a specified table

- **TABLES Keyword**

Use TABLES to display a list of the tables in the current database, not including system catalog tables, in one of the following formats:

- If you are the owner of the **cust_calls** table, it appears as **cust_calls**.
- If you are *not* the owner of the **cust_calls** table, the name of the owner precedes the table name, such as **'june'.cust_calls**.

- **COLUMNS Keyword**

Use COLUMNS to display the names and data types of the columns in a specified table, showing for each column whether NULL values are allowed.

- **INDEXES Keyword**

Use INDEXES to display the name, owner, and type of each index in a specified table, the clustered status, and listing the indexed columns.

- **FRAGMENTS Keyword**

Use FRAGMENTS to display the names of dbspaces storing fragments of a table. If the table is fragmented with an expression-based distribution scheme, the INFO statement also shows the expressions.

- **ACCESS or PRIVILEGES Keyword**

Use ACCESS or PRIVILEGES to display user-access privileges for a specified table. (These two keywords are synonyms in this context.)

- **REFERENCES Keyword**

Use REFERENCES to display the References privilege for users for the columns of a specified table. For database-level privileges, use a SELECT statement to query the **sysusers** system catalog table.

- **STATUS Keyword**

Use STATUS to display information about the owner, row length, number of rows and columns, creation date, and status of audit trails for a specified table.

An alternative to using the INFO statement of SQL is to use the **Info** command of the **SQL** menu or of the **Table** menu of DB–Access to display the same and additional information.

Examples

Use the following example to list the user tables in a database:

```
INFO TABLES;
```

To display information about a specific table, use the syntax:

```
INFO info_keyword FOR table
```

Here *table* is the table name and *info_keyword* is one of the seven keyword options, besides TABLES, to the INFO statement. For example, to display the names of the columns of the table **customer**, use this statement:

```
INFO COLUMNS FOR customer;
```

This example produces the following output:

Column name	Type	Nulls
customer_num	serial	no
fname	char(15)	yes
lname	char(15)	yes
company	char(20)	yes
address1	char(20)	yes
address2	char(20)	yes
city	char(15)	yes
state	char(2)	yes
zipcode	char(5)	yes
phone	char(18)	yes

Related Information

Related statements: "GRANT " on page 2-395 and "REVOKE" on page 2-494

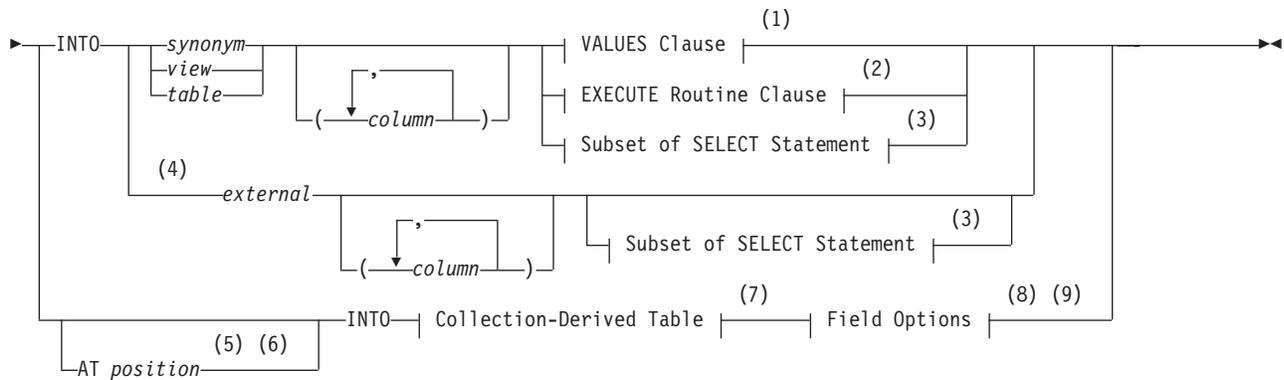
For a description of the **Info** option on the SQL menu or the **TABLE** menu in DB–Access, see the *IBM Informix DB–Access User's Guide*.

INSERT

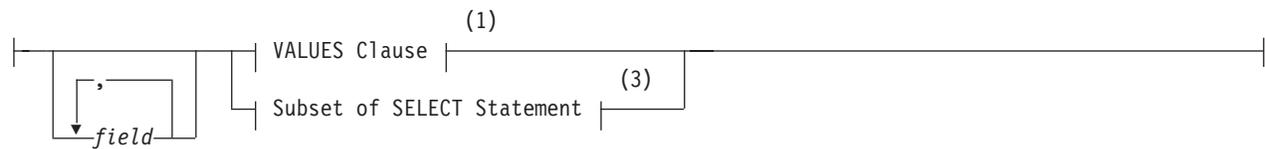
Use the INSERT statement to insert one or more new rows into a table or view, or to insert one or more elements into an SPL or Informix ESQL/C collection variable.

Syntax

▶▶—INSERT—————▶



Field Options:



Notes:

- 1 See "VALUES Clause" on page 2-434
- 2 See "Execute Routine Clause" on page 2-441
- 3 See "Subset of SELECT Statement" on page 2-441
- 4 Extended Parallel Server only
- 5 Stored Procedure Language only
- 6 ESQL/C only
- 7 See "Collection-Derived Table" on page 5-4
- 8 Informix extension
- 9 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>column</i>	Column to receive new value	See "Specifying Columns" on page 2-432.	"Identifier " on page 5-21
<i>external</i>	External table into which to insert data	Must exist	"Database Object Name " on page 5-16
<i>field</i>	Field of a named or unnamed ROW data type	Must already be defined in the database	"Field Definition" on page 2-170
<i>position</i>	Position at which to insert an element of a LIST data type	Literal integer or an INT or SMALLINT type SPL variable.	"Literal Number " on page 4-185
<i>synonym, table, view</i>	Table, view, or synonym in which to insert data	Synonym or view and the table to which it points must exist	"Database Object Name " on page 5-16

Usage

To insert data into a table, you must either own the table or have the Insert privilege for the table. (see "GRANT " on page 2-395.) To insert data into a view,

you must have the required Insert privilege, and the view must meet the requirements explained in “Inserting Rows Through a View.”

If the table or view has data integrity constraints, the inserted rows must meet the constraint criteria. If they do not, the database server returns an error. If the checking mode is set to IMMEDIATE, all specified constraints are checked at the end of each INSERT statement. If the checking mode is set to DEFERRED, all specified constraints are *not* checked until the transaction is committed.

Specifying Columns

If you do not explicitly specify one or more columns, data is inserted into columns using the column order that was established when the table was created or last altered. The column order is listed in the **syscolumns** system catalog table.

In Informix ESQL/C, you can use the DESCRIBE statement with an INSERT statement to identify the column order and the data type of the columns in a table.

The number of columns specified in the INSERT INTO clause must equal the number of values supplied in the VALUES clause or by the SELECT statement, either implicitly or explicitly. If you specify a column list, the columns receive data in the order in which you list the columns. The first value following the VALUES keyword is inserted into the first column listed, the second value is inserted into the second column listed, and so on.

If you omit a column from the column list, and the column does not have a default value associated with it, the database server places a NULL value in the column when the INSERT statement is executed.

Using the AT Clause (ESQL/C, IDS, and SPL)

Use the AT clause to insert LIST elements at a specified position in a collection variable. By default, Dynamic Server adds a new element at the end of a LIST collection.

If you specify a position greater than the number of elements in the list, the database server adds the element to the end of the list. You must specify a position value of at least 1 because the first element in the list is at position 1.

The following SPL example inserts a value at a specific position in a list:

```
CREATE PROCEDURE test3()  
  DEFINE a_list LIST(SMALLINT NOT NULL);  
  SELECT list_col INTO a_list FROM table1 WHERE id = 201;  
  INSERT AT 3 INTO TABLE(a_list) VALUES( 9 );  
  UPDATE table1 VALUES list_col = a_list WHERE id = 201;  
END PROCEDURE;
```

Suppose that before this INSERT, **a_list** contained the elements {1,8,4,5,2}. After this INSERT, **a_list** contains the elements {1,8,9,4,5,2}. The new element 9 was inserted at position 3 in the list. For more information on inserting values into collection variables, see “Collection-Derived Table” on page 5-4.

Inserting Rows Through a View

You can insert data through a *single-table* view if you have the Insert privilege on the view. To do this, the defining SELECT statement can select from only one table, and it cannot contain any of the following components:

- DISTINCT keyword

- GROUP BY clause
- Derived value (also referred to as a virtual column)
- Aggregate value

Columns in the underlying table that are unspecified in the view receive either a default value or a NULL value if no default is specified. If one of these columns has no default value, and a NULL value is not allowed, the INSERT fails.

You can use data-integrity constraints to prevent users from inserting values into the underlying table that do not fit the view-defining SELECT statement. For further information, see “WITH CHECK OPTION Keywords” on page 2-277.

With Dynamic Server, you can insert rows through a *single-table* or a *multiple-table* view if an INSTEAD OF trigger specifies valid INSERT operations in its Action clause. See “INSTEAD OF Triggers on Views (IDS)” on page 2-270 for information on how to create INSTEAD OF triggers that insert through views.

If several users are entering sensitive information into a single table, the built-in USER function can limit their view to only the specific rows that each user inserted. The following example contains a view and an INSERT statement that achieves this effect:

```
CREATE VIEW salary_view AS
  SELECT lname, fname, current_salary FROM salary WHERE entered_by = USER;

INSERT INTO salary VALUES ('Smith', 'Pat', 75000, USER);
```

Inserting Rows with a Cursor

In Informix ESQL/C, if you associate a cursor with an INSERT statement, you must use the OPEN, PUT, and CLOSE statements to carry out the INSERT operation. For databases that have transactions but are not ANSI-compliant, you must issue these statements within a transaction.

If you are using a cursor that is associated with an INSERT statement, the rows are buffered before they are written to the disk. The insert buffer is flushed under the following conditions:

- The buffer becomes full.
- A FLUSH statement executes.
- A CLOSE statement closes the cursor.
- In a database that is not ANSI-compliant, an OPEN statement implicitly closes and then reopens the cursor.
- A COMMIT WORK statement ends the transaction.

When the insert buffer is flushed, the client processor performs appropriate data conversion before it sends the rows to the database server. When the database server receives the buffer, it converts any user-defined data types and then begins to insert the rows one at a time into the database. If an error is encountered while the database server inserts the buffered rows into the database, any buffered rows that follow the last successfully inserted rows are discarded.

Inserting Rows into a Database Without Transactions

If you are inserting rows into a database with no transaction logging, you must take explicit action to restore inserted rows if the operation fails. For example, if

INSERT fails after entering some rows, the successfully inserted rows remain in the table. You cannot recover automatically from a failed insert into a database for which no transaction log exists

Inserting Rows into a Database with Transactions

If you are inserting rows into a database and you are using explicit transactions, use the ROLLBACK WORK statement to undo the INSERT. If you do not execute BEGIN WORK before the INSERT, and the INSERT fails, the database server automatically rolls back any data modifications made since the beginning of the INSERT. If you are using an explicit transaction, and the INSERT fails, the database server automatically undoes the effects of the INSERT.

In an ANSI-compliant database, transactions are implicit, and all database modifications take place within a transaction. In this case, if an INSERT statement fails, use the ROLLBACK WORK statement to undo the insertions.

Tables that you create with the RAW logging type are not logged. Thus, raw tables are not recoverable, even if the database uses logging.

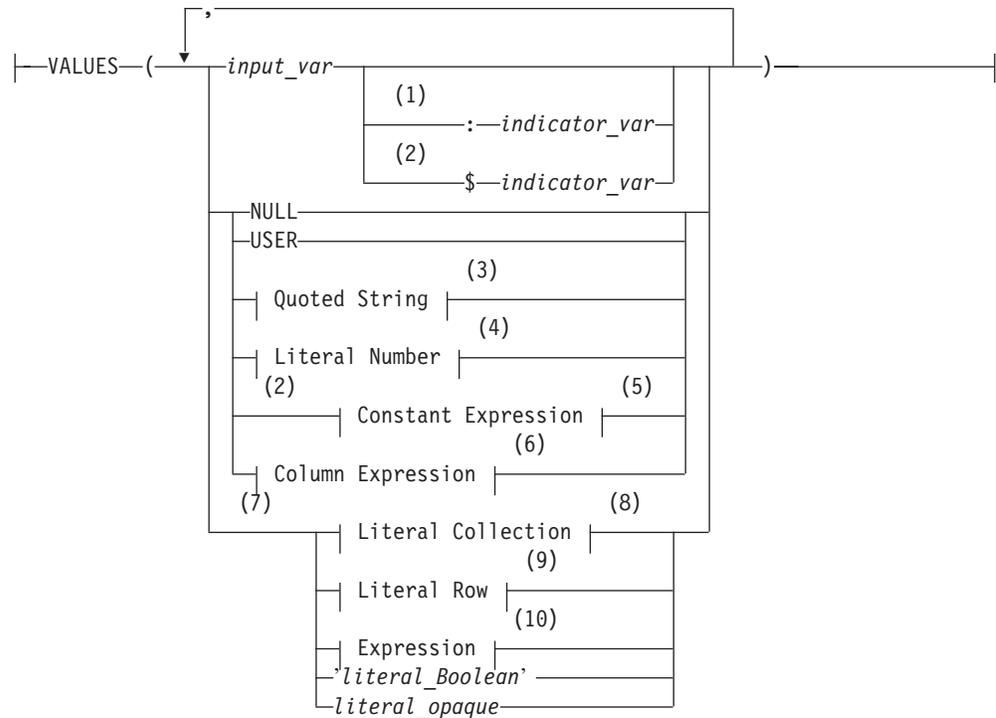
Rows that you insert with a transaction remain locked until the end of the transaction. The end of a transaction is either a COMMIT WORK statement, where all modifications are made to the database, or a ROLLBACK WORK statement, where none of the modifications are made to the database. If many rows are affected by a *single* INSERT statement, you can exceed the maximum number of simultaneous locks permitted. To prevent this situation, either insert fewer rows per transaction, or lock the page (or the entire table) before you execute the INSERT statement.

VALUES Clause

The VALUES clause can specify values to insert into one or more columns. When you use the VALUES clause, you can insert only one row at a time.

Each value that follows the VALUES keyword is assigned to the corresponding column listed in the INSERT INTO clause (or in column order, if a list of columns is not specified). If you are inserting a quoted string into a column, the maximum length that can be inserted without error is 256 bytes.

VALUES Clause:



Notes:

- 1 ESQL/C only
- 2 Informix extension
- 3 See "Quoted String" on page 4-189
- 4 See "Literal Number " on page 4-185
- 5 See "Constant Expressions" on page 4-70
- 6 See "Column Expressions" on page 4-59
- 7 Dynamic Server only
- 8 See "Literal Collection" on page 4-178
- 9 See "Literal Row" on page 4-186
- 10 See "Expression" on page 4-40

Element	Description	Restrictions	Syntax
<i>indicator_var</i>	Variable to show if SQL statement returns NULL to <i>input_var</i>	See the <i>IBM Informix ESQL/C Programmer's Manual</i> .	Language specific
<i>input_var</i>	Variable that holds value to insert. This can be a collection variable.	Can contain any value option of VALUES clause	Language specific
<i>literal_opaque</i>	Literal representation for an opaque data type	Must be recognized by the input support function of the opaque data type	See documentation of the opaque type.
<i>literal_Boolean</i>	Literal representation of a BOOLEAN value as a single character	Either 't' (TRUE) or 'f' (FALSE)	"Quoted String" on page 4-189

In Informix ESQL/C, if you use an *input_var* variable to specify the value, you can insert character strings longer than 256 bytes into a table.

For the keywords and the types of literal values that are valid in the VALUES clause, refer to “Constant Expressions” on page 4-70.

Considering Data Types

The value that the INSERT statement puts into a column does not need to be of the same data type as the column that receives it. These two data types, however, must be compatible. Two data types are *compatible* if the database server has some way to cast one data type to another. A *cast* is the mechanism by which the database server converts one data type to another.

The database server makes every effort to perform data conversion. If the data cannot be converted, the INSERT operation fails. Data conversion also fails if the target data type cannot hold the value that is specified. For example, you cannot insert the integer 123456 into a column defined as a SMALLINT data type because this data type cannot hold a number that large.

For a summary of the casting that the database server provides, see the *IBM Informix Guide to SQL: Reference*. For information on how to create a user-defined cast, see the CREATE CAST statement in this document and *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

Inserting Values into Serial Columns

You can insert successive numbers, explicit values, or explicit values that reset the value in a SERIAL, BIGSERIAL, or SERIAL8 column:

- To insert a consecutive serial value
Specify a zero (0) for the serial column in the INSERT statement. In this case, the database server assigns the next highest value.
- To insert an explicit value
Specify the nonzero value after first verifying that it does not duplicate one already in the table. If the serial column is uniquely indexed or has a unique constraint, and your value duplicates one already in the table, an error results. If the value is greater than the current maximum value, you will create a gap in the series.
- To create a gap in the series (that is, to reset the serial value)
Specify a positive value that is greater than the current maximum value in the column.
Alternatively, you can use the MODIFY clause of the ALTER TABLE statement to reset the next value of a serial column.

For more information, see “Altering the Next Serial Value” on page 2-58.

NULL values are not valid in serial columns.

In Dynamic Server, inserting a serial value into a table that is part of a table hierarchy updates all tables in the hierarchy that contain the serial counter with the value that you insert. You can express this value either as zero (0) for the next highest value, or as a specific positive integer.

Inserting Values into Opaque-Type Columns (IDS)

Dynamic Server supports INSERT operations that specify literal values of opaque data types as quoted strings in the VALUES clause. You can use this syntax to

insert values of opaque UDTs into columns of tables in the local database, or into columns of tables in other databases of the local Dynamic Server instance.

Some opaque data types require special processing when they are inserted. For example, if an opaque data type contains spatial or multirepresentational data, it might provide a choice of how to store the data: inside the internal structure or, for large objects, in a smart large object.

This is accomplished by calling a user-defined support function called **assign()**. When you execute INSERT on a table whose rows contains one of these opaque types, the database server automatically invokes the **assign()** function for the type. The **assign()** function can make the decision of how to store the data. For more information about the **assign()** support function, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

Inserting Values into Collection Columns (IDS)

You can use the VALUES clause to insert values into a collection column. For more information, see “Collection Constructors” on page 4-83.

For example, suppose you define the **tab1** table as follows:

```
CREATE TABLE tab1
(
  int1 INTEGER,
  list1 LIST(ROW(a INTEGER, b CHAR(5)) NOT NULL),
  dec1 DECIMAL(5,2)
);
```

The following INSERT statement inserts a row into **tab1**:

```
INSERT INTO tab1 VALUES
(
  10,
  LIST{ROW(1, 'abcde'),
        ROW(POW(3,3), '=27'),
        ROW(ROUND(ROOT(126)), '=11')},
  100
);
```

The collection column, **list1**, in this example, has three elements. Each element is an unnamed row type with an INTEGER field and a CHAR(5) field. The first element is composed of two literal values, an integer (1) and a quoted string (abcde). The second and third elements also use a quoted string to indicate the second field, but specify the value for the first field with an expression.

Regardless of what method you use to insert values into a collection column, you cannot insert NULL elements into the column. Thus expressions that you use cannot evaluate to NULL. If the collection that you are attempting to insert contains a NULL element, the database server returns an error.

You can also use a collection variable to insert the values of one or more collection elements into a collection column. For more information, see “Collection-Derived Table” on page 5-4.

Inserting Values into ROW-Type Columns (IDS)

Use the VALUES clause to insert values into a named or unnamed ROW type column, as in the following example:

```
CREATE ROW TYPE address_t
(
  street CHAR(20),
```

```

    city CHAR(15),
    state CHAR(2),
    zipcode CHAR(9)
);
CREATE TABLE employee
(
    name ROW ( fname CHAR(20), lname CHAR(20)),
    address address_t
);

```

The next example inserts literal values in the **name** and **address** columns:

```

INSERT INTO employee VALUES
(
    ROW('John', 'Williams'),
    ROW('103 Baker St', 'Tracy', 'CA', 94060)::address_t
);

```

INSERT uses ROW constructors to generate values for the **name** column (an unnamed ROW data type) and the **address** column (a named ROW data type). When you specify a value for a named ROW data type, you must use the CAST AS keywords or the double colon (::) operator, with the name of the ROW data type, to cast the value to the named ROW data type.

For the syntax of ROW constructors, see “Constructor Expressions (IDS)” on page 4-82 in the Expression segment. For information on literal values for named ROW and unnamed ROW data types, see “Literal Row” on page 4-186.

When you use a ROW variable in the VALUES clause, the ROW variable must contain values for each field value. For more information, see “Inserting into a Row Variable (ESQL/C, IDS, SPL)” on page 2-442.

You can use Informix ESQL/C host variables to insert *nonliteral* values in two ways:

- An entire ROW type into a column. Use a **row** variable in the VALUES clause to insert values for all fields in a ROW column at one time.
- Individual fields of a ROW type. To insert nonliteral values in a ROW-type column, insert the elements into a **row** variable and then specify the **collection** variable in the SET clause of an UPDATE statement.

Data Types in Distributed INSERT Operations (IDS)

In Dynamic Server, an INSERT (or any other SQL data-manipulation language statement) that accesses a database of another Dynamic Server instance can reference only the following data types:

- Built-in data types that are not opaque
- BOOLEAN
- LVARCHAR
- DISTINCT of the built-in data types that are not opaque
- DISTINCT of BOOLEAN
- DISTINCT of LVARCHAR
- DISTINCT of any of the DISTINCT data types that appear in this list.

Cross-server distributed INSERT operations can support these DISTINCT types only if the DISTINCT types are cast explicitly to built-in types, and all of the DISTINCT types, their data type hierarchies, and their casts are defined exactly the same way in each participating database.

Cross-server DML operations cannot reference a column or expression of a complex, large-object, nor user-defined data type (UDT), nor of an unsupported DISTINCT or built-in opaque type. For additional information about the data types that Dynamic Server supports in cross-server DML operations, see “Data Types in Cross-Server Transactions” on page 2-537.

Distributed operations that access other databases of the local Dynamic Server instance, however, can access the cross-server data types that are listed above, and also the following data types:

- Most of the *built-in opaque data types*, as listed in “Data Types in Cross-Database Transactions” on page 2-536
- DISTINCT of the built-in types that are referenced in the line above
- DISTINCT of any of the data types that are listed in either of the two lines above
- Opaque user-defined data types (UDTs) that can be cast explicitly to built-in data types.

Cross-database distributed INSERT operations can support these DISTINCT and opaque UDTs only if all the opaque UDTs and DISTINCT types are cast explicitly to built-in types, and all of the opaque UDTs, DISTINCT types, data type hierarchies, and casts are defined exactly the same way in each participating database.

INSERT operations cannot access the database of another database server unless both servers define TCP/IP connections in the DBSERVERNAME or DBSERVERALIASES configuration parameters. This applies to any communication between Dynamic Server instances, even if both database servers reside on the same computer.

Using Expressions in the VALUES Clause

With IBM Informix Dynamic Server, you can insert any type of expression except a column expression into a column. For example, you can insert built-in functions that return the current date, date and time, login name of the current user, or database server name where the current database resides.

The TODAY keyword returns the system date. The CURRENT or SYSDATE keyword returns the system date and time. The USER keyword returns a string that contains the login account name of the current user. The SITENAME or DBSERVERNAME keyword returns the database server name where the current database resides. The following example uses built-in functions to insert data:

```
INSERT INTO cust_calls (customer_num, call_dtime, user_id,  
                      call_code, call_descr)  
VALUES (212, CURRENT, USER, 'L', '2 days');
```

For more information, see “Expression” on page 4-40.

Inserting NULL Values

When you execute the INSERT statement, the database server inserts a NULL value into any column for which you provide no value, as well as for all columns that have no default values and that are not listed explicitly. You also can specify the NULL keyword in the VALUES clause to indicate that a column should be assigned a NULL value.

The following example inserts values into three columns of the **orders** table:

```
INSERT INTO orders (orders_num, order_date, customer_num) VALUES (0, NULL, 123);
```

In this example, a NULL value is explicitly entered in the **order_date** column, and all other columns of the **orders** table that are *not* explicitly listed in the INSERT INTO clause are also filled with NULL values.

Inserting Values into Protected Tables

In a database that uses label-based access control (LBAC), the INTO clause of the INSERT statement can reference a table that is protected by a security policy if the user holds sufficient credentials for the security policy of the label that protects the table, as well as holding the Insert privilege on the table.

A user who holds no security label, however, cannot insert data into a table that has LBAC row protection, even if the user has been granted the required exemptions from rules of the security policy, unless the row label of the protected table is specified in the VALUES clause of the INSERT statement. Data manipulation language statements can provide the row label of a protected table by calling any of three built-in functions whose first argument specifies the name of the security policy, and whose additional arguments are one of the following:

- name of the security label
- name of the IDSSECURITYLABEL column in the table.
- names of the security policy components in the label and the values of their elements

For example, the following INSERT statement calls the built-in **SECLABEL_BY_NAME** function in order to insert a new row into a table called **tab002** that is protected by a row label called **Decca** of the **MegaCorp** security policy:

```
INSERT INTO tab002
  VALUES (SECLABEL_BY_NAME('Megacorp', 'Decca'), 45, 'A.C.Debussy');
```

Whether this INSERT operation succeeds depends on whether the security credentials of the user are sufficient, relative to the component values of the **Decca** label, to enable write access to the **tab002** table.

For additional examples of INSERT statements that access protected tables by calling **SECLABEL_BY_NAME** or similar built-in functions, see “Security Label Support Functions (IDS)” on page 4-119. For general information about LBAC security policies, security labels, read and write access rules, and exemptions from those rules, see your IBM Informix Dynamic Server Security Guide.

Truncated CHAR Values

In a database that is not ANSI-compliant, if you assign a value to a CHAR(*n*) column or variable and the length of that value exceeds *n* characters, the database server truncates the last characters without raising an error. For example, suppose that you define this table:

```
CREATE TABLE tab1 (col_one CHAR(2));
```

The database server truncates the data values in the following INSERT statements to "jo" and "sa" respectively, but does not return a warning:

```
INSERT INTO tab1 VALUES ("john");
INSERT INTO tab1 VALUES ("sally");
```

Thus, in a database that is not ANSI-compliant, the semantic integrity of data for a CHAR(*n*) column or variable is not enforced when the value inserted or updated exceeds the declared length *n*. (But in an ANSI-compliant database, the database server issues error -1279 when truncation of character data occurs.)

Subset of SELECT Statement

As indicated in the diagram for “INSERT ” on page 2-430, not all clauses and options of the SELECT statement are available for you to use in a query within an INSERT statement. The following SELECT clauses and options are not supported by Dynamic Server in an INSERT statement:

- FIRST and INTO TEMP
- ORDER BY and UNION

Extended Parallel Server supports the ORDER BY and UNION options, but not FIRST nor INTO TEMP in the INSERT statement.

In an ANSI-compliant database, if this statement has a WHERE clause that does not return rows, **sqlca** returns SQLNOTFOUND (100).

If an INSERT statement that is part of a multistatement prepared object inserts no rows, **sqlca** returns SQLNOTFOUND (100) for both ANSI-compliant databases and databases that are not ANSI-compliant. In databases that are not ANSI-compliant, **sqlca** returns zero (0) if no rows satisfy the WHERE clause.

In Dynamic Server, if you are inserting values into a supertable in a table hierarchy, the subquery can reference a subtable. If you are inserting values into a subtable in a table hierarchy, the subquery can reference the supertable if it references only the supertable. That is, the subquery must use the SELECT...FROM ONLY (*supertable*)...syntax.

Using External Tables (XPS)

In Extended Parallel Server, a SELECT statement that is a part of a LOAD or UNLOAD operation involving an external table is subject to these restrictions:

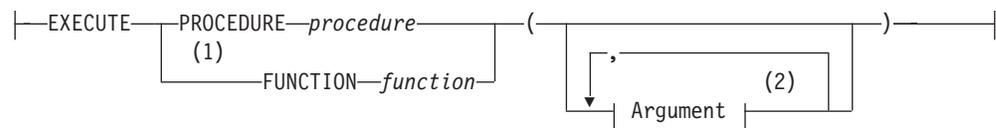
- Only one external table is allowed in the FROM clause.
- The SELECT subquery cannot contain an INTO clause, but it can include any valid SQL expression.

When you move data from a database into an external table, the SELECT statement must define all columns in the external table. The SELECT statement must not contain a FIRST, FOR UPDATE, INTO, INTO SCRATCH, or INTO TEMP clause. You can use an ORDER BY clause, however, to produce files that are ordered within themselves.

Execute Routine Clause

You can specify the EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement to insert values that a user-defined function returns.

Execute Routine Clause:



Notes:

- 1 Dynamic Server only
- 2 See “Arguments ” on page 5-1

Element	Description	Restrictions	Syntax
<i>function, procedure</i>	User-defined function or procedure to insert the data	Must exist	"Database Object Name " on page 5-16

When you use a user-defined function to insert column values, the return values of the function must have a one-to-one correspondence with the listed columns. That is, each value that the function returns must be of the data type expected by the corresponding *column* in the column list.

For backward compatibility, Dynamic Server can use the EXECUTE PROCEDURE keywords to execute an SPL function that was created with the CREATE PROCEDURE statement.

If the called SPL routine scans or updates the target table of the INSERT statement, the database returns an error. That is, the SPL routine cannot select data from the table into which you are inserting rows.

If a called SPL routine contains certain SQL statements, the database server returns an error. For information on which SQL statements cannot be used in an SPL routine that is called within a data manipulation statement, see "Restrictions on SPL Routines in Data-Manipulation Statements" on page 5-78.

Number of Values Returned by SPL, C, and Java Functions

An SPL function can return one or more values. Make sure that the number of returned values matches the number of columns in the table or the number of columns in the column list of the INSERT statement. These columns must have data types that are compatible with the values that the SPL function returns.

An external function written in the C or Java language can only return *one* value. Make sure that you specify only one column in the column list of the INSERT statement. This column must have a compatible data type with the value that the external function returns. The external function can be an iterator function.

The following example shows how to insert data into a temporary table called **result_tmp** in order to output to a file the results of a user-defined function (**f_one**) that returns multiple rows:

```
CREATE TEMP TABLE result_tmp( ... );
INSERT INTO result_tmp EXECUTE FUNCTION f_one();
UNLOAD TO 'file' SELECT * FROM foo_tmp;
```

Inserting into a Row Variable (ESQL/C, IDS, SPL)

The INSERT statement does not support a row variable in the Collection-Derived-Table segment. You can use the UPDATE statement, however, to insert new field values into a row variable. For example, the following Informix ESQL/C code fragment inserts a new row into the **rectangles** table (which "Inserting Values into ROW-Type Columns (IDS)" on page 2-437 defines):

```
EXEC SQL BEGIN DECLARE SECTION;
    row (x int, y int, length float, width float) myrect;
EXEC SQL END DECLARE SECTION;

...
EXEC SQL update table(:myrect)
    set x=7, y=3, length=6, width=2;
EXEC SQL insert into rectangles values (12, :myrect);
```

For more information, see “ Updating a Row Variable (IDS, ESQL/C)” on page 2-703.

Using INSERT as a Dynamic Management Statement

In Informix ESQL/C, you can use the INSERT statement to handle situations where you need to write code that can insert data whose structure is unknown at the time you compile. For more information, refer to the dynamic management section of the *IBM Informix ESQL/C Programmer’s Manual*.

Related Information

Related statements: “CLOSE ” on page 2-77, “CREATE EXTERNAL TABLE (XPS) ” on page 2-104, “DECLARE ” on page 2-286, “DESCRIBE” on page 2-311, “EXECUTE ” on page 2-348, “FLUSH ” on page 2-377, “FOREACH ” on page 3-22, “OPEN ” on page 2-459, “PREPARE ” on page 2-468, “PUT ” on page 2-479, and “SELECT” on page 2-528

For a task-oriented discussion of inserting data into tables and for information on how to access row and collections with SPL variables, see the *IBM Informix Guide to SQL: Tutorial*.

For a discussion of the GLS aspects of the INSERT statement, see the *IBM Informix GLS User’s Guide*.

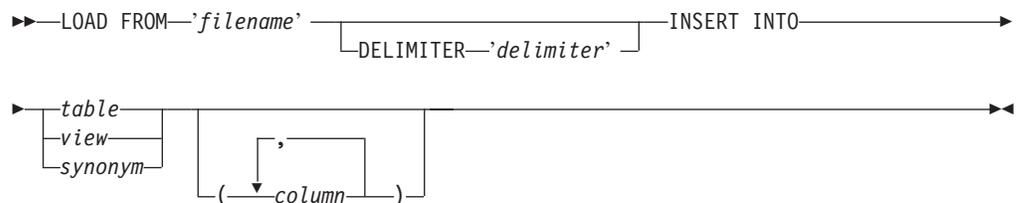
For information on how to access row and collection values with Informix ESQL/C host variables, see the chapter on complex data types in the *IBM Informix ESQL/C Programmer’s Manual*.

LOAD

Use the LOAD statement to insert data from an operating-system file into an existing table or view.

This statement is an extension to the ANSI/ISO standard for SQL. You can use this statement only with DB-Access.

Syntax



Element	Description	Restrictions	Syntax
<i>column</i>	Column to receive data values from	See “INSERT INTO Clause” on page 2-448.	“Identifier ” on page 5-21
<i>delimiter</i>	Character to separate data values in each line of the load file. Default <i>delimiter</i> is the pipe () symbol.	See “DELIMITER Clause” on page 2-448.	“Quoted String” on page 4-189
<i>filename</i>	Path and filename of file to read. Default pathname is current directory	See “LOAD FROM File” on page 2-444.	Specific to operating system rules

Element	Description	Restrictions	Syntax
<i>synonym, table, view</i>	Synonym for the table in which to insert data from <i>filename</i>	<i>Synonym</i> and <i>table</i> or <i>view</i> to which it points must exist	"Database Object Name" " on page 5-16

Usage

The LOAD statement appends new rows to the table. It does not overwrite existing data. You cannot add a row that has the same key as an existing row.

To use the LOAD statement, you must have Insert privileges for the table where you want to insert data. For information on database-level and table-level privileges, see the GRANT statement.

LOAD FROM File

The LOAD FROM file contains the data to be loaded into the specified table or view. The default pathname for the load file is the current directory.

You can use the file that the UNLOAD statement creates as the LOAD FROM file. (See "UNLOAD TO File" on page 2-684 for a description of how values of various data types are represented within the UNLOAD TO file.)

If you do not include a list of columns in the INSERT INTO clause, the fields in the file must match the columns that are specified for the table in number, order, and data type.

Each line of the file must have the same number of fields. You must define field lengths that are less than or equal to the length that is specified for the corresponding column. Specify only values that can convert to the data type of the corresponding column. The following table indicates how the database server expects you to represent the data types in the LOAD FROM file (when you use the default locale, U.S. English).

Type of Data	Input Format
blank	One or more blank characters between delimiters You can include leading blanks in fields that do not correspond to character columns.
BOOLEAN	A t or T indicates a TRUE value, and an f or F indicates a FALSE value.
COLLECTIONS	Collection must have its values surrounded by braces ({ }) and a field delimiter separating each element. For more information, see " Loading Complex Data Types (IDS)" on page 2-447.
DATE	Character string in the following format: <i>mm/dd/year</i> You must state the month as a two-digit number. You can use a two-digit number for the year if the year is in the 20th century. (You can specify another century algorithm with the DBCENTURY environment variable.) The value must be an actual date; for example, February 30 is illegal. You can use a different date format if you indicate this format with the GL_DATE or DBDATE environment variable. For more information about environment variables, see the <i>IBM Informix Guide to SQL: Reference</i> and the <i>IBM Informix GLS User's Guide</i> .

Type of Data	Input Format
DECIMAL, MONEY, FLOAT	Value that can include a leading and/or trailing currency symbol and thousands and decimal separators Your locale files or the DBMONEY environment variable can specify a currency format.
NULL	Nothing between the delimiters
ROW types (named or unnamed)	ROW type must have its values surrounded by parentheses and a field delimiter that separates each element. For more information, see “ Loading Complex Data Types (IDS)” on page 2-447.
Simple large objects (TEXT, BYTE)	TEXT and BYTE columns are loaded directly from the LOAD TO file. For more information, see “Loading Simple Large Objects” on page 2-446.
Smart large objects (CLOB, BLOB)	CLOB and BLOB columns are loaded from a separate operating-system file. The field for the CLOB or BLOB column in the LOAD FROM file contains the name of this separate file. For more information, see “Loading Smart Large Objects (IDS)” on page 2-447.
Time	Character string in <i>year-month-day hour:minute:second.fraction</i> format You cannot use data type keywords or qualifiers for DATETIME or INTERVAL values. The year must be a 4-digit number, and the month must be a 2-digit number. The DBTIME or GL_DATETIME environment variable can specify other formats.
User-defined data formats (opaque types)	Associated opaque type must have an import support function defined if special processing is required to copy the data in the LOAD FROM file to the internal format of the opaque type. An import binary support function might also be required for data in binary format. The LOAD FROM file data must be in the format that the import or import binary support function expects. The associated opaque type must have an assign support function if special processing is required before writing the data in the database. See “ Loading Opaque-Type Columns (IDS)” on page 2-448.

For more information on **DB*** environment variables, refer to the *IBM Informix Guide to SQL: Reference*. For more information on **GL*** environment variables, refer to the *IBM Informix GLS User’s Guide*.

If you are using a nondefault locale, the formats of DATE, DATETIME, MONEY, and numeric column values in the LOAD FROM file must be compatible with the formats that the locale supports for these data types. For more information, see the *IBM Informix GLS User’s Guide*.

The following example shows the contents of an input file named **new_custs**:

```
0|Jeffery|Padgett|Wheel Thrills|3450 El Camino|Suite 10|Palo Alto|CA|94306||
0|Linda|Lane|Palo Alto Bicycles|2344 University||Palo Alto|CA|94301|
(415)323-6440
```

This data file conveys the following information:

- Indicates a serial field by specifying a zero (0)
- Uses the pipe (|), the default delimiter
- Assigns NULL values to the **phone** field for the first row and the **address2** field for the second row

The NULL values are shown by two delimiters with nothing between them.

The following statement loads the values from the **new_custs** file into the **customer** table that **jason** owns:

```
LOAD FROM 'new_custs' INSERT INTO jason.customer;
```

If you include any of the following special characters as part of the value of a field, you must precede the character with a backslash (\) escape symbol:

- Backslash
- Delimiter
- Newline character anywhere in the value of a VARCHAR or NVARCHAR column
- Newline character at end of a value for a TEXT value

Do not use the backslash character (\) as a field separator. It serves as an escape character to inform the LOAD statement that the next character is to be interpreted as part of the data, rather than as having special significance.

Fields that correspond to character columns can contain more characters than the defined maximum allows for the field. The extra characters are ignored.

If you are loading files that contain VARCHAR data types, note the following information:

- If you give the LOAD statement data in which the character fields (including VARCHAR) are longer than the column size, the excess characters are disregarded.
- Use the backslash (\) to escape embedded delimiter and backslash characters in all character fields, including VARCHAR.
- Do not use the following characters as delimiting characters in the LOAD FROM file: digits (0 to 9), the letters a to f, and A to F, the backslash (/) character, or the NEWLINE character.

Loading Simple Large Objects

The database server loads simple large objects (BYTE and TEXT columns) directly from the LOAD FROM file. Keep the following restrictions in mind when you load BYTE and TEXT data:

- You cannot have leading and trailing blanks in BYTE fields.
- Use the backslash (\) to escape the special significance of literal delimiter and backslash characters in TEXT fields.
- Data being loaded into a BYTE column must be in ASCII-hexadecimal form. BYTE columns cannot contain preceding blanks.
- Do not use the following characters as delimiting characters in the LOAD FROM file: digits (0 to 9), the letters a to f, and A to F, the backslash (/) character, or the NEWLINE character.

For loading TEXT columns in a non-default locale, the database server handles any required code-set conversions for the data. See also the *IBM Informix GLS User's Guide*.

If you are unloading files that contain BYTE or TEXT data types, objects smaller than 10 kilobytes are stored temporarily in memory. You can adjust the 10-kilobyte setting to a larger setting with the **DBBLOBBUF** environment variable. Simple large objects that are larger than the default or the setting of **DBBLOBBUF** are

stored in a temporary file. For more information about the **DBBLOBBUF** environment variable, see the *IBM Informix Guide to SQL: Reference*.

Loading Smart Large Objects (IDS)

The database server loads smart large objects (BLOB and CLOB columns) from a separate operating-system file on the client computer. For information on the structure of this file, see “Unloading Smart Large Objects (IDS)” on page 2-686.

In a LOAD FROM file, a CLOB or BLOB column value appears as follows:

```
start_off,length,client_path
```

In this format, *start_off* is the starting offset (in hexadecimal) of the smart-large-object value within the client file, *length* is the length (in hexadecimal) of the BLOB or CLOB value, and *client_path* is the pathname for the client file. No blank spaces can appear between these values.

For example, to load a CLOB value that is 512 bytes long and is at offset 256 in the **/usr/apps/clob9ce7.318** file, the database server expects the CLOB value to appear as follows in the LOAD FROM file:

```
|100,200,/usr/apps/clob9ce7.318|
```

If the whole client file is to be loaded, a CLOB or BLOB column value appears as follows in the LOAD FROM file:

```
client_path
```

For example, to load a CLOB value that occupies the entire file **/usr/apps/clob9ce7.318**, the database server expects the CLOB value to appear as follows in the LOAD FROM file:

```
|/usr/apps/clob9ce7.318|
```

In DB-Access, the USING clause is valid within files executed from DB-Access. In interactive mode, DB-Access prompts you for a password, so the USING keyword and *validation_var* are not used.

For CLOB columns, the database server handles any required code-set conversions for the data. See also the *IBM Informix GLS User's Guide*.

Loading Complex Data Types (IDS)

In a LOAD FROM file, complex data types appear as follows:

- Collections are introduced with the appropriate constructor (SET, MULTISSET, or LIST), and their elements are enclosed in braces ({}) and separated with a comma, as follows:

```
constructor{val1 , val2 , ... }
```

For example, to load the SET values {1, 3, 4} into a column whose data type is SET(INTEGER NOT NULL), the corresponding field of the LOAD FROM file appears as:

```
|SET{1 , 3 , 4}|
```

- Row types (named and unnamed) are introduced with the ROW constructor and their fields are enclosed with parentheses and separated with a comma, as follows:

```
ROW(val1 , val2 , ... )
```

For example, to load the ROW values (1, 'abc'), the corresponding field of the LOAD FROM file appears as:

```
|ROW(1 , abc)|
```

Loading Opaque-Type Columns (IDS)

Some opaque data types require special processing when they are inserted. For example, if an opaque data type contains spatial or multirepresentational data, it might provide a choice of how to store the data: inside the internal structure or, for large objects, in a smart large object.

This processing is accomplished by calling a user-defined support function called **assign()**. When you execute the LOAD statement on a table whose rows contain one of these opaque types, the database server automatically invokes the **assign()** function for the type. The **assign()** function can make the decision of how to store the data. For more information about the **assign()** support function, see the *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

DELIMITER Clause

Use the DELIMITER clause to specify the delimiter that separates the data contained in each column in a row in the input file. You can specify TAB (CTRL-I) or a blank space (= ASCII 32) as the delimiter symbol. You cannot use the following items as the delimiter symbol:

- Backslash (\)
- NEWLINE character (CTRL-J)
- Hexadecimal numbers (0 to 9, a to f, A to F)

If you omit this clause, the database server checks the **DBDELIMITER** environment variable. For information about how to set the **DBDELIMITER** environment variable, see the *IBM Informix Guide to SQL: Reference*.

If the **DBDELIMITER** environment variable has not been set, the default delimiter is the pipe (|).

The following example specifies the semicolon (;) as the delimiting character. The example uses Windows file-naming conventions.

```
LOAD FROM 'C:\data\loadfile' DELIMITER ';'
INSERT INTO orders;
```

INSERT INTO Clause

Use the INSERT INTO clause to specify the table, synonym, or view in which to load the new data. You must specify the column names only if one of the following conditions is true:

- You are not loading data into all columns.
- The input file does not match the default order of the columns (the order specified when the table was created).

The following example identifies the **price** and **discount** columns as the only columns in which to add data. The example uses Windows file naming conventions.

```
LOAD FROM 'C:\tmp\prices' DELIMITER ','
INSERT INTO norman.worktab(price,discount)
```

Related Information

Related statements: “UNLOAD ” on page 2-683 and “INSERT ” on page 2-430

For a task-oriented discussion of the LOAD statement and other utilities for moving data, see the *IBM Informix Migration Guide*.

For a discussion of the GLS aspects of the LOAD statement, see the *IBM Informix GLS User's Guide*.

LOCK TABLE

Use the LOCK TABLE statement to control access to a table by other processes.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax

```
▶▶ LOCK TABLE owner. table IN SHARE MODE ▶▶  
                  └─synonym┘          └─EXCLUSIVE┘
```

Element	Description	Restrictions	Syntax
<i>owner</i>	Owner of <i>synonym</i> or <i>table</i>	Must be the owner of the specified object	“ Owner Name ” on page 5-45
<i>synonym</i>	Synonym for the table to be locked	Synonym and the table to which it points must exist	“Identifier ” on page 5-21
<i>table</i>	Table to be locked	See first paragraph of “Usage.”	“Identifier ” on page 5-21

Usage

You can use LOCK TABLE to lock a table if either of the following is true:

- You are the owner of the table.
- You have Select privilege on the table or on a column in the table, either from a direct grant or from a grant to PUBLIC or to your current role.

The LOCK TABLE statement fails if the table is already locked in EXCLUSIVE mode by another process, or if you request an EXCLUSIVE lock while another user has locked the same table in SHARE mode.

The SHARE keyword locks a table in *shared mode*. Shared mode gives other processes *read* access to the table but denies *write* access. Other processes cannot update or delete data if a table is locked in shared mode.

The EXCLUSIVE keyword locks a table in *exclusive mode*. This mode denies other processes both read and write access to the table. Exclusive-mode locking automatically occurs during the ALTER INDEX, ALTER TABLE, CREATE INDEX, DROP INDEX, RENAME COLUMN, RENAME TABLE, START VIOLATIONS TABLE, STOP VIOLATIONS TABLE, and TRUNCATE statements.

Concurrent Access to Tables with Shared Locks

After the LOCK TABLE statement that specifies the IN SHARE MODE keywords executes successfully, other users can read the table but cannot modify its data until the lock is released. In databases that support transaction logging, the SELECT statement can implicitly place a shared lock on each tables that is listed in the FROM clause, in order to prevent other users from modifying those tables until the query is committed or rolled back.

Concurrent Access to Tables with Exclusive Locks

After the LOCK TABLE statement with the IN EXCLUSIVE MODE option executes successfully, no other user can obtain a lock on the specified table. When you attempt a DDL operation on that table, however, you might receive RSAM error -106 if the same table is being accessed by a concurrent session (for example, by opening a cursor). This error can also affect implicit locks that certain DDL statements place on tables automatically.

This is possible because table locks do not preclude table access. An exclusive lock prevents other users from obtaining a lock, but it cannot prevent them from opening the table for write operations that wait for the exclusive lock to be released, or for Dirty Read operations on the table. You can set the **IFX_DIRTY_WAIT** environment variable to specify that the DDL operation wait for a specified number of seconds for Dirty Read operations to commit or rollback.

When one or more rows in a table are locked by an exclusive lock, the effect on other users partly depends on their transaction isolation level. Other users in all isolation levels except the Dirty Read isolation level might encounter locking errors, such as transactions that fail because the lock was not released within a specified time limit, or a deadlock situation.

On tables where row-level locking affects some of the rows, the risk of locking conflicts can be reduced by enabling transactions to read the most recently committed version of the data in the row-level locked table, rather than waiting for the transaction that holds the lock on that row to be committed or rolled back. This can be accomplished in several different ways, including these:

- From an individual session, issue this SQL statement
`SET ISOLATION TO COMMITTED READ LAST COMMITTED;`
- For all sessions using Committed Read or Read Committed isolation levels, set the USELASTCOMMITTED configuration parameter to 'ALL' or to 'COMMITTED READ', or else issue the SET ENVIRONMENT USELASTCOMMITTED statement with 'ALL' or 'COMMITTED READ' as the session environment option.
- For all sessions using Dirty Read or Read Uncommitted isolation levels, set the USELASTCOMMITTED configuration parameter to 'ALL' or to 'DIRTY READ', or else issue the SET ENVIRONMENT USELASTCOMMITTED statement with 'ALL' or 'DIRTY READ' as the session environment option.
- For users for whom a *user.sysdbopen()* procedure is defined in the database, the DBA can define that procedure to include the SET ENVIRONMENT USELASTCOMMITTED statement with 'ALL' or 'COMMITTED READ' as the session environment option, and also issue the SET ISOLATION statement to set Committed Read as the isolation level.
- For users for whom no *user.sysdbopen()* procedure exists in the database, the DBA can define a *PUBLIC.sysdbopen* procedure that specifies the same SET ENVIRONMENT USELASTCOMMITTED and SET ISOLATION statements.

This LAST COMMITTED isolation feature is useful only when row-level locking is in effect, rather than when another session holds an exclusive lock on the entire table. This feature is disabled for the specified table when LOCK TABLE applies a table-level lock. See “The LAST COMMITTED Option to Committed Read” on page 2-634 for more information about this LAST COMMITTED feature for concurrent access to tables in which some rows are locked by exclusive locks, and for restrictions on the kinds of tables that can support this feature.

Databases with Transactions

If your database was created with transactions, the LOCK TABLE statement succeeds only if it executes within a transaction. You must issue a BEGIN WORK statement before you can execute a LOCK TABLE statement.

Transactions are implicit in an ANSI-compliant database. The LOCK TABLE statement succeeds if the specified table is not already locked by another process.

The following guidelines apply to the use of the LOCK TABLE statement within transactions:

- You cannot lock system catalog tables.
- You cannot switch between shared and exclusive table locking within a transaction. For example, once you lock the table in shared mode, you cannot upgrade the lock mode to exclusive.
- If you issue a LOCK TABLE statement before you access a row in the table, and PDQ is not in effect, no row locks are set for the table. In this way, you can override row-level locking and avoid exceeding the maximum number of locks that are defined in the database server configuration. (But if PDQ is not in effect, you might run out of locks with error -134 unless the LOCKS parameter of your ONCONFIG file specifies a large enough number of locks.)
- All row and table locks release automatically after a transaction is completed. The UNLOCK TABLE statement fails in a database that uses transaction logging.
- The same user can explicitly use LOCK TABLE to lock up to 32 tables concurrently. (Use SET ISOLATION to specify an appropriate isolation level, such as Repeatable Read, if you need to lock rows from more than 32 tables during a single transaction.)

The following example shows how to change the locking mode of a table in a database that was created with transaction logging:

```
BEGIN WORK;  
LOCK TABLE orders IN EXCLUSIVE MODE;  
...  
COMMIT WORK;  
BEGIN WORK;  
LOCK TABLE orders IN SHARE MODE;  
...  
COMMIT WORK;
```

Warning: It is recommended that you not use nonlogging tables in a transaction. If you need to use a nonlogging table in a transaction, either lock the table in exclusive mode or set the isolation level to Repeatable Read to prevent concurrency problems.

Databases Without Transactions

In a database that was created without transactions, table locks that were set by the LOCK TABLE statement are released after any of the following events:

- An UNLOCK TABLE statement executes.
- The user closes the database.
- The user exits the application.

To change the lock mode on a table, release the lock with the UNLOCK TABLE statement and then issue a new LOCK TABLE statement.

The following example shows how to change the lock mode of a table in a database that was created without transactions:

```
LOCK TABLE orders IN EXCLUSIVE MODE;
.
.
.
UNLOCK TABLE orders;
.
.
.
LOCK TABLE orders IN SHARE MODE;
```

Locking Granularity

The default granularity for locking a table is at the *page* level, or whatever you specify (either PAGE or ROW) in the **IFX_TABLE_LOCKMODE** environment variable, or if that is not set, by setting DEF_TABLE_LOCKMODE in the ONCONFIG file. The LOCK MODE clause of the CREATE TABLE or ALTER TABLE statement can override the default locking granularity by specifying PAGE, ROW, or (for Extended Parallel Server only) TABLE for a specified table. Only row-level locks support the LAST COMMITTED feature of Dynamic Server.

The LOCK TABLE statement, however, always locks the entire table, overriding any other locking granularity specification for the table. For Extended Parallel Server, the LOCK MODE clause of the CREATE INDEX statement can specify COARSE or NORMAL locking granularity for an index.

In all of these contexts, the term "lock mode" means the *locking granularity*. In the context of the SET LOCK MODE statement, however, "lock mode" refers to the behavior of the database server when a process attempts to access a row or a table that another process has locked.

Related Information

Related statements: "BEGIN WORK " on page 2-74, "COMMIT WORK " on page 2-81, "SAVE EXTERNAL DIRECTIVES " on page 2-522, "SET ISOLATION" on page 2-631, "SET LOCK MODE" on page 2-637, and "UNLOCK TABLE " on page 2-688

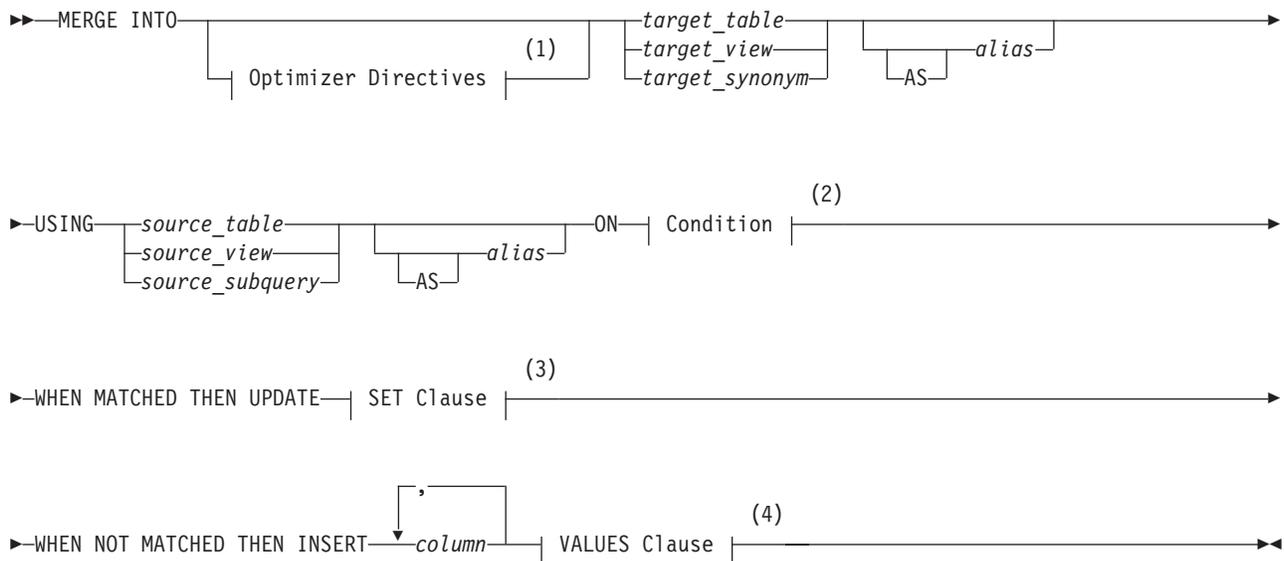
For a discussion of concurrency and locks, see the *IBM Informix Guide to SQL: Tutorial*.

MERGE

The MERGE statement provides an efficient way to merge two tables by combining UPDATE and INSERT operations into a single statement.

Only Extended Parallel Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 See "Optimizer Directives" on page 5-35
- 2 See "Condition" on page 4-5
- 3 See "SET Clause" on page 2-693
- 4 See "VALUES Clause" on page 2-434

Element	Description	Restrictions	Syntax
<i>alias</i>	Temporary name that you declare here for the <i>target</i> table object or <i>source</i> table object	If potentially ambiguous, you must precede <i>alias</i> with the AS keyword	"Identifier" on page 5-21
<i>column</i>	Name of a column in the <i>target</i> table object in which to insert <i>source</i> data	Must exist in the <i>target</i> object	"Identifier" on page 5-21
<i>source_table</i> , <i>source_view</i> , <i>source_subquery</i>	Name of a table object (or the result of a query) that contains the data that you are relocating	Must exist	"Database Object Name" on page 5-16; SELECT, p. "SELECT" on page 2-528
<i>target_table</i> , <i>target_view</i> , <i>target_synonym</i>	Name or synonym of a table or updatable view in which to insert the <i>source</i> data	See "Target Table Restrictions" on page 2-454	"Database Object Name" on page 5-16

Usage

The MERGE statement enables you to merge data of a source table into a destination table, here called the *target table*. It provides an efficient way to update and insert the rows of the source table into the target table, based on search conditions that you specify within the statement.

The *target* object can be in a different database from the *source* object, but it must be a database managed by the currently running Extended Parallel Server instance.

The *source* object, however, can be in a different database from the *target* object, and need not be a database managed by the current running Extended Parallel Server instance.

Rows in the *target* object that match the search conditions are updated; otherwise, new rows are inserted into the *target* object.

The MERGE statement combines some of the functionality of the UPDATE and INSERT statements into a single statement. The SET clause of this statement is identical to that of the UPDATE statement. For more information, see “SET Clause” on page 2-693.

The VALUES clause of the MERGE is similar to the INSERT statement, but only accepts column names and constant values, not expressions. For more information, see “VALUES Clause” on page 2-434.

The following example uses the transaction table **new_sale** to merge rows into the fact table **sale**, updating **sale_count** if there are already records in the **sale** table. If not, the rows are inserted into the **sale** table.

```
MERGE INTO sale USING new_sale AS n ON sale.cust_id = n.cust_id
  WHEN MATCHED THEN UPDATE SET sale.salecount = sale.salecount + n.salecount
  WHEN NOT MATCHED THEN INSERT (cust_id, salecount)
  VALUES (n.cust_id,n.salecount);
```

If an error occurs while MERGE is executing, the entire statement is rolled back.

If the constraints checking mode for the target table is set to IMMEDIATE, then unique and referential constraints of the target object are checked after all the UPDATE and INSERT operations are complete. The NOT NULL and check constraints are checked during the UPDATE and INSERT operations. If the checking mode is set to DEFERRED, the constraints are not checked until after the transaction is committed.

Target Table Restrictions

The following restrictions apply to the target table of the MERGE statement. If any of the following conditions exist, MERGE returns an error.

- The target table of MERGE cannot have any violations tables defined
- The target table of MERGE cannot be an external table nor a remote table
- The target table of MERGE cannot be a system catalog table
- The target table of MERGE cannot be a static table nor a duplicated table
- The target table of MERGE cannot be a pseudo-table (a memory-resident object in a system database) nor a read-only view
- If an Update or Insert trigger is defined on the target table of MERGE, an error is issued. (Here MERGE does not act as a triggering event.)

Handling Duplicate Rows

While MERGE is executing, the same row in the target table cannot be updated more than once. Inserted rows are not updated by the MERGE statement. Using the previous example, if the **sale** table contains the two records shown at the left, then the **new_sale** table contains the three records shown at the right:

Records in 'sale' Table		Records in 'new_sale' Table	
cust_id	sale_count	cust_id	sale_count
Tom	129	Tom	20
Julie	230	Julie	3
		Julie	10

When merging **new_sale** into **sale** using the expression **sale.cust_id = new_sale.cust_id**, the MERGE statement returns an error, because it attempts to update one of the records in the source table more than once.

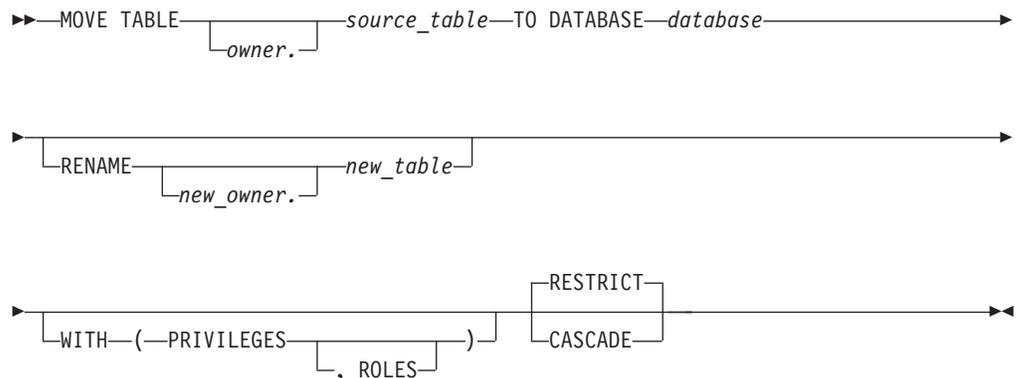
Related Information

Related statements: "INSERT " on page 2-430 and "UPDATE" on page 2-689

MOVE TABLE

The MOVE TABLE statement moves the schema of a database table to another database of the same database server instance. Only Extended Parallel Server supports this syntax, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Element	Description	Restrictions	Syntax
<i>database</i>	Name of the destination database	Must exist and must have the same database server as the source database	"Database Name" on page 5-15
<i>owner</i>	Current table owner	Must be the owner of <i>source_table</i>	"Owner Name" on page 5-45
<i>new_owner</i>	Owner designated here for <i>new_table</i>	Must be a valid owner name in the destination <i>database</i>	"Owner Name" on page 5-45
<i>new_table</i>	Name declared here for relocated table	Must be unique among table, view, and synonym names in destination <i>database</i>	"Identifier" on page 5-21
<i>source_table</i>	Current table name	Must exist in the source database	"Identifier" on page 5-21

Usage

The MOVE TABLE statement provides a way to move a permanent table from one database to another on the same Extended Parallel Server instance. When the table

is moved, many of its objects, including any indexes and fragmentation strategy, are preserved, but the data values in the table are not physically moved to the destination database. The database server transfers only the system catalog information describing *source_table* to the destination database. For this reason, the performance of MOVE TABLE depends on the time that is required to transfer the system catalog data, but not on the size of *source table*.

If indexes exist on *source table*, the data values of the index are not moved. Only the system catalog information is moved to the destination database. By default, all roles assigned to *source table* are dropped from the source database.

The destination table has, by default, restricted privileges. Only the owner of *source table* has privileges on *new table*. Also by default, all roles assigned to *source table* are dropped from *new table*. In the following basic example, the privileges of *source table* are not preserved in the destination database:

```
MOVE TABLE employee TO DATABASE payroll;
```

You can preserve the privileges and roles of the original table using the WITH clause. For more information, see “Preserving Table Privileges” on page 2-458.

Only users with DBA privileges on both the source and destination databases can use the MOVE TABLE statement.

The following types of tables cannot be moved using the MOVE TABLE statement:

- System catalog tables
- Temporary tables (TEMP or SCRATCH)
- Tables not currently in the database
- Active violations tables
- Tables appearing in the FROM clause of a GK index
- System database tables (tables in **sysmaster** and **sysutils**)

If the destination database already contains a table with the same name as *source table*, the MOVE TABLE statement fails with an error. Similarly, if any dependent objects of source table have names that already exist in the destination database, the MOVE TABLE statement fails and returns an error.

The source table can be either fragmented or non-fragmented. You must keep the dbspace and fragmentation schema in mind, however, because the data values of the destination table still reside in the same dbspace. MOVE TABLE does not provide a mechanism to move data from one dbspace to another.

The MOVE TABLE statement can only move tables between databases. Use the RENAME TABLE statement to move a table within the same database.

Considerations When Using the MOVE TABLE Statement

All user-defined indexes on *source table*, including detached and globally detached indexes, are preserved in the destination table. All system-generated indexes are dropped, except those whose constraints are preserved. For more information on constraints, see “Moving Tables with Constraints” on page 2-457.

By default, if any views exist that depend on *source table*, MOVE TABLE fails with an error. But if you specify CASCADE mode, then those views are dropped and MOVE TABLE succeeds. Remote views referring to *source table* are not checked. This functionality is similar to that of the DROP TABLE statement.

Like views, remote synonyms for *source table* are not checked. Any statement that refers to such synonyms after *source table* is moved returns an error message, indicating that the specified table does not exist.

Stored procedures referring to *source table* are not modified. If a stored procedure referencing *source table* is executed after the table is moved, an error is returned saying the table does not exist.

All triggers associated with the source table are dropped.

If the source table has duplicate tables defined, the duplicate tables are dropped when the table is moved.

All local synonyms of *source table* are dropped when the table is moved. This includes synonyms created on the views that are dependent on *source table*.

All column distributions of *source table* are preserved in the destination table. Entries for *source table* in the **sysdistrib** system catalog table are moved. This ensures that the destination table has the same statistics as the source table. For this reason, you do not need to run the UPDATE STATISTICS statement after moving a table.

A *source table* of the MOVE TABLE statement receives a new **tabid** as if it were a newly created table. This means that system-defined names of objects that are dependent on *source table* are also affected. User-defined object names are preserved after the table is moved.

Moving Tables with Constraints

The following table-level constraints are preserved in MOVE TABLE operations:

- UNIQUE
- NOT NULL
- CHECK
- PRIMARY KEY

All source table defaults are also preserved.

If you specify CASCADE mode, all referential constraints associated with *source table* are dropped when the MOVE TABLE statement executes. As part of this, all system-generated indexes for these constraints are also dropped. If an index has been explicitly created, however, it is preserved in the destination table.

By default, MOVE TABLE statement returns an error if any existing referential constraint references *source table*.

If a child table of *source table* has ON DELETE CASCADE set, the data of the child table is not affected, but its corresponding referential constraint is dropped. This behavior is the same as with the DROP TABLE statement.

By default, if *source table* is a parent table, then the MOVE TABLE statement fails. If *source table* is a child table, however, the MOVE TABLE operation succeeds.

Renaming the Table

If you want to designate a different identifier or owner for the destination table, you can specify the RENAME clause of MOVE TABLE. If the destination database

already has that identifier registered for any table, view or synonym, MOVE TABLE fails and returns an error. Also, if the new owner specified in this clause is an existing role name in the destination database, MOVE TABLE fails and returns an error. The RENAME clause that also specifies a new owner name does not change owner of the table objects, such as indexes or constraints. MOVE TABLE maintains the original owners of those objects after the table has moved. That is, the RENAME clause renames only the table, and not its objects.

Preserving Table Privileges

By default, no privileges that users or roles hold on *source table* persist on the new table. You can preserve these privileges in the new table, however, by using the WITH (PRIVILEGES) and WITH (PRIVILEGES, ROLES) clauses.

The WITH (PRIVILEGES) Option

This clause moves all access privileges on *source table* to the destination table. All entries for *source table* in the **sysstabaauth** and **syscolauth** system catalog tables are transferred to *new table* in the destination database. If these entries specify a role, however, the role is ignored, and is also dropped from the source database. The following MOVE TABLE example preserves the user privileges of *source table* and changes the owner of *new table*:

```
MOVE TABLE customer TO DATABASE sales RENAME mary.customer WITH (PRIVILEGES);
```

The WITH (PRIVILEGES, ROLES) Option

This clause specifies that all privileges on *source table* to be moved to *new table*, including any privileges of roles. All entries for *source table* in the **sysstabaauth** and **syscolauth** system catalog tables are transferred to the destination table. The following example preserves both user and role privileges in *new table*:

```
MOVE TABLE mytab TO DATABASE newdb WITH (PRIVILEGES, ROLES);
```

The MOVE TABLE statement returns an error if there is a conflict between user names and role names in the source and destination databases. If any roles that are preserved in the destination table are not defined in the destination database, an error is returned. It is the responsibility of the DBA to ensure that role names and user names are correctly defined in the source and destination databases.

The RESTRICT and CASCADE Options

You can use one of the RESTRICT or CASCADE keywords to specify how the MOVE TABLE statement treats dependencies among database objects. If you include neither of these keywords, the default mode is RESTRICT.

Specifying CASCADE Mode

The CASCADE keyword of MOVE TABLE removes the following database objects, if they exist, that are related to *source table*:

- Referential constraints that are defined on the table
- Views that are defined on the table
- Any violations table or diagnostics table that is associated with the table

Specifying RESTRICT Mode

The RESTRICT keyword has no effect unless *source table* has one or more referential constraints or views defined on it, or unless a violations or diagnostics table is associated with it. If you explicitly specify RESTRICT (or equivalently, if you do not specify CASCADE, thereby making RESTRICT the default mode), the MOVE TABLE operation fails with an error if any of the following conditions are true for *source table*:

- One or more referential constraints are defined on the table
- One or more views are defined on the table
- A violations table or diagnostics table is associated with the table

Unlike the DROP TABLE statement, whose default mode is CASCADE, the RESTRICT mode is the default mode for the MOVE TABLE statement.

Physically Moving Data to Another Database

MOVE TABLE cannot be used if you need to physically move the table data from one database to another. In this case, you need to perform the following steps:

1. Use CREATE TABLE to create a new table, with an appropriate schema, in the destination database.
2. Use the INSERT INTO ... SELECT FROM statement to transfer the data from the source table to the destination table
3. Use the DROP TABLE statement to drop the original table.

You must also recreate any table objects, including indexes and fragmentation strategies, and verify that there is sufficient disk space to move the data.

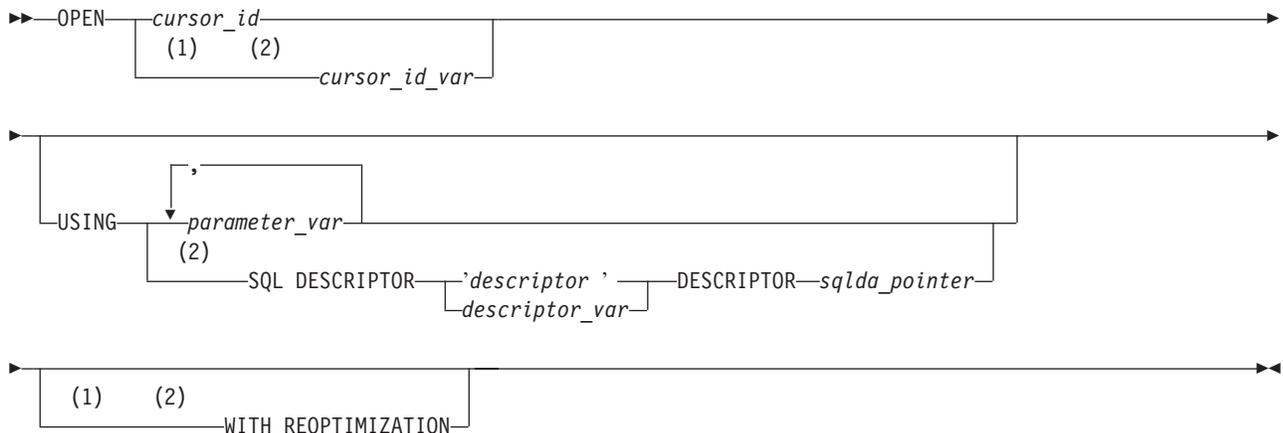
Related Information

Related statements: "CREATE TABLE " on page 2-191, "DROP TABLE " on page 2-342, "INSERT " on page 2-430, "RENAME TABLE" on page 2-492

OPEN

Use the OPEN statement to activate a cursor. Use this statement with Informix ESQL/C or with SPL.

Syntax



Notes:

- 1 Informix extension
- 2 ESQL/C only

Element	Description	Restrictions	Syntax
<i>cursor_id</i>	Name of a cursor	Must have been declared by the DECLARE statement	"Identifier " on page 5-21
<i>cursor_id_var</i>	Host variable = <i>cursor_id</i>	Must be a character data type	Language specific
<i>descriptor</i>	Name of a system-descriptor area	Must have been allocated	"Quoted String" on page 4-189
<i>descriptor_var</i>	Host variable that identifies the system-descriptor area	System-descriptor area must have been allocated	"Quoted String" on page 4-189
<i>parameter_var</i>	Host variable whose contents replace a question (?) mark placeholder in a prepared SQL statement	Must be a character or collection data type	Language specific
<i>sqlda_pointer</i>	Pointer to sqlda structure defining data type and memory location of values to replace question (?) marks in a prepared statement	Cannot begin with a dollar (\$) sign nor with a colon (:). You must use an sqlda structure with dynamic SQL statements.	"DESCRIBE" on page 2-311

Usage

A *cursor* is an identifier associated with an SQL statement that returns an ordered set of values. The OPEN statement activates a cursor that the DECLARE statement defined.

Cursor can be classified by their associated SQL statements:

- A *Select cursor*: a cursor that is associated with a SELECT statement
- A *Function cursor*: a cursor that is associated with the EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement
- An *Insert cursor*: a cursor that is associated with the INSERT statement
- A *Collection cursor*: a Select or Insert cursor that operates on a collection variable.

In UDRs written in the SPL language, the OPEN statement can reference only Select or Function cursors, and these must specify the identifier of the cursor, rather than a variable that stores the *cursor_id*. The OPEN statement cannot reference a direct cursor that the FOREACH statement of SPL has declared.

The specific actions that the database server takes differ, depending on the statement with which the cursor is associated. In ESQL/C, when you associate one of the previous statements with a cursor directly (that is, you do not prepare the statement and associate the statement identifier with the cursor), the OPEN statement implicitly prepares the statement. (This is not a feature of OPEN in SPL routines, where the DECLARE statement associates a cursor with the identifier of an existing prepared statement, rather than directly with SQL statement text.)

In an ANSI-compliant database, you receive an error code if you try to open a cursor that is already open.

Opening a Select Cursor

When you open either a Select cursor or an update cursor that is created with the SELECT... FOR UPDATE syntax, the SELECT statement is passed to the database server with any values that are specified in the USING clause. The database server processes the query to the point of locating or constructing the first row of the active set. The following example illustrates a simple OPEN statement in Informix ESQL/C:

```
EXEC SQL declare s_curs cursor for select * from orders;
EXEC SQL open s_curs;
```

An SPL routine cannot reference an update cursor in the OPEN statement.

Opening an Update Cursor Inside a Transaction

If you are working in a database with explicit transactions, you must open an update cursor within a transaction. This requirement is waived if you declared the cursor using the WITH HOLD option.

Opening a Function Cursor

When you open a Function cursor, the EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement is passed to the database server with any values that are specified in the USING clause.

The values in the USING clause are passed as arguments to the user-defined function. This user-defined function must be declared to accept values. (If the statement was previously prepared, the statement was passed to the database server when it was prepared.) The database server executes the function to the point where it returns the first set of values.

The following example illustrates a simple OPEN statement in Informix ESQL/C:

```
EXEC SQL declare s_curs cursor for
    execute function new_func(arg1,arg2)
    into :ret_val1, :ret_val2;
EXEC SQL open s_curs;
```

In Extended Parallel Server, to re-create this example, use the CREATE PROCEDURE statement instead of the CREATE FUNCTION statement.

Reopening a Select or Function Cursor

The database server evaluates the values that are specified in the USING clause of the OPEN statement only when it opens a Select cursor or Function cursor. While the cursor is open, subsequent changes to program variables in the USING clause do not change the active set of the cursor.

In a database that is ANSI-compliant, you receive an error code if you try to open a cursor that is already open.

In a database that is not ANSI-compliant, a subsequent OPEN statement closes the cursor and then reopens it. When the database server reopens the cursor, it creates a new active set, based on the current values of the variables in the USING clause. If the variables have changed since the previous OPEN statement, reopening the cursor can generate an entirely different active set.

Even if the values of the variables are unchanged, the values in the active set can be different, as in the following situations:

- If the user-defined function takes a different execution path from the previous OPEN statement on a Function cursor
- If data in the table was modified since the previous OPEN statement on a Select cursor

The database server can process most queries dynamically, without pre-fetching all rows when it opens the Select or Function cursor. Therefore, if other users are

modifying the table at the same time that the cursor is being processed, the active set might reflect the results of these actions.

For some queries, the database server evaluates the entire active set when it opens the cursor. These queries include those with the following features:

- Queries that require sorting: those with an ORDER BY clause or with the DISTINCT or UNIQUE keyword
- Queries that require hashing: those with a join or with the GROUP BY clause

For these queries, any changes that other users make to the table while the cursor is being processed are not reflected in the active set.

Errors Associated with Select and Function Cursors

Because the database server is seeing the query for the first time, it might detect errors. In this case, it does not actually return the first row of data, but it resets the **SQLCODE** variable and the **sqlca.sqlcode** field of the **sqlca**. The value is either negative or zero, as the following table describes.

Code Value

Significance

Negative

An error was detected in the SELECT statement

Zero The SELECT statement is valid

Unlike ESQL/C routines, SPL routines do not have direct access to the **sqlca** structure. An ESQL/C routine must invoke the built-in **SQLCODE** function explicitly to access the return code of the SELECT, EXECUTE FUNCTION, or EXECUTE PROCEDURE statement associated with the cursor that OPEN references.

If the SELECT, SELECT...FOR UPDATE, EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement is valid, but no rows match its criteria, the first FETCH statement returns a value of 100 (SQLNOTFOUND), meaning that no rows were found.

Tip: When you encounter an **SQLCODE** error, a corresponding **SQLSTATE** error value also exists. For information about how to view the message text, refer to "Using the SQLSTATE Error Status Code" on page 2-387.

Opening an Insert Cursor (ESQL/C)

When you open an Insert cursor, the cursor passes the INSERT statement to the database server, which checks the validity of the keywords and column names. The database server also allocates memory for an insert buffer to hold new data. (See "DECLARE " on page 2-286.)

An OPEN statement for a cursor that is associated with an INSERT statement cannot include a USING clause.

Example of Opening an Insert Cursor

The following Informix ESQL/C example illustrates an OPEN statement with an Insert cursor:

```
EXEC SQL prepare s1 from
    'insert into manufact values ('npr', 'napier)';
EXEC SQL declare in_curs cursor for s1;
EXEC SQL open in_curs;
EXEC SQL put in_curs;
EXEC SQL close in_curs;
```

Reopening an Insert Cursor

When you reopen an Insert cursor that is already open, you effectively flush the insert buffer; any rows that are stored in the insert buffer are written into the database table. The database server first closes the cursor, which causes the flush and then reopens the cursor. For information about how to check errors and count inserted rows, see “Error Checking” on page 2-484.

In an ANSI-compliant database, you receive an error code if you try to open a cursor that is already open.

Opening a Collection Cursor (ESQL/C)

You can declare both Select and Insert cursors on collection variables. Such cursors are called *Collection cursors*. You must use the OPEN statement to activate these cursors.

Use the name of a collection variable in the USING clause of the OPEN statement. For more information on the use of OPEN ... USING with a collection variable, see “Fetching from a Collection Cursor (IDS)” on page 2-374 and “Inserting into a Collection Cursor (IDS)” on page 2-483.

USING Clause

The USING clause is required when the cursor is associated with a prepared statement that includes question-mark (?) placeholders, as follows:

- A SELECT statement with input parameters in its WHERE clause
- An EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement with input parameters as arguments to its user-defined function
- An INSERT statement with input parameters in its VALUES clause (in ESQL/C).

In SPL routines, you must specify these parameters as SPL variables.

In ESQL/C, you can supply values for these parameters in one of the following ways:

- You can specify one or more host variables.
- You can specify a system-descriptor area.
- You can specify a pointer to an **sqllda** structure.

For more information, see “PREPARE ” on page 2-468.

If you know the number and the order of parameters to be supplied at runtime and their data types, you can define the parameters that are needed by the statement as host variables in your program. You pass parameters to the database server positionally, by opening the cursor with the USING keyword, followed by the names of the variables in their sequential order. These variables are matched with the SELECT or EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement question-mark (?) placeholders in a one-to-one correspondence, from left to right.

You cannot include indicator variables of ESQL/C in the list of variables. To use an indicator variable, you must include the SELECT or EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement text as part of the DECLARE statement, rather than the identifier of a prepared statement.

You must supply one host variable name for each placeholder. The data type of each variable must be compatible with the corresponding type that the prepared statement requires. The following Informix ESQL/C code fragment opens a Select cursor and specifies host variables in the USING clause:

```
sprintf (select_1, "%s %s %s %s %s",
        "SELECT o.order_num, sum(total price)",
        "FROM orders o, items i",
        "WHERE o.order_date > ? AND o.customer_num = ?",
        "AND o.order_num = i.order_num",
        "GROUP BY o.order_num");
EXEC SQL prepare statement_1 from :select_1;
EXEC SQL declare q_curs cursor for statement_1;
EXEC SQL open q_curs using :o_date, :o.customer_num;
```

The following example illustrates the USING clause of the OPEN statement with an EXECUTE FUNCTION statement in an Informix ESQL/C code fragment:

```
stcopy ("EXECUTE FUNCTION one_func(?, ?)", exfunc_stmt);
EXEC SQL prepare exfunc_id from :exfunc_stmt;
EXEC SQL declare func_curs cursor for exfunc_id;
EXEC SQL open func_curs using :arg1, :arg2;
```

In Extended Parallel Server, to re-create this example use the CREATE PROCEDURE statement instead of the CREATE FUNCTION statement.

Specifying a System Descriptor Area (ESQL/C)

If you do not know the number of parameters to be supplied at runtime or their data types, you can associate input values from a system-descriptor area. A system-descriptor area describes the data type and memory location of one or more values to replace question-mark (?) placeholders.

A system-descriptor area conforms to the X/Open standards.

Use the SQL DESCRIPTOR keywords to introduce the name of a system descriptor area as the location of the parameters.

The COUNT field in the system-descriptor area corresponds to the number of dynamic parameters in the prepared statement. The value of COUNT must be less than or equal to the number of item descriptors that were specified when the system-descriptor area was allocated with the ALLOCATE DESCRIPTOR statement. You can obtain the value of a field with the GET DESCRIPTOR statement and set the value with the SET DESCRIPTOR statement.

The following example shows the OPEN ... USING SQL DESCRIPTOR statement:

```
EXEC SQL allocate descriptor 'desc1';
...
EXEC SQL open selcurs using sql descriptor 'desc1';
```

As the example indicates, the system descriptor area must be allocated before you reference it in the OPEN statement.

Specifying a Pointer to an `sqllda` Structure (ESQL/C)

If you do not know the number of parameters to be supplied at runtime, or their data types, you can associate input values from an `sqllda` structure. An `sqllda` structure lists the data type and memory location of one or more values to replace question-mark (?) placeholders.

Use the `DESCRIPTOR` keyword to introduce a pointer to the `sqllda` structure as the location of the parameters.

The `sqllda` value specifies the number of input values that are described in occurrences of `sqlvar`. This number must correspond to the number of dynamic parameters in the prepared statement.

Example of Specifying a Pointer to an `sqllda` Structure

The following example shows an `OPEN ... USING DESCRIPTOR` statement:

```
struct sqllda *sdp;
...
EXEC SQL open selcurs using descriptor sdp;
```

Using the `WITH REOPTIMIZATION` Option (ESQL/C)

Use the `WITH REOPTIMIZATION` keywords to reoptimize your query plan. When you prepare `SELECT`, `EXECUTE FUNCTION`, or `EXECUTE PROCEDURE` statements, the database server uses a query plan to optimize the query. If you later modify the data associated with the prepared statement, you can compromise the effectiveness of the query plan for that statement. In other words, if you change the data, you might deoptimize your query. To ensure optimization of your query, you can prepare the statement again, or open the cursor again using the `WITH REOPTIMIZATION` option.

You should generally use the `WITH REOPTIMIZATION` option, because it provides the following advantages over preparing a statement again:

- Rebuilds only the query plan, rather than the entire statement
- Uses fewer resources
- Reduces overhead
- Requires less time

The `WITH REOPTIMIZATION` option forces the database server to optimize the query-design plan before it processes the `OPEN` cursor statement.

The following examples use the `WITH REOPTIMIZATION` keywords:

```
EXEC SQL open selcurs using descriptor sdp with reoptimization;
```

Relationship Between `OPEN` and `FREE`

The database server allocates resources to prepared statements and open cursors. If you execute a `FREE statement_id` or `FREE statement_id_var` statement, you can still open the cursor associated with the freed statement ID. If you release resources with a `FREE cursor_id` or `FREE cursor_id_var` statement, however, you cannot use the cursor unless you declare the cursor again.

Similarly, if you use the `SET AUTOFREE` statement for one or more cursors, when the program closes the specific cursor, the database server automatically frees the cursor-related resources. In this case, you cannot use the cursor unless you declare the cursor again.

DDL Operations on Tables Referenced by Cursors

Various DDL statements can drop, rename, or alter the schema of a table that is referenced directly (or indirectly, by the identifier of a prepared statement) in the DECLARE statement that defines a cursor. Subsequent OPEN operations on the cursor might fail with error -710, or might produce unexpected results. Changing the number of columns or the data type of a column has this effect, and the user typically must reissue the DESCRIBE statement, the PREPARE statement, and (for cursors associated with routines) the UPDATE STATISTICS statement for any SPL routines that reference a table whose schema has been modified.

These restrictions do not apply, however, if an index is added or dropped when automatic recompilation is enabled for prepared objects and for SPL routines that reference tables that ALTER TABLE, CREATE INDEX, or DROP INDEX operations have modified. This is the default behavior of Dynamic Server. For more information about enabling or disabling automatic recompilation after schema changes, see the description of the IFX_AUTO_REPREPARE option to the SET ENVIRONMENT statement. For more information about the AUTO_REPREPARE configuration parameter, see your *IBM Informix Administrator's Reference*.

When the AUTO_REPREPARE configuration parameter and the IFX_AUTO_REPREPARE session environment variable are set to disable recompilation of prepared objects, however, adding an index to a table that is referenced directly or indirectly in a DECLARE statement can similarly invalidate the associated cursor. Subsequent OPEN statements that specify the invalid cursor fail, even if they include the WITH REOPTIMIZATION keywords. If an index is added to the table that is associated with a cursor while automatic recompilation is disabled, you must prepare the statement again and declare the cursor again before you can open the cursor. For cursors associated with calls to an SPL routines, you must run the UPDATE STATISTICS statement for routines that reference tables to which an index has been added or dropped. You cannot simply reopen a cursor that is based on a prepared statement that is no longer valid.

Related Information

Related statements: "ALLOCATE DESCRIPTOR " on page 2-2, "DEALLOCATE DESCRIPTOR " on page 2-284, "DESCRIBE" on page 2-311, "CLOSE " on page 2-77, "DECLARE " on page 2-286, "EXECUTE " on page 2-348, "FETCH " on page 2-367, "FLUSH " on page 2-377, "FREE " on page 2-379, "GET DESCRIPTOR " on page 2-381, "PREPARE " on page 2-468, "PUT " on page 2-479, " SET AUTOFREE " on page 2-582, " SET DEFERRED_PREPARE" on page 2-603, and "SET DESCRIPTOR " on page 2-605

For a task-oriented discussion of the OPEN statement, see the *IBM Informix Guide to SQL: Tutorial*.

For more information on system-descriptor areas and the **sqllda** structure, refer to the *IBM Informix ESQL/C Programmer's Manual*.

OUTPUT

The OUTPUT statement writes query results in an operating-system file, or pipes query results to another program

This statement is an extension to the ANSI/ISO standard for SQL. You can use this statement only with DB–Access.

Syntax

►► OUTPUT TO *filename* | *program* | WITHOUT HEADINGS | SELECT Statement (1) ►►

Notes:

1 See “SELECT” on page 2-528

Element	Description	Restrictions	Syntax
<i>filename</i>	Path and filename where query results are written. The default path is the current directory.	Can specify a new or existing file. If the file exists, query results overwrite the current contents of the file.	Must conform to the rules of your operating system.
<i>program</i>	Name of a program to receive the query results as input	Program must exist, must be known to the operating system, and must be able to read the results of a query.	Must conform to the rules of your operating system.

Usage

You can use the OUTPUT statement to direct the results of a query to an operating-system file or to a program. You can also specify whether column headings should be omitted from the query output.

Sending Query Results to a File

To send the results of a query to an operating-system file, specify the full pathname for the file. If the file already exists, the output overwrites the current contents.

The following examples show how to send the result of a query to an operating-system file. The example uses UNIX file naming conventions.

```
OUTPUT TO /usr/april/query1
SELECT * FROM cust_calls WHERE call_code = 'L'
```

Displaying Query Results Without Column Headings

To display the results of a query without column headings, use the WITHOUT HEADINGS keywords.

Sending Query Results to Another Program

You can use the keyword PIPE to send the query results to another program, as the following example shows:

```
OUTPUT TO PIPE more
SELECT customer_num, call_dtime, call_code
FROM cust_calls;
```

Related Information

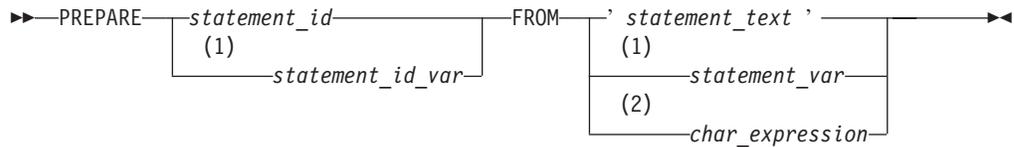
Related statements: “SELECT” on page 2-528 and “UNLOAD ” on page 2-683

PREPARE

Use the PREPARE statement to parse, validate, and generate an execution plan for one or more SQL statements at runtime.

Use this statement in ESQL/C or SPL routines.

Syntax



Notes:

- 1 ESQL/C only
- 2 SPL only

Element	Description	Restrictions	Syntax
<i>char_expression</i>	Expression that evaluates to the text of a single SQL statement	Statement must be a SELECT, EXECUTE FUNCTION, or EXECUTE PROCEDURE	"Expression" on page 4-40
<i>statement_id</i>	Identifier declared here for the prepared object	Must be unique in the routine among names of cursors and prepared objects (and in SPL, among variables)	"Identifier " on page 5-21
<i>statement_id_var</i>	Host variable storing <i>statement_id</i>	Must have been previously declared as a character data type	Language specific
<i>statement_text</i>	Text of the SQL statement(s) to prepare	See "Preparing Multiple SQL Statements" on page 2-476 and "Statement Text" on page 2-470.	"Quoted String" on page 4-189.
<i>statement_var</i>	Host variable storing the text of one or more SQL statements	Must be a character data type. Not valid if the SQL statement(s) contains the Collection-Derived Table segment.	Language specific

Usage

The PREPARE statement enables your program to assemble the text of one (or for ESQL/C, more than one) SQL statement at runtime, to declare an identifier for the resulting *prepared object*, and to make it executable. This dynamic form of SQL is accomplished in three steps:

1. The PREPARE statement accepts statement text as input, either as a quoted string, or an ESQL/C character variable, or (in SPL) as the value to which a character expression evaluates. Statement text can contain question-mark (?) placeholders to represent values that are to be defined when the statement is executed.
2. The OPEN statement (and in ESQL/C routines, the EXECUTE statement) can supply the required input values and execute the prepared statement once or many times.
3. Resources allocated to the prepared statement can be released later using the FREE statement.

For more information about the replacement of placeholders in prepared statements with runtime values, see the section “Preparing Statements That Receive Parameters” on page 2-473.

The same collating order that is current when you create a prepared object is also used when that object is executed, even if the execution-time collation of the session (or of `DB_LOCALE`) is different.

Restrictions

The number of prepared objects in a single program is limited by available memory. These include statement identifiers declared in `PREPARE` statements (*statement_id* or *statement_id_var*) and declared cursors. To avoid exceeding the limit, use the `FREE` statement to release some statements or cursors.

The maximum length of a `PREPARE` statement is 64 kilobytes in `ESQL/C` routines, or 32 kilobytes in `SPL` routines.

In `SPL` routines, a prepared object can include the text of no more than one `SQL` statement, and that statement must be either an `EXECUTE FUNCTION`, `EXECUTE PROCEDURE`, or `SELECT` statement, but the `SELECT` statement cannot include the `INTO variable`, `INTO TEMP`, or `FOR UPDATE` clause.

An expression that specifies the statement text in an `SPL` routine must evaluate to a `CHAR`, `LVARCHAR`, `NCHAR`, `NVARCHAR`, or `VARCHAR` data type. You must explicitly cast to one of these types an expression of any other text data type, such as a `UDT`.

For restrictions in `ESQL/C` routines on the `SQL` statements in the character string, see “Restricted Statements in Single-Statement Prepares” on page 2-471 and “Restricted Statements in Multistatement Prepared Objects” on page 2-477.

Declaring a Statement Identifier

`PREPARE` sends the statement text to the database server, which analyzes the statement text. If the text contains no syntax errors, the database server translates it to an internal form. This translated statement is saved for later execution in a data structure that the `PREPARE` statement allocates. The name of the structure is the value that is assigned to the statement identifier in the `PREPARE` statement. Subsequent `SQL` statements can refer to the structure by using the same statement identifier that was used in the `PREPARE` statement.

A subsequent `FREE` statement releases the database server resources that were allocated to the statement. After you release these resources with `FREE`, you cannot use the statement identifier in a `DECLARE` statement or (in `ESQL/C`) with the `EXECUTE` statement until you prepare the statement again.

The database server resources for the prepared objects that an `SPL` routine defines are released automatically when the routine exits.

Scope of Statement Identifiers

An `ESQL/C` program can consist of one or more source-code files. By default, the scope of reference of a statement identifier is global to the program. Therefore, a statement identifier that is prepared in one file can be referenced from another file.

In a multiple-file program, if you want to limit the scope of reference of a statement identifier to the file in which it is prepared, preprocess all the files with the `-local` command-line option.

Releasing a Statement Identifier

A statement identifier can represent only one SQL statement or (in ESQL/C) one semicolon-separated list of SQL statements at a time. A new PREPARE statement can specify an existing statement identifier if you want to bind the identifier to a different SQL statement text.

The PREPARE statement supports dynamic statement-identifier names, which allow you to prepare a statement identifier as an identifier or (in ESQL/C) as a host variable of a data type that can contain a character string. The first example that follows shows a statement identifier that was specified as a host variable. The second specifies a statement identifier as a character string.

```
stcopy ("query2", stmtid);  
EXEC SQL prepare :stmtid from 'select * from customer';  
  
EXEC SQL prepare query2 from 'select * from customer';
```

The variable must be a character data type. In C, it must be declared as **char**.

In an SPL routine, statement identifiers that the PREPARE statement declares are automatically defined in the local scope. Do not attempt to declare a statement identifier as having local or global scope. A statement identifier defined in one SPL routine is not visible to any other SPL routine that the same session calls. SPL statement identifiers share the same namespace as SPL variables and cursor names.

Statement Text

Statement text can be specified in the PREPARE statement

- as a quoted string
- or as text that is stored in an ESQL/C program variable
- or (in SPL routines) as a character expression.

as a quoted string or as text that is stored in a program variable. The following restrictions apply to the statement text:

- The text can contain only SQL statements. It cannot contain statements or comments from the host programming language.
- The text can contain comments preceded by a double hyphen (--), or that are enclosed in braces ({ }) or in C-style slash and asterisk (/* */) delimiters.

These symbols introduce or enclose SQL comments. For more information on SQL comment symbols, see “How to Enter SQL Comments” on page 1-3.

- The text can contain either a single SQL statement or (in ESQL/C routines) a series of statements that are separated by semicolon (;) symbols.

For a list of SQL statements that cannot be prepared, see “Restricted Statements in Single-Statement Prepares” on page 2-471. For more information on how to prepare multiple SQL statements, see “Preparing Multiple SQL Statements” on page 2-476.

- The text cannot include an embedded SQL statement prefix or terminator, such as a dollar sign (\$) or the words EXEC SQL.
- Host-language variables are not recognized as such in prepared text.

Therefore, you cannot prepare a SELECT (or EXECUTE FUNCTION or EXECUTE PROCEDURE) statement that includes an INTO clause, because the INTO clause requires a host-language variable.

- The only identifiers that you can use are names that are defined in the database, such as names of tables and columns. For more information on how to use identifiers in statement text, see “Preparing Statements with SQL Identifiers” on page 2-474.
- Use a question mark (?) as a placeholder to indicate where data is supplied when the statement executes, as in this Informix ESQL/C example:

```
EXEC SQL prepare new_cust from
      'insert into customer(fname,lname) values(?,?)';
```

For more information on how to use question marks as placeholders, see “Preparing Statements That Receive Parameters” on page 2-473.

If the prepared statement contains the Collection-Derived Table segment or an Informix ESQL/C collection variable, some additional limitations exist on how you can assemble the text for the PREPARE statement. For information about dynamic SQL, see the *IBM Informix ESQL/C Programmer's Manual*. SPL routines cannot use dynamic SQL statements to process prepared statements that contain the Collection-Derived Table segment.

Preparing and Executing User-Defined Routines

The way to prepare a user-defined routine (UDR) depends on whether the UDR is a user-defined procedure or a user-defined function:

- To prepare a user-defined procedure, prepare the EXECUTE PROCEDURE statement that executes the procedure.

To execute the prepared procedure, use the EXECUTE statement.

- To prepare a user-defined function, prepare the EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement that executes the function.

You cannot include the INTO clause of EXECUTE FUNCTION (or EXECUTE PROCEDURE) in the PREPARE statement.

How to execute a prepared user-defined function depends on whether it returns only one group or multiple groups of values. Use the EXECUTE statement for user-defined functions that return only one group of values.

To execute user-defined functions that return more than one group of return values, you must associate the EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement with a cursor.

Restricted Statements in Single-Statement Prepares

In general, you can prepare any data manipulation language (DML) statement.

In Dynamic Server, you can prepare any single SQL statement except the following statements:

- ALLOCATECOLLECTION
- ALLOCATEDESCRIPTOR
- ALLOCATEROW
- CLOSE
- CONNECT
- CREATEFUNCTIONFROM

- CREATEPROCEDUREFROM
- CREATEROUTINEFROM
- DEALLOCATECOLLECTION
- DEALLOCATEDESCRIPTOR
- DEALLOCATEROW
- DECLARE
- DESCRIBE
- DISCONNECT
- EXECUTE
- EXECUTEIMMEDIATE
- FETCH
- FLUSH
- FREE
- GETDESCRIPTOR
- GETDIAGNOSTICS
- INFO
- LOAD
- OPEN
- OUTPUT
- PREPARE
- PUT
- SETAUTOFREE
- SETCONNECTION
- SETDEFERRED_PREPARE
- SETDESCRIPTOR
- UNLOAD
- WHENEVER

In Extended Parallel Server, you can prepare any single SQL statement except the following statements:

- ALLOCATE DESCRIPTOR
- CLOSE CONNECT
- CREATE PROCEDURE FROM
- DEALLOCATE DESCRIPTOR
- DECLARE
- DESCRIBE
- DISCONNECT
- EXECUTE
- EXECUTE IMMEDIATE
- FETCH
- FLUSH
- FREE
- GET DESCRIPTOR
- GET DIAGNOSTICS
- INFO

- LOAD
- OPEN
- OUTPUT
- PREPARE
- PUT
- SET CONNECTION
- SET DEFERRED_PREPARE
- SET DESCRIPTOR
- UNLOAD
- WHENEVER

You can prepare a SELECT statement. If SELECT includes the INTO TEMP clause, an ESQL/C program can execute the prepared statement with an EXECUTE statement. If it does not include the INTO TEMP clause, the statement returns rows of data. Use DECLARE, OPEN, and FETCH cursor statements to retrieve the rows.

In ESQL/C, a prepared SELECT statement can include a FOR UPDATE clause. This clause is used with the DECLARE statement to create an update cursor. The next example shows a SELECT statement with a FOR UPDATE clause in Informix ESQL/C:

```
sprintf(up_query, "%s %s %s",
        "select * from customer ",
        "where customer_num between ? and ? ",
        "for update");
EXEC SQL prepare up_sel from :up_query;
EXEC SQL declare up_curs cursor for up_sel;
EXEC SQL open up_curs using :low_cust,:high_cust;
```

Preparing Statements When Parameters Are Known

In some prepared statements, all necessary information is known at the time the statement is prepared. The following example in Informix ESQL/C shows two statements that were prepared from constant data:

```
sprintf(redo_st, "%s %s",
        "drop table workt1; ",
        "create table workt1 (wtk serial, wtv float)" );
EXEC SQL prepare redotab from :redo_st;
```

Preparing Statements That Receive Parameters

In some statements, parameters are unknown when the statement is prepared because a different value can be inserted each time the statement is executed. In these statements, you can use a question-mark (?) placeholder where a parameter must be supplied when the statement is executed.

The PREPARE statements in the following Informix ESQL/C examples show some uses of question-mark (?) placeholders:

```
EXEC SQL prepare s3 from
        'select * from customer where state matches ?';
EXEC SQL prepare in1 from 'insert into manufact values (?,?,:)';
sprintf(up_query, "%s %s",
        "update customer set zipcode = ?"
```

```

    "where current of zip_cursor");
EXEC SQL prepare update2 from :up_query;
EXEC SQL prepare exfunc from
    'execute function func1 (?, ?)';

```

You can use a placeholder to defer evaluation of a value until runtime only for an expression, but not for an SQL identifier, except as noted in “Preparing Statements with SQL Identifiers.”

The following example of an Informix ESQL/C code fragment prepares a statement from a variable that is named **demoquery**. The text in the variable includes one question-mark (?) placeholder. The prepared statement is associated with a cursor and, when the cursor is opened, the USING clause of the OPEN statement supplies a value for the placeholder:

```

EXEC SQL BEGIN DECLARE SECTION;
    char queryvalue [6];
    char demoquery [80];
EXEC SQL END DECLARE SECTION;

EXEC SQL connect to 'stores_demo';
sprintf(demoquery, "%s %s",
    "select fname, lname from customer ",
    "where lname > ? ");
EXEC SQL prepare quid from :demoquery;
EXEC SQL declare democursor cursor for quid;
strcpy("C", queryvalue);
EXEC SQL open democursor using :queryvalue;

```

The USING clause is available in both OPEN statements that are associated with a cursor and EXECUTE statements (all other prepared statements).

You can use a question-mark (?) placeholder to represent the name of an Informix ESQL/C or SPL collection variable.

Preparing Statements with SQL Identifiers

In general, you must specify SQL identifiers explicitly in the statement text when you prepare the statement. In a few special cases, however, you can use the question-mark (?) placeholder for an SQL identifier:

- For the database name in the DATABASE statement.
- For the dbspace name in the IN *dbspace* clause of the CREATE DATABASE statement.
- For the cursor name in statements that use cursor names.

Obtaining SQL Identifiers from User Input

If a prepared statement requires identifiers, but the identifiers are unknown when you write the prepared statement, you can construct a statement that receives SQL identifiers from user input.

The following Informix ESQL/C example prompts the user for the name of a table and uses that name in a SELECT statement. Because this name is unknown until runtime, the number and data types of the table columns are also unknown. Therefore, the program cannot allocate host variables to receive data from each row in advance. Instead, this program fragment describes the statement into an **sqlda** descriptor and fetches each row with the descriptor. The fetch puts each row into memory locations that the program provides dynamically.

If a program retrieves all the rows in the active set, the FETCH statement would be placed in a loop that fetched each row. If the FETCH statement retrieves more than one data value (column), another loop exists after the FETCH, which performs some action on each data value:

```
#include <stdio.h>
EXEC SQL include sqllda;
EXEC SQL include sqltypes;
char *malloc ( );

main()
{
    struct sqllda *demodesc;
    char tablename[19];
    int i;
EXEC SQL BEGIN DECLARE SECTION;
    char demoselect[200];
EXEC SQL END DECLARE SECTION;

/* This program selects all the columns of a given tablename.
   The tablename is supplied interactively. */

EXEC SQL connect to 'stores_demo';
printf( "This program does a select * on a table\n" );
printf( "Enter table name: " );
scanf( "%s", tablename );
sprintf(demoselect, "select * from %s", tablename );

EXEC SQL prepare iid from :demoselect;
EXEC SQL describe iid into demodesc;

/* Print what describe returns */

for ( i = 0; i < demodesc->sqld; i++ )
    prsqllda (demodesc->sqlvar + i);
/* Assign the data pointers. */

for ( i = 0; i < demodesc->sqld; i++ )
{
    switch (demodesc->sqlvar[i].sqltype & SQLTYPE)
    {
        case SQLCHAR:
            demodesc->sqlvar[i].sqltype = CCHARTYPE;
            /* make room for null terminator */
            demodesc->sqlvar[i].sqlllen++;
            demodesc->sqlvar[i].sqldata =
                malloc( demodesc->sqlvar[i].sqlllen );
            break;

        case SQLSMINT: /* fall through */
        case SQLINT: /* fall through */
        case SQLSERIAL:
            demodesc->sqlvar[i].sqltype = CINTTYPE;
            demodesc->sqlvar[i].sqldata =
                malloc( sizeof( int ) );
            break;
        /* And so on for each type. */
    }
}

/* Declare and open cursor for select . */
EXEC SQL declare d_curs cursor for iid;
EXEC SQL open d_curs;

/* Fetch selected rows one at a time into demodesc. */
for( ; ; )
{
```

```

printf( "\n" );
EXEC SQL fetch d_curs using descriptor demodesc;
if ( sqlca.sqlcode != 0 )
    break;
for ( i = 0; i < demodesc->sqld; i++ )
{
    switch ( demodesc->sqlvar[i].sqltype)
    {
        case CCHARTYPE:
            printf( "%s: \"%s\n", demodesc->sqlvar[i].sqlname,
                demodesc->sqlvar[i].sqldata );
            break;
        case CINTTYPE:
            printf( "%s: %d\n", demodesc->sqlvar[i].sqlname,
                *((int *) demodesc->sqlvar[i].sqldata) );
            break;
        /* And so forth for each type... */
    }
}
EXEC SQL close d_curs;
EXEC SQL free d_curs;
/* Free the data memory. */

for ( i = 0; i < demodesc->sqld; i++ )
    free( demodesc->sqlvar[i].sqldata );
free( demodesc );

printf ( "Program Over.\n" );
}

prsqlda(sp)
    struct sqlvar_struct *sp;

{
    printf ( "type = %d\n", sp->sqltype);
    printf ( "len = %d\n", sp->sqllen);
    printf ( "data = %lx\n", sp->sqldata);
    printf ( "ind = %lx\n", sp->sqlind);
    printf ( "name = %s\n", sp->sqlname);
}

```

Preparing Multiple SQL Statements

In ESQL/C, you can execute several SQL statements as one action if you include them in the same PREPARE statement. Multistatement text is processed as a unit; actions are not treated sequentially. Therefore, multistatement text cannot include statements that depend on actions that occur in a previous statement in the text. For example, you cannot create a table and insert values into that table in the same prepared statement block.

If a statement in a multistatement prepare returns an error, the whole prepared statement stops executing. The database server does not execute any remaining statements. In most situations, compiled products return error-status information on the error, but do not indicate which statement in the text causes an error. You can use the `sqlca.sqlerrd[4]` field in the `sqlca` to find the offset of the errors.

In a multistatement prepare, if no rows are returned from a WHERE clause in the following statements, the database server returns SQLNOTFOUND (100):

- UPDATE ... WHERE ...
- SELECT INTO TEMP ... WHERE ...
- INSERT INTO ... WHERE ...

- DELETE FROM ... WHERE ...

In the next example, four SQL statements are prepared into a single Informix ESQL/C string called **query**. Individual statements are delimited with semicolons.

A single PREPARE statement can prepare the four statements for execution, and a single EXECUTE statement can execute the statements that are associated with the **qid** statement identifier:

```

sprintf (query, "%s %s %s %s %s %s %s",
    "update account set balance = balance + ? ",
    "where acct_number = ?;",
    "update teller set balance = balance + ? ",
    "where teller_number = ?;",
    "update branch set balance = balance + ? ",
    "where branch_number = ?;",
    "insert into history values (?, ?);");
EXEC SQL prepare qid from :query;

EXEC SQL begin work;
EXEC SQL execute qid using
    :delta, :acct_number, :delta, :teller_number,
    :delta, :branch_number, :timestamp, :values;
EXEC SQL commit work;

```

Here the semicolons (;) are required as SQL statement-terminator symbols between each SQL statement in the text that **query** holds.

Restricted Statements in Multistatement Prepared Objects

In addition to the statements listed as exceptions in “Restricted Statements in Single-Statement Prepares” on page 2-471, you cannot use the following statements in the text of a multiple-statement prepared object:

- CLOSE DATABASE
- CREATE DATABASE
- DATABASE
- DROP DATABASE
- RENAME DATABASE
- SELECT (*with one exception*)

The following types of statements are also not valid in a multistatement prepare:

- Statements that can cause the current database to close during the execution of the multistatement sequence
- Statements that include references to TEXT or BYTE host variables

In general, you cannot use the SELECT statement in a multistatement prepare. The only form of the SELECT statement allowed in a multistatement prepare is a SELECT statement with an INTO temporary table clause.

Using Prepared Statements for Efficiency

To increase performance efficiency, you can use the PREPARE statement and an EXECUTE statement in a loop to eliminate overhead that redundant parsing and optimizing cause. For example, an UPDATE statement that is located within a WHILE loop is parsed each time the loop runs. If you prepare the UPDATE statement outside the loop, the statement is parsed only once, eliminating overhead and speeding statement execution. The following example shows how to prepare an Informix ESQL/C statement to improve performance:

```

EXEC SQL BEGIN DECLARE SECTION;
    char disc_up[80];
    int cust_num;
EXEC SQL END DECLARE SECTION;
main()
{
    sprintf(disc_up, "%s %s", "update customer ",
        "set discount = 0.1 where customer_num = ?");
    EXEC SQL prepare up1 from :disc_up;
    while (1)
    {
        printf("Enter customer number (or 0 to quit): ");
        scanf("%d", cust_num);
        if (cust_num == 0)
            break;
        EXEC SQL execute up1 using :cust_num;
    }
}

```

Like the SQL statement cache, prepared statements can reduce how often the same query plan is reoptimized, thereby conserving resources in some contexts. The section “Prepared Statements and the Statement Cache” on page 2-655 discusses the use of prepared DML statements, cursors, and the SQL statement cache as combined or alternative techniques for improving query performance.

DDL Operations on Tables Referenced in Prepared Objects

These restrictions do not apply, however, if an index is added or dropped when automatic recompilation is enabled for prepared objects and routines that directly reference tables that ALTER TABLE, CREATE INDEX, or DROP INDEX operations have modified. This is the default behavior of Dynamic Server. For more information about enabling or disabling automatic recompilation after changes to the schema of a table, see “IFX_AUTO_REPREPARE Environment Option (IDS)” on page 2-620.

When the AUTO_REPREPARE configuration parameter and the IFX_AUTO_REPREPARE session environment variable are set to disable automatic recompilation, however, adding an index to a table that a prepared statement references indirectly can similarly invalidate the prepared statement. A subsequent OPEN statement fails if the cursor refers to the invalid prepared statement, even if the OPEN statement includes the WITH REOPTIMIZATION keywords. If an index on an indirectly referenced table is added after the statement was prepared while automatic recompilation is disabled, you must prepare the statement again and declare the cursor again. You cannot simply reopen the cursor if it is based on a prepared statement that is no longer valid.

Related Statements

Related statements: “CLOSE ” on page 2-77, “DECLARE ” on page 2-286, “DESCRIBE” on page 2-311, “EXECUTE ” on page 2-348, “FREE ” on page 2-379, “OPEN ” on page 2-459, “ SET AUTOFREE ” on page 2-582, and “ SET DEFERRED_PREPARE” on page 2-603

For information about basic concepts that relate to the PREPARE statement, see the *IBM Informix Guide to SQL: Tutorial*.

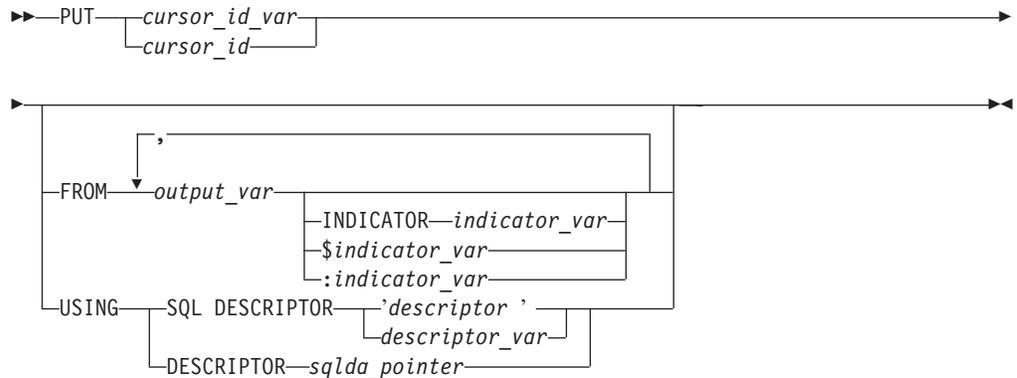
For information about more advanced concepts that relate to the PREPARE statement, see the *IBM Informix ESQL/C Programmer’s Manual*.

PUT

Use the PUT statement to store a row in an insert buffer for later insertion into the database.

This statement is an extension to the ANSI/ISO standard for SQL. You can use this statement with ESQL/C.

Syntax



Element	Description	Restrictions	Syntax
<i>cursor_id</i>	Name of a cursor	Must be open	"Identifier" on page 5-21
<i>cursor_id_var</i>	Host variable = <i>cursor_id</i>	Must be a character type; cursor must be open	Language specific
<i>descriptor</i>	Name of a system-descriptor area	Must already be allocated	"Quoted String" on page 4-189
<i>descriptor_var</i>	Host-variable that contains <i>descriptor</i>	Must already be allocated	"Quoted String" on page 4-189
<i>indicator_var</i>	Host variable to receive a return code if corresponding <i>output_var</i> receives a NULL value	Cannot be a DATETIME or INTERVAL data type	Language specific
<i>output_var</i>	Host variable whose contents replace a question-mark (?) placeholder in a prepared INSERT statement	Must be a character data type	Language specific
<i>sqlda_pointer</i>	Pointer to an sqlda structure	First character cannot be the (\$) or (:) symbol	"DESCRIBE" on page 2-311

Usage

PUT stores a row in an *insert buffer* that is created when the cursor is opened.

If the buffer has no room for the new row when the statement executes, the buffered rows are written to the database in a block, and the buffer is emptied. As a result, some PUT statement executions cause rows to be written to the database, and some do not. You can use the FLUSH statement to write buffered rows to the database without adding a new row. The CLOSE statement writes any remaining rows before it closes an Insert cursor.

If the current database uses explicit transactions, you must execute a PUT statement within a transaction.

The following example uses a PUT statement in Informix ESQL/C:

```
EXEC SQL prepare ins_mcode from
    'insert into manufact values(?,?)';
EXEC SQL declare mcode cursor for ins_mcode;
EXEC SQL open mcode;
EXEC SQL put mcode from :the_code, :the_name;
```

The PUT statement is not an X/Open SQL statement. Therefore, you get a warning message if you compile a PUT statement in X/Open mode.

Supplying Inserted Values

The values in the inserted row can come from one of the following sources:

- Constant values that are written into the INSERT statement
- Program variables that are named in the INSERT statement
- Program variables in the FROM clause of the PUT statement
- Values that are prepared in memory addressed by an **sqllda** structure or a system-descriptor area and then specified in the USING clause of the PUT statement

The system descriptor area or **sqllda** structure that *descriptor* or *sqllda_pointer* references must define a data type and memory location of each value that corresponds to a question-mark (?) placeholder in a prepared INSERT statement.

Using Constant Values in INSERT

The VALUES clause lists the values for the inserted columns. One or more of these values can be constants (that is, numbers or character strings).

When *all* the inserted values are constants, the PUT statement has a special effect. Instead of creating a row and putting it in the buffer, the PUT statement merely increments a counter. When you use a FLUSH or CLOSE statement to empty the buffer, one row and a repetition count are sent to the database server, which inserts that number of rows. In the following Informix ESQL/C example, 99 empty customer records are inserted into the **customer** table. Because all values are constants, no disk output occurs until the cursor closes. (The constant zero for **customer_num** causes generation of a SERIAL value.) The following example inserts 99 empty customer records into the customer table:

```
int count;
EXEC SQL declare fill_c cursor for
    insert into customer(customer_num) values(0);
EXEC SQL open fill_c;
for (count = 1; count <= 99; ++count)
    EXEC SQL put fill_c;
EXEC SQL close fill_c;
```

Naming Program Variables in INSERT

When you associate the INSERT statement with a cursor (in the DECLARE statement), you create an Insert cursor. In the INSERT statement, you can name program variables in the VALUES clause. When each PUT statement is executed, the contents of the program variables at that time are used to populate the row that is inserted into the buffer.

If you are creating an Insert cursor (using DECLARE with INSERT), you must use only program variables in the VALUES clause. Variable names are not recognized in the context of a prepared statement; you associate a prepared statement with a cursor through its statement identifier.

The following Informix ESQL/C example illustrates the use of an Insert cursor. The code includes the following statements:

- The DECLARE statement associates a cursor called **ins_curs** with an INSERT statement that inserts data into the **customer** table.
The VALUES clause specifies a data structure that is called **cust_rec**; the Informix ESQL/C preprocessor converts **cust_rec** to a list of values, one for each component of the structure.
- The OPEN statement creates a buffer.
- A user-defined function (not defined within this example) obtains customer information from user input and stores it in **cust_rec**.
- The PUT statement composes a row from the current contents of the **cust_rec** structure and sends it to the row buffer.
- The CLOSE statement inserts into the **customer** table any rows that remain in the row buffer and closes the Insert cursor:

```
int keep_going = 1;
EXEC SQL BEGIN DECLARE SECTION
    struct cust_row { /* fields of a row of customer table */ } cust_rec;
EXEC SQL END DECLARE SECTION
EXEC SQL declare ins_curs cursor for
    insert into customer values (:cust_row);
EXEC SQL open ins_curs;
while ( (sqlca.sqlcode == 0) && (keep_going) )

    {
    keep_going = get_user_input(cust_rec); /* ask user for new customer */
    if (keep_going ) /* user did supply customer info
    */
        {
        cust_rec.customer_num = 0; /* request new serial value */
        EXEC SQL put ins_curs;
        }
    if (sqlca.sqlcode == 0) /* no error from PUT */
        keep_going = (prompt_for_y_or_n("another new customer") == 'Y')
    }
EXEC SQL close ins_curs;
```

Use an indicator variable if the data to be inserted might be NULL.

Naming Program Variables in PUT

When the INSERT statement is prepared (see "PREPARE " on page 2-468), you cannot use program variables in its VALUES clause, but you can represent values by a question-mark (?) placeholder. List the program variables in the FROM clause of the PUT statement to supply the missing values.

The following Informix ESQL/C example lists host variables in a PUT statement:

```
char answer [1] = 'y';
EXEC SQL BEGIN DECLARE SECTION;
    char ins_comp[80];
    char u_company[20];
EXEC SQL END DECLARE SECTION;

main()
{
    EXEC SQL connect to 'stores_demo';
```

```

EXEC SQL prepare ins_comp from
    'insert into customer (customer_num, company) values (0, ?)';
EXEC SQL declare ins_curs cursor for ins_comp;
EXEC SQL open ins_curs;

while (1)
{
    printf("\nEnter a customer: ");
    gets(u_company);
    EXEC SQL put ins_curs from :u_company;
    printf("Enter another customer (y/n) ? ");
    if (answer = getch() != 'y')
        break;
}
EXEC SQL close ins_curs;
EXEC SQL disconnect all;
}

```

Indicator variables are optional, but you should use an indicator variable if the possibility exists that *output_var* might contain a NULL value. If you specify the indicator variable without the INDICATOR keyword, you cannot put a blank space between *output_var* and *indicator_var*.

Using the USING Clause

If you do not know the number of parameters to be supplied at runtime or their data types, you can associate input values from a system-descriptor area or an **sqlda** structure. Both of these descriptor structures describe the data type and memory location of one or more values to replace question-mark (?) placeholders.

Each time the PUT statement executes, the values that the descriptor structure describes are used to replace question-mark (?) placeholders in the INSERT statement. This process is similar to using a FROM clause with a list of variables, except that your program has full control over the memory location of the data values.

Specifying a System-Descriptor Area

The SQL DESCRIPTOR option specifies the name of a system-descriptor area.

The **COUNT** field in the system-descriptor area corresponds to the number of dynamic parameters in the prepared statement. The value of **COUNT** must be less than or equal to the number of item descriptors that were specified when the system-descriptor area was allocated with the ALLOCATE DESCRIPTOR statement. You can obtain the value of a field with the GET DESCRIPTOR statement and set the value with the SET DESCRIPTOR statement.

A system-descriptor area conforms to the X/Open standards.

The following Informix ESQL/C example shows how to associate values from a system-descriptor area:

```

EXEC SQL allocate descriptor 'desc1';
...
EXEC SQL put selcurs using sql descriptor 'desc1';

```

Specifying an sqlda Structure

Use the **DESCRIPTOR** option to introduce the name of a pointer to an **sqlda** structure. The following Informix ESQL/C example shows how to associate values from an **sqlda** structure:

```

EXEC SQL put selcurs using descriptor pointer2;

```

Inserting into a Collection Cursor (IDS)

A Collection cursor allows you to access the individual elements of a collection variable. To declare a Collection cursor, use the DECLARE statement and include the Collection-Derived Table segment in the INSERT statement that you associate with the cursor. Once you open the Collection cursor with the OPEN statement, the cursor can put elements in the collection variable.

To put elements, one at a time, into the Insert cursor, use the PUT statement and the FROM clause. The PUT statement identifies the Collection cursor that is associated with the collection variable. The FROM clause identifies the element value to be inserted into the cursor. The data type of any host variable in the FROM clause must match the element type of the collection.

Important: The collection variable stores the elements of the collection. However, it has no intrinsic connection with a database column. Once the collection variable contains the correct elements, you must then save the variable into the actual collection column with the INSERT or UPDATE statement.

Suppose you have a table called **children** with the following schema:

```
CREATE TABLE children
(
  age      SMALLINT,
  name     VARCHAR(30),
  fav_colors SET(VARCHAR(20) NOT NULL)
);
```

The following Informix ESQL/C program fragment shows how to use an Insert cursor to put elements into a collection variable called **child_colors**:

```
EXEC SQL BEGIN DECLARE SECTION;
  client collection child_colors;
  char *favorites[]
  (
    "blue",
    "purple",
    "green",
    "white",
    "gold",
    0
  );
  int a = 0;
  char child_name[21];
EXEC SQL END DECLARE SECTION;

EXEC SQL allocate collection :child_colors;

/* Get structure of fav_colors column for untyped
 * child_colors collection variable */
EXEC SQL select fav_colors into :child_colors
  from children where name = :child_name;
/* Declare insert cursor for child_colors collection
 * variable and open this cursor */
EXEC SQL declare colors_curs cursor for
  insert into table(:child_colors)
  values (?);
EXEC SQL open colors_curs;
/* Use PUT to gather the favorite-color values
 * into a cursor */
while (fav_colors[a])
{
  EXEC SQL put colors_curs from :favorites[:a];
  a++
```

```

    ...
}
/* Flush cursor contents to collection variable */
EXEC SQL flush colors_curs;
EXEC SQL update children set fav_colors = :child_colors;

EXEC SQL close colors_curs;
EXEC SQL deallocate collection :child_colors;

```

After the FLUSH statement executes, the collection variable, **child_colors**, contains the elements {"blue", "purple", "green", "white", "gold"}. The UPDATE statement at the end of this program fragment saves the new collection into the **fav_colors** column of the database. Without this UPDATE statement, the new collection would not be added to the collection column.

Writing Buffered Rows

To open an Insert cursor, the OPEN statement creates an insert buffer. The PUT statement puts a row into this insert buffer. The buffered rows are inserted into the database table as a block only when necessary; this process is called *flushing the buffer*. The buffer is flushed after any of the following events:

- Buffer is too full to hold the new row at the start of a PUT statement.
- A FLUSH statement executes.
- A CLOSE statement closes the cursor.
- An OPEN statement specifies an already open cursor, closing it before reopening it. (This implicit CLOSE statement flushes the buffer.)
- A COMMIT WORK statement executes.
- Buffer contains BYTE or TEXT data (flushed after a single PUT statement).

If the program terminates without closing an Insert cursor, the buffer remains unflushed. Rows that were inserted into the buffer since the last flush are lost. Do not rely on the end of the program to close the cursor and flush the buffer.

Error Checking

The **sqlca** structure contains information on the success of each PUT statement as well as information that lets you count the rows that were inserted. The result of each PUT statement is contained in the following fields of the **sqlca**: **sqlca.sqlcode**, **SQLCODE**, and **sqlca.sqlerrd[2]**.

Data buffering with an Insert cursor means that errors are not discovered until the buffer is flushed. For example, an input value that is incompatible with the data type of the column for which it is intended is discovered only when the buffer is flushed. When an error is discovered, buffered rows that were not inserted before the error are *not* inserted; they are lost from memory.

The **SQLCODE** field is set to 0 if no error occurs; otherwise, it is set to an error code. The third element of the **sqlerrd** array is set to the number of rows that were successfully inserted into the database:

- If any row is put into the insert buffer, but *not* written to the database, **SQLCODE** and **sqlerrd** are set to 0 (**SQLCODE** because no error occurred, and **sqlerrd** because no rows were inserted).
- If a block of buffered rows is written to the database during the execution of a PUT statement, **SQLCODE** is set to 0 and **sqlerrd** is set to the number of rows that was successfully inserted into the database.

- If an error occurs while the buffered rows are written to the database, **SQLCODE** indicates the error, and **sqlerrd** contains the number of successfully inserted rows. (The uninserted rows are discarded from the buffer.)

Tip: When you encounter an **SQLCODE** error, a **SQLSTATE** error value also exists. See the GET DIAGNOSTICS statement for details of how to obtain the message text.

To count the number of pending and inserted rows in the database

1. Prepare two integer variables (for example, **total** and **pending**).
2. When the cursor is opened, set both variables to 0.
3. Each time a PUT statement executes, increment both **total** and **pending**.
4. Whenever a PUT or FLUSH statement executes or the cursor closes, subtract the third field of the **SQLERRD** array from **pending**.

At any time, (**total** - **pending**) represents the number of rows actually inserted. If no statements fail, **pending** contains zero after the cursor is closed. If an error occurs during a PUT, FLUSH, or CLOSE statement, the value that remains in **pending** is the number of uninserted (discarded) rows.

Related Statements

Related statements: "ALLOCATE DESCRIPTOR " on page 2-2, "CLOSE " on page 2-77, "DEALLOCATE DESCRIPTOR " on page 2-284, "FLUSH " on page 2-377, "DECLARE " on page 2-286, "GET DESCRIPTOR " on page 2-381, "OPEN " on page 2-459, "PREPARE " on page 2-468, and "SET DESCRIPTOR " on page 2-605

For a task-oriented discussion of the PUT statement, see the *IBM Informix Guide to SQL: Tutorial*.

For more information about error checking, the system-descriptor area, and the **sqllda** structure, see the *IBM Informix ESQL/C Programmer's Manual*.

RELEASE SAVEPOINT

Use the RELEASE SAVEPOINT statement to destroy a specified savepoint. After this statement executes successfully, rollback to the specified savepoint (or to any other savepoint between the RELEASE SAVEPOINT statement and the specified savepoint) is no longer possible.

This statement is compliant with the ANSI/ISO standard for SQL.

Syntax

►►—RELEASE SAVEPOINT—*savepoint*—◄◄

Element	Description	Restrictions	Syntax
<i>savepoint</i>	Name of the savepoint to be destroyed	Must exist in the current savepoint level	"Identifier " on page 5-21

Usage

The RELEASE SAVEPOINT statement destroys the specified savepoint. Any savepoints set between that savepoint and the RELEASE SAVEPOINT statement in the current savepoint level are also destroyed. Any other savepoints, however, that were set earlier than the specified savepoint in the current savepoint level continue to be active.

The RELEASE SAVEPOINT statement fails with an error in the following contexts:

- No SQL transaction is open.
- No savepoint with the specified name exists in the current savepoint level.
- The statement is part of a triggered action.
- The statement is part of an XA transaction.
- The autocommit transaction mode of the client API is enabled.
- The statement is part of a cross-server distributed SQL transaction in which one of the participating database servers does not support savepoints.
- The statement is issued within a UDR that is invoked within a DML statement.

The identifier of any savepoint that RELEASE SAVEPOINT destroys can be reused in a subsequent SAVEPOINT statement of the same savepoint level, even if the released savepoint was set by a SAVEPOINT statement that included the UNIQUE keyword.

Because savepoints are program objects, not database objects, the RELEASE SAVEPOINT statement has no direct effect on the database or on its system catalog tables. RELEASE SAVEPOINT can affect user tables and the system catalog indirectly, however, if it changes the scope of a subsequent ROLLBACK TO SAVEPOINT operation that cancels uncommitted changes to the database within a different portion of the current savepoint level, as the next example illustrates.

The following program fragment sets and then releases a savepoint called **sp45**:

```
BEGIN WORK;
CREATE DATABASE third_base IN db3 WITH BUFFERED LOG;
SAVEPOINT sp46;
CREATE TABLE tab1 ( col1 INT, col2 CHAR(24));
SAVEPOINT sp45 UNIQUE;
...
CREATE TABLE tab2 ( col1 INT8, col2 LVARCHAR(24000));
SAVEPOINT sp44;
...
RELEASE SAVEPOINT sp45;
ROLLBACK TO SAVEPOINT;
```

The effect of the RELEASE SAVEPOINT statement in this example is to destroy two savepoint objects, **sp45** and **sp44**. If the only remaining savepoint in the current savepoint level is sp46, the subsequent ROLLBACK TO SAVEPOINT statement cancels the DDL statements that created **tab1** and **tab2**, and any DML operations on those tables that preceded the ROLLBACK statement. The rollback does not, however, cancel the CREATE DATABASE statement that created the **third_base** database. Without the RELEASE SAVEPOINT statement, the CREATE TABLE statement that created **tab1** would not have been cancelled, because Dynamic Server would have treated **sp44** as the default.savepoint object that the TO SAVEPOINT clause of the ROLLBACK statement references.

Related Statements

Related statements: “SAVEPOINT ” on page 2-526 and “ROLLBACK WORK” on page 2-520.

RENAME COLUMN

Use the RENAME COLUMN statement to change the name of a column.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax

```
➤➤—RENAME COLUMN owner.—table—.—old_column—TO—new_column—➤➤
```

Element	Description	Restrictions	Syntax
<i>new_column</i>	Name that you declare here to replace <i>old_column</i>	Must be unique among column names in <i>table</i> . See also “How Triggers Are Affected.”	“Identifier ” on page 5-21
<i>old_column</i>	Column to rename	Must exist within <i>table</i>	“Identifier ” on page 5-21
<i>owner</i>	Owner of the table	Must be the owner of <i>table</i>	“ Owner Name ” on page 5-45
<i>table</i>	Table that contains <i>old_column</i>	Must exist in the current database	“Identifier ” on page 5-21

Usage

You can rename a column of a table if any of the following conditions are true:

- You own the table or have Alter privilege on the table.
- You have the DBA privilege on the database.

In Extended Parallel Server, you cannot rename the columns of a fragmented table if the table is fragmented by range. For more information, see “RANGE Method Clause (XPS)” on page 2-218.

How Views and Check Constraints Are Affected

If you rename a column that appears in a view, the text of the view definition in the **sysviews** system catalog table is updated to reflect the new column name. If you rename a column that appears in a check constraint, the text of the check constraint in the **syschecks** system catalog table is updated to reflect the new column name.

How Triggers Are Affected

If you rename a column that appears within the definition a trigger, it is replaced with the new name only in the following instances:

- When it appears as part of a correlation name inside the FOR EACH ROW action clause of a trigger
- When it appears as part of a correlation name in the INTO clause of an EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement

- When it appears as a triggering column in the UPDATE clause

When the trigger executes, if the database server encounters a column name that no longer exists in the table, an error is returned.

Example of RENAME COLUMN

The following example assigns the new name of `c_num` to the `customer_num` column in the `customer` table:

```
RENAME COLUMN customer.customer_num TO c_num;
```

Related Information

Related statements: “ALTER TABLE” on page 2-44, “CREATE TABLE ” on page 2-191, “CREATE TRIGGER” on page 2-239, “CREATE VIEW ” on page 2-274, and “RENAME TABLE” on page 2-492

RENAME DATABASE

Use the RENAME DATABASE statement to change the name of a database.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax

```

▶▶ RENAME DATABASE owner. old_database TO new_database

```

Element	Description	Restrictions	Syntax
<i>new_database</i>	New name that you declare here for <i>old_database</i>	Must be unique among database names of the current database server; must not be opened by any user when this statement is issued	“Database Name” on page 5-15
<i>old_database</i>	Name that <i>new_database</i> replaces	Must exist on current database server, but it cannot be the name of the current database	“Database Name” on page 5-15
<i>owner</i>	Owner of <i>old_database</i>	Must be the owner of the database	“ Owner Name ” on page 5-45

Usage

You can rename a database if either of the following is true:

- You created the database.
- You have the DBA privilege on the database.

The RENAME DATABASE statement fails with error -9874, however, if the specified database contains any of the following objects:

- a virtual table
- a virtual index
- an R-tree index
- a DataBlade that references the current name of the database in a user-defined primary access method or in a user-defined secondary access method.

You can only rename databases of the database server to which you are currently connected.

You cannot rename a database from inside an SPL routine.

Related Information

Related statement: “CREATE DATABASE ” on page 2-97

For information on how to update the three-part names of jar files after you rename the database, see the *J/Foundation Developer’s Guide*.

RENAME INDEX

Use the RENAME INDEX statement to change the name of an existing index.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```

▶▶—RENAME INDEX—owner.—old_index—TO—new_index—▶▶

```

Element	Description	Restrictions	Syntax
<i>new_index</i>	New name that you declare here for the index	Name must be unique to the database (or to the session, if <i>old_index</i> is on a temporary table)	“Identifier ” on page 5-21
<i>old_index</i>	Index name that <i>new_index</i> replaces	Must exist, but it cannot be any of the following: -- An index on a system catalog table -- A system-generated constraint index -- A Virtual-Index Interface (VII)	“Identifier ” on page 5-21
<i>owner</i>	Owner of index	Must be the owner of <i>old_index</i>	“ Owner Name ” on page 5-45

Usage

You can rename an index if you are the owner of the index or have the DBA privilege on the database.

When you rename an index, the database server changes the index name in the **sysindexes**, **sysconstraints**, **sysobjstate**, and **sysfragments** system catalog tables. (But for an index on a temporary table, no system catalog tables are updated.)

Indexes on system catalog tables cannot be renamed. If you want to change the name of a system-generated index that implements a constraint, use the ALTER TABLE ... DROP CONSTRAINT statement to drop the constraint, and then use the ALTER TABLE ... ADD CONSTRAINT statement to define a new constraint that has the same definition as the constraint that you dropped, but for which you declare the new name.

By default, SPL routines that use the renamed index are reoptimized when they are next executed after the index is renamed. When automatic recompilation is

disabled, however, SPL routines that use the renamed index are automatically recompiled on their next use if the renamed index is associated with a directly referenced table. If the table is only referenced indirectly, however, execution can fail with error -710. For more information about enabling or disabling automatic recompilation after changes to the schema of a referenced table, see the “IFX_AUTO_REPREPARE Environment Option (IDS)” on page 2-620. For more information about the AUTO_REPREPARE configuration parameter, see your *IBM Informix Administrator’s Reference*.

Related Information

Related statements: ALTER INDEX, CREATE INDEX, and DROP INDEX

For a discussion of SPL-routine reoptimization, see your *IBM Informix Performance Guide*.

RENAME SECURITY

Use the RENAME SECURITY statement to change the name of an existing security object. The object can be a security policy, or a security label, or a security label component.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```

▶▶ RENAME SECURITY { POLICY | LABEL policy | LABEL COMPONENT } old_name TO new_name

```

Element	Description	Restrictions	Syntax
<i>new_name</i>	New name that you declare here for the security object	Must be unique among identifiers of security objects in the database, and must be different from <i>old_name</i>	“Identifier ” on page 5-21
<i>old_name</i>	Current name that <i>new_name</i> replaces	Must exist in the database as the identifiers of a security object	“Identifier ” on page 5-21
<i>policy</i>	Security policy of the <i>old_name</i> label	Must be the security policy for which <i>old_name</i> is a security label	“Identifier ” on page 5-21

Usage

Only DBSECADM can issue this statement. The RENAME SECURITY statement replaces the *old_name* with the specified *new_name* in the table of the system catalog in which the renamed security object is registered:

- **syssecpolicies.secpolicyname** for security policies
- **sysseclabels.seclabelname** for security labels
- **sysseclabelcomponents.compname** for security label components.

This statement does not, however, change the numeric value of the **syssecpolicies.secpolicyid**, **sysseclabels.seclabelid**, nor **sysseclabelcomponents.compid** of the renamed security object.

The keyword or keywords that follow the SECURITY keyword identify the type of security object that is being renamed. In the following example, the new identifier **honesty** replaces **best** as the name of a security policy:

```
RENAME SECURITY POLICY best TO honesty;
```

In the following example, the new identifier **transparent** replaces **opaque** as the name of a label for the **honesty** security policy:

```
RENAME SECURITY LABEL honesty.opaque TO transparent;
```

In the next example, the new identifier **accountant** replaces **architect** as the name of a security label component:

```
RENAME SECURITY LABEL COMPONENT architect TO accountant;
```

Related Information

Related statements: "ALTER SECURITY LABEL COMPONENT " on page 2-36, ALTER TABLE, "CREATE SECURITY LABEL " on page 2-175, "CREATE SECURITY LABEL COMPONENT " on page 2-177, "CREATE SECURITY POLICY " on page 2-181, CREATE TABLE, "DROP SECURITY " on page 2-338, GRANT EXEMPTION, GRANT SECURITY LABEL, REVOKE EXEMPTION, and REVOKE SECURITY LABEL

For a discussion of LBAC security objects, see your IBM Informix Dynamic Server Security Guide

RENAME SEQUENCE

Use the RENAME SEQUENCE statement to change the name of a sequence.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```

▶▶ RENAME SEQUENCE owner. old_sequence TO new_sequence

```

Element	Description	Restrictions	Syntax
<i>new_sequence</i>	New name that you declare here for an existing sequence	Must be unique among the names of sequences, tables, views, and synonyms in the database	"Identifier " on page 5-21
<i>old_sequence</i>	Current name of a sequence	Must exist in the current database	"Identifier " on page 5-21
<i>owner</i>	Owner of the sequence	Must be the owner of the sequence	" Owner Name " on page 5-45

Usage

To rename a sequence, you must be the owner of the sequence, have the ALTER privilege on the sequence, or have the DBA privilege on the database.

You cannot use a synonym to specify the name of the sequence.

In a database that is not ANSI compliant, the name of *new_sequence* (or in an ANSI-compliant database, the combination of *owner.new_sequence*) must be unique among sequences, tables, views, and synonyms in the database.

Related Information

Related statements: “ALTER SEQUENCE ” on page 2-40, “CREATE SEQUENCE” on page 2-184, “DROP SEQUENCE ” on page 2-340, “CREATE SYNONYM ” on page 2-187, “DROP SYNONYM ” on page 2-341, “GRANT ” on page 2-395, “REVOKE” on page 2-494, “INSERT ” on page 2-430, “UPDATE” on page 2-689, and “SELECT” on page 2-528

For information about generating values from a sequence, see “NEXTVAL and CURRVAL Operators (IDS)” on page 4-79.

RENAME TABLE

Use the RENAME TABLE statement to change the name of a table.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax

```

▶▶ RENAME TABLE [owner.] old_table TO [ (1) new_owner. ] new_table

```

Notes:

- 1 Extended Parallel Server only

Element	Description	Restrictions	Syntax
<i>new_table</i>	New name for <i>old_table</i>	Must be unique among the names of sequences, tables, views, and synonyms in the database	“Identifier ” on page 5-21
<i>old_table</i>	Name that <i>new_table</i> replaces	Must be the name (not the synonym) of a table that exists in the current database	“Identifier ” on page 5-21
<i>owner</i>	Current owner of the table	Must be the owner of the table.	“ Owner Name ” on page 5-45
<i>new_owner</i>	New owner of the table (XPS only)	Must have DBA privilege on the database	“ Owner Name ” on page 5-45

Usage

To rename a table, you must be the owner of the table, or have the ALTER privilege on the table, or have the DBA privilege on the database.

An error occurs if *old_table* is a synonym, rather than the name of a table.

The renamed table remains in the current database. You cannot use the RENAME TABLE statement to move a table from the current database to another database, nor to rename a table that resides in another database.

In Dynamic Server, you cannot change the table *owner* by renaming the table. An error occurs if you try to specify an *owner*. qualifier for the new name of the table.

In Extended Parallel Server, a user with DBA privilege on the database can change the owner of a table, if the table is local. Both the table name and owner can be changed by a single statement. The following example uses the RENAME TABLE statement to change the owner of a table:

```
RENAME TABLE tro.customer TO mike.customer;
```

When the table owner is changed, you must specify both the old owner and new owner.

Important: When the owner of a table is changed, the existing privileges granted by the original owner are retained.

In Extended Parallel Server, you cannot rename a table that contains a dependent GK index.

In an ANSI-compliant database, if you are not the owner of *old_table*, you must specify *owner.old_table* as the old name of the table.

If *old_table* is referenced by a view in the current database, the view definition is updated in the **sysviews** system catalog table to reflect the new table name. For further information on the **sysviews** system catalog table, see the *IBM Informix Guide to SQL: Reference*.

If *old_table* is a triggering table, the database server takes these actions:

- Replaces the name of the table in the trigger definition but does *not* replace the table name where it appears inside any triggered actions
- Returns an error if the new table name is the same as a correlation name in the REFERENCING clause of the trigger definition

When the trigger executes, the database server returns an error if it encounters a table name for which no table exists.

You can reorganize the **items** table to move the **quantity** column from the fifth position to the third position by following these steps:

1. Create a new table, **new_table**, that contains the column **quantity** in the third position.
2. Fill the table with data from the current **items** table.
3. Drop the old **items** table.
4. Rename **new_table** with the name **items**.

The following example uses the RENAME TABLE statement as the last step:

```
CREATE TABLE new_table
(
  item_num      SMALLINT,
  order_num     INTEGER,
  quantity      SMALLINT,
  stock_num     SMALLINT,
  manu_code     CHAR(3),
  total_price   MONEY(8)
);
INSERT INTO new_table
```

```

SELECT item_num, order_num, quantity, stock_num,
       manu_code, total_price FROM items;
DROP TABLE items;
RENAME TABLE new_table TO items;

```

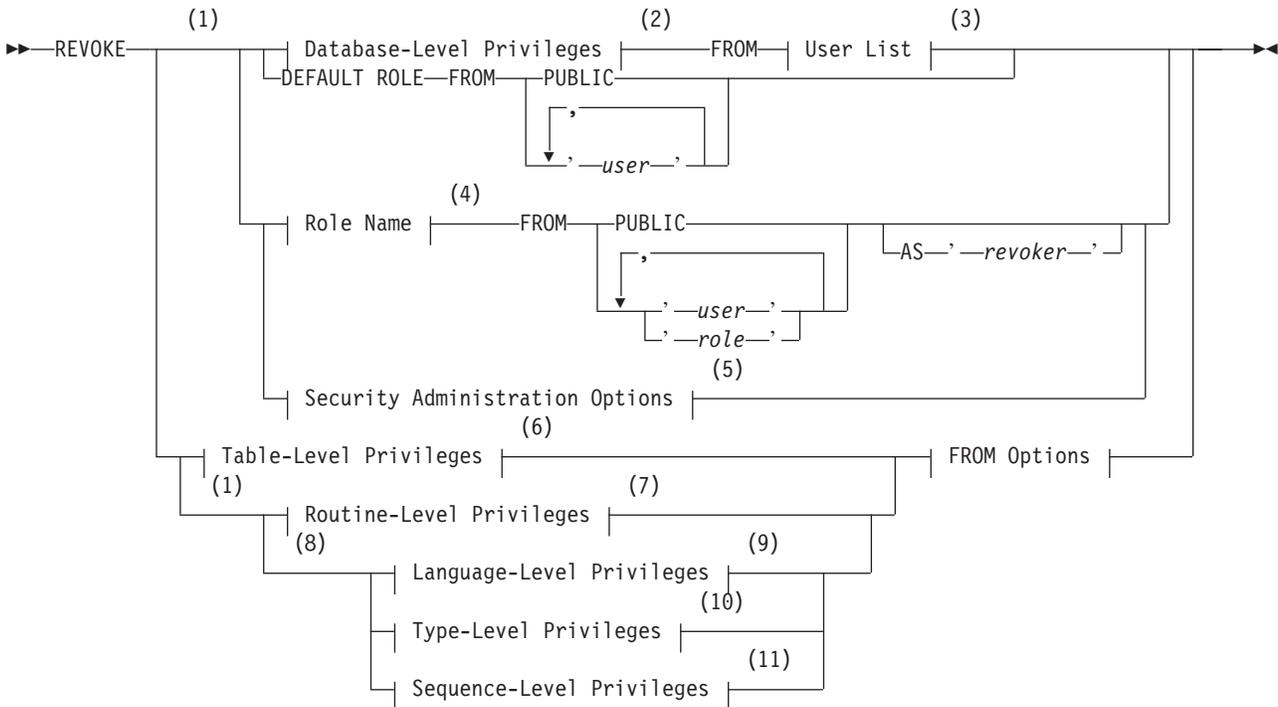
Related Information

Related statements: “ALTER TABLE” on page 2-44, “CREATE TABLE ” on page 2-191, “DROP TABLE ” on page 2-342, and “RENAME COLUMN ” on page 2-487

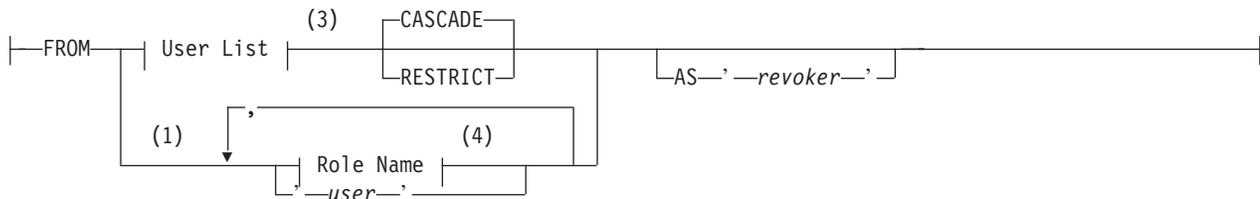
REVOKE

Use the REVOKE statement to cancel access privileges or roles that are held by users, by roles, or by PUBLIC, or to cancel user security labels or exemptions from the rules of security policies.

Syntax



FROM Options:



Notes:

- 1 Informix extension
- 2 See “Database-Level Privileges” on page 2-496

- 3 See "User List" on page 2-504
- 4 See "Role Name" on page 2-505
- 5 See "Security Administration Options" on page 2-508
- 6 See "Table-Level Privileges" on page 2-497
- 7 See "Routine-Level Privileges" on page 2-501
- 8 Dynamic Server only
- 9 See "Language-Level Privileges (IDS)" on page 2-503
- 10 See "Type-Level Privileges (IDS)" on page 2-500
- 11 See "Sequence-Level Privileges (IDS)" on page 2-503

Element	Description	Restrictions	Syntax
<i>revoker</i>	Authorization identifier of the grantor of the privileges to be revoked	Must be grantor of the specified privileges	" Owner Name " on page 5-45
<i>role</i>	Role from which you revoke another role	Must exist	" Owner Name " on page 5-45
<i>user</i>	User whose role (or default role) you cancel	Must exist	" Owner Name " on page 5-45

Usage

To cancel privileges on one or more fragments of a table that has been fragmented by expression, see "REVOKE FRAGMENT " on page 2-517.

You can revoke privileges if any of the following conditions is true for the privileges that you are attempting to revoke on some database object:

- You granted them and did not designate another user as grantor.
- The GRANT statement specified you as grantor.
- You are revoking privileges from PUBLIC on an object that you own, and those privileges were granted by default when you created the object.
- You have database-level DBA privileges and you specify in the AS clause the name of a user who was grantor of the privilege.

The REVOKE statement can cancel any of the following access privileges or roles that a user, or PUBLIC, or a role currently holds:

- Privileges on the database (but a role cannot hold database-level privileges)
- Privileges on a table, synonym, view, or sequence object
- Privileges on a user-defined data type (UDT), a user-defined routine (UDR), or on the SPL language
- A non-default role, or the default role of PUBLIC or of a user.

You cannot revoke privileges from yourself. You cannot revoke *grantor* status from another user. To revoke a privilege that was granted to another user by the AS *grantor* clause of the GRANT statement, you must have the DBA privilege, and you must use the AS clause to specify that user as *revoker*.

If you enclose *revoker*, *role*, or *user* in quotation marks, the name is case sensitive and is stored exactly as you typed it. In an ANSI-compliant database, if you do not use quotation marks as delimiters, the name is stored in uppercase letters.

Database-Level Privileges

Three concentric layers of database-level privileges, Connect, Resource, and DBA, authorize increasing power over database access and control. Only a user with the DBA privilege can grant or revoke database-level privileges.

Database-Level Privileges:



Because of the hierarchical organization of the privileges (as outlined in the privilege definitions that are described later in this section), if you revoke either the Resource or the Connect privilege from a user with the DBA privilege, the statement has no effect. If you revoke the DBA privilege from a user who has the DBA privilege, the user retains the Connect privilege on the database. To deny database access to a user with the DBA or Resource privilege, you must first revoke the DBA or the Resource privilege and then revoke the Connect privilege in a separate REVOKE statement.

Similarly, if you revoke the Connect privilege from a user who has the Resource privilege, the statement has no effect. If you revoke the Resource privilege from a user, the user retains the Connect privilege on the database.

Only users or PUBLIC can hold database-level privileges. You cannot revoke these privileges from a role, because a role cannot hold database level privileges.

The following table lists the keyword for each database-level privilege.

Privilege	Effect
DBA	<p>Has all the capabilities of the Resource privilege and can perform the following additional operations:</p> <ul style="list-style-type: none">• Grant any database-level privilege, including the DBA privilege, to another user.• Grant any table-level privilege to another user or to a role.• Grant a role to a user or to another role.• Revoke a privilege whose grantor you specify as the <i>revoker</i> in the AS clause of the REVOKE statement.• Restrict the Execute privilege to DBAs when registering a UDR.• Execute the SET SESSION AUTHORIZATION statement.• Create any database object.• Create tables, views, and indexes, designating another user as owner of these objects.• Alter, drop, or rename database objects, regardless of who owns it.• Execute the DROP DISTRIBUTIONS option of the UPDATE STATISTICS statement.• Execute DROP DATABASE and RENAME DATABASE statements.

Privilege	Effect
RESOURCE	<p>Lets you extend the structure of the database. In addition to the capabilities of the Connect privilege, the holder of the Resource privilege can perform the following operations:</p> <ul style="list-style-type: none"> • Create new tables. • Create new indexes. • Create new user-defined routines. • Create new data types.
CONNECT	<p>If you have this privilege, you can query and modify data, and modify the database schema if you own the database object that you want to modify. A user holding the Connect privilege can perform the following operations:</p> <ul style="list-style-type: none"> • Connect to the database with the CONNECT statement or another connection statement. • Execute SELECT, INSERT, UPDATE, and DELETE statements, provided that the user has the necessary table-level privileges. • Create views, provided that the user has the Select privilege on the underlying tables. • Create synonyms. • Create temporary tables and create indexes on temporary tables. • Alter or drop a table or an index, if the user owns the table or index (or has the Alter, Index, or References privilege on the table). • Grant privileges on a table, if the user owns the table (or was given privileges on the table with the WITH GRANT OPTION keyword).

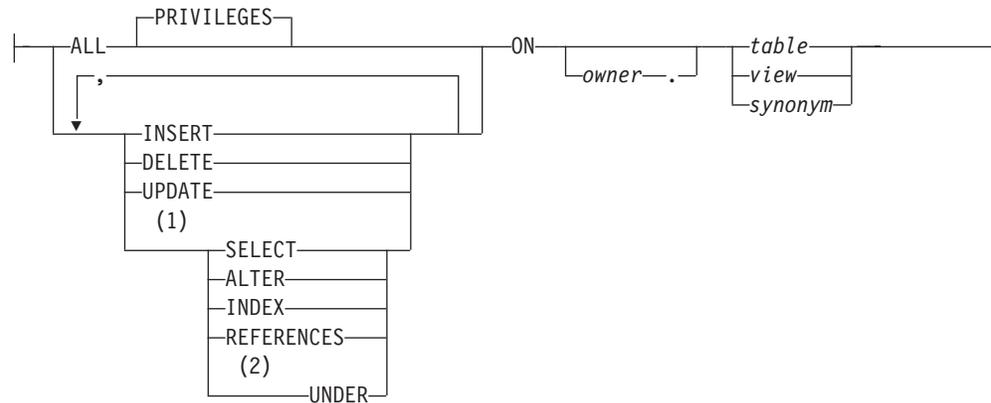
Table-Level Privileges

Table-level privileges, also called *table privileges*, specify which operations a user or role can perform on a table or view in the database. You can use a synonym to specify the table or view on which you grant or revoke table privileges.

Select, Update, and References privileges can be granted on a subset of the columns of a table or view, but can be revoked only for all columns. If Select privileges are revoked from a user for a table that is referenced in the SELECT statement defining a view that the same user owns, then that view is dropped, unless it also includes columns from tables in another database.

Use the following syntax to specifying which table-level privileges to revoke from one or more users or roles:

Table-Level Privileges:



Notes:

- 1 Informix extension
- 2 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>owner</i>	Name of the user who owns the <i>table</i> , <i>view</i> , or <i>synonym</i>	Must be a valid authorization identifier	"Owner Name" on page 5-45
<i>synonym, table, view</i>	Synonym, table, or view on which privileges are granted	Must exist in the current database	"Identifier" on page 5-21

In one REVOKE statement, you can list one or more of the following keywords to specify the privileges on the specified table to be revoked from the users or roles.

Privilege	Effect after REVOKE
INSERT	User cannot insert rows.
DELETE	User cannot delete rows.
SELECT	User cannot display data retrieved by a SELECT statement.
UPDATE	User cannot change column values.
INDEX	User cannot create permanent indexes. You must have the Resource privilege to take advantage of the Index privilege. (But any user who has the Connect privilege can create indexes on temporary tables.)
ALTER	The holder cannot add or delete columns, modify column data types, add or delete constraints, change the locking mode of a table from PAGE to ROW, nor add or drop a corresponding named ROW type table. The user also cannot enable or disable indexes, constraints, nor triggers, as described in "SET Database Object Mode" on page 2-591.
Privilege	Effect after REVOKE
REFERENCES	User cannot reference columns in referential constraints. You must also have the Resource privilege on the database to take advantage of the References privilege on tables. (You can add, however, a referential constraint during an ALTER TABLE statement. without holding the Resource privilege on the database.) Revoking the References privilege disallows cascading DELETE operations.
UNDER (IDS)	User cannot create subtables under a typed table.

Privilege	Effect after REVOKE
ALL	This removes all of the table privileges that are listed above. (Here the PRIVILEGES keyword is optional.)

See also “Table-Level Privileges” on page 2-399.

If a user receives the same privilege from two different grantors and one grantor revokes the privilege, the grantee still has the privilege until the second grantor also revokes the privilege. For example, if both you and a DBA grant the Update privilege on your table to **ted**, both you and the DBA must revoke the Update privilege to prevent **ted** from updating your table.

If user **ted** holds the same privileges through a role or as PUBLIC, however, this REVOKE operation does not prevent **ted** from exercising the Update privilege.

When to Use REVOKE Before GRANT

You can use combinations of REVOKE and GRANT to replace PUBLIC with specific users as grantees, and to remove table-level privileges on some columns.

Replacing PUBLIC with Specified Users: If a table owner grants a privilege to PUBLIC, the owner cannot revoke the same privilege from any specific user. For example, assume PUBLIC has default Select privileges on your **customer** table. Suppose that you issue the following statement in an attempt to exclude **ted** from accessing your table:

```
REVOKE ALL ON customer FROM ted;
```

This statement results in ISAM error message 111, No record found, because the system catalog tables (**syscolauth** or **sysstabauth**) contain no table-level privilege entry for a user named **ted**. This REVOKE operation does not prevent **ted** from keeping all the table-level privileges given to PUBLIC on the **customer** table.

To restrict table-level privileges, first revoke the privileges with the PUBLIC keyword, then re-grant them to some appropriate list of users and roles. The following statements revoke the Index and Alter privileges from all users for the **customer** table, and then grant these privileges specifically to user **mary**:

```
REVOKE INDEX, ALTER ON customer FROM PUBLIC;
GRANT INDEX, ALTER ON customer TO mary;
```

Restricting Access to Specific Columns: Unlike GRANT, the REVOKE statement has no syntax to specify privileges on a subset of columns in a table. To revoke the Select, Update, or References privilege on a column from a user, you must revoke the privilege for all the columns of the table. To provide access to some of the columns on which you previously had granted privileges, issue a new GRANT statement to restore the appropriate privilege on specific columns.

The next example cancels Select privileges for PUBLIC on certain columns:

```
REVOKE SELECT ON customer FROM PUBLIC;
GRANT SELECT (fname, lname, company, city) ON customer TO PUBLIC;
```

In the next example, **mary** first receives the ability to reference four columns in **customer**, then the table owner restricts references to two columns:

```
GRANT REFERENCES (fname, lname, company, city) ON customer TO mary;
REVOKE REFERENCES ON customer FROM mary;
GRANT REFERENCES (company, city) ON customer TO mary;
```

Effect of the ALL Keyword

The ALL keyword revokes all table-level privileges. If any or all of the table-level privileges do not exist for the revokee, REVOKE with the ALL keyword executes successfully but returns the following **SQLSTATE** code:

```
01006--Privilege not revoked
```

For example, assume that user **hal** has the Select and Insert privileges on the **customer** table. User **jocelyn** wants to revoke all table-level privileges from user **hal**. So user **jocelyn** issues the following REVOKE statement:

```
REVOKE ALL ON customer FROM hal;
```

This statement executes successfully but returns **SQLSTATE** code 01006. The **SQLSTATE** warning is returned because both of the following are true:

- The statement succeeds in revoking the Select and Insert privileges from user **hal** because user **hal** had those privileges.
- **SQLSTATE** code 01006 is returned because user **hal** lacked other privileges implied by the ALL keyword, but these privileges were not revoked.

The ALL keyword instructs the database server to revoke everything possible, including nothing. If the user from whom privileges are revoked has no privileges on the table, the REVOKE ALL statement still succeeds, because it revokes everything possible from the user (in this case, no privileges at all).

Effect of the ALL Keyword on UNDER Privilege (IDS): If you revoke ALL privileges on a typed table, the Under privilege is included in the privileges that are revoked. If you revoke ALL privileges on a table that is not based on a ROW type, the Under privilege is not included among the privileges that are revoked. (The Under privilege can be granted only on a typed table.)

Effect of Uncommitted Transactions

The REVOKE statement places an exclusive row lock on the entry in the **systables** system catalog table for the table on which privileges are revoked. This lock is not released until the transaction that contains the REVOKE statement terminates. When another transaction attempts to prepare a SELECT statement against this table while the first transaction is open, the concurrent transaction fails, because the **systables** row for the specified table remains exclusively locked. The attempt to prepare the SELECT statement cannot succeed until after the first transaction is either committed or rolled back.

Type-Level Privileges (IDS)

You can revoke two privileges on data types:

- The Usage privilege on a user-defined data type
- The Under privilege on a named ROW type

Type-Level Privileges:

```
┌───┴───┐
│ USAGE ON TYPE—type_name ───────────────────────────────────┐
│ UNDER ON TYPE—row_type ───────────────────────────────────┘
```

Element	Description	Restrictions	Syntax
<i>row_type</i>	Named ROW type for which to revoke Under privilege	Must exist	Data Type, p. "Data Type" on page 4-21

Element	Description	Restrictions	Syntax
<i>type_name</i>	User-defined type for which to revoke Usage privilege	Must exist	Data Type, p. "Data Type" on page 4-21

Usage Privilege

Any user can reference a built-in data type in an SQL statement, but not a DISTINCT data type that is based on a built-in data type. The creator of a user-defined data type or a DBA must explicitly grant the Usage privilege on the UDT, including a DISTINCT data type based on a built-in data type.

REVOKE with the USAGE ON TYPE keywords removes the Usage privilege that you granted earlier to another user, to PUBLIC, or to a role.

The following statement revokes from user **mark** the privilege of using the **widget** user-defined type:

```
REVOKE USAGE ON TYPE widget FROM mark;
```

Under Privilege

You own any named ROW data type that you create. If you want other users to be able to create subtypes under this named ROW type, you must grant these users the Under privilege on your named ROW type. If you later want to remove the ability of these users to create subtypes under the named ROW type, you must revoke the Under privilege from these users. A REVOKE statement with the UNDER ON TYPE keywords removes the Under privilege that you granted earlier to these users.

For example, suppose that you created a ROW type named **rtype1**:

```
CREATE ROW TYPE rtype1 (cola INT, colb INT);
```

If you want another user named **kathy** to be able to create a subtype under this named ROW type, you must grant the Under privilege on this named ROW type to user **kathy**:

```
GRANT UNDER ON TYPE rtype1 TO kathy;
```

Now user **kathy** can create another ROW type under the **rtype1** ROW type even though **kathy** is not the owner of the **rtype1** ROW type:

```
CREATE ROW TYPE rtype2 (colc INT, cold INT) UNDER rtype1;
```

If you later want to remove the ability of user **kathy** to create subtypes under the **rtype1** ROW type, enter the following statement:

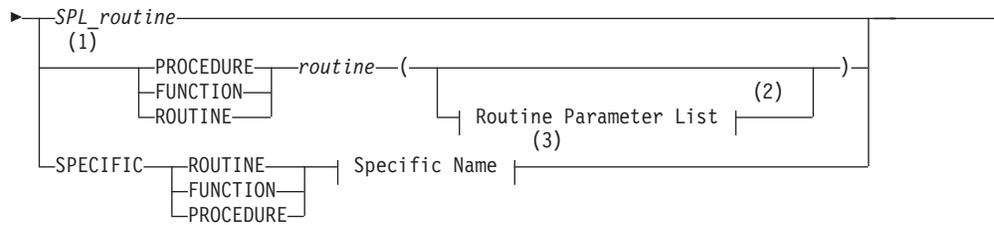
```
REVOKE UNDER ON TYPE rtype1 FROM kathy;
```

Routine-Level Privileges

If you revoke the Execute privilege on a UDR from a user, that user can no longer execute that UDR in any way. For details of how a user can execute a UDR, see "Routine-Level Privileges" on page 2-405.

Routine-Level Privileges:

|—EXECUTE ON—————→



Notes:

- 1 Dynamic Server only
- 2 See “Routine Parameter List” on page 5-67
- 3 See “Specific Name ” on page 5-73

Element	Description	Restrictions	Syntax
<i>routine</i>	A user-defined routine	Must exist	“Identifier ” on page 5-21
<i>SPL_routine</i>	An SPL routine	Must be unique in the database	“Identifier ” on page 5-21

In an ANSI-compliant database, the *owner* name must qualify the *routine* name, unless the user who issues the REVOKE statement is the owner of the routine.

The following example cancels the Execute privilege of user **mark** on the **delete_order** routine that is owned by **luke**:

```
REVOKE EXECUTE on luke.delete_order FROM mark;
```

In Dynamic Server, any negator function for which you grant the Execute privilege requires a separate, explicit, REVOKE statement.

When you create a UDR under any of the following circumstances, PUBLIC is not granted Execute privilege by default. Therefore you must explicitly grant the Execute privilege before you can revoke it:

- You create the UDR in an ANSI-compliant database.
- You have DBA privilege and specify DBA after the CREATE keyword to restrict the Execute privilege to users with the DBA database-level privilege.
- The **NODEFDAC** environment variable is set to yes to prevent PUBLIC from receiving any privileges that are not explicitly granted.

But if you create a UDR with none of those conditions in effect, PUBLIC can execute your UDR without the GRANT EXECUTE statement. To limit who can executes your UDR, revoke Execute privilege FROM PUBLIC, and grant it to users (see “User List” on page 2-504) or roles (see “Role Name” on page 2-505).

In Dynamic Server, if two or more UDRs have the same name, use a keyword from this list to specify which of those UDRs a user list can no longer execute.

Keyword

UDR for Which Execution by the User is Prevented

SPECIFIC

The UDR identified by *specific name*

FUNCTION

Any function with the specified *routine name* (and parameter types that match *routine parameter list*, if specified)

PROCEDURE

Any procedure with the specified *routine name* (and parameter types that match *routine parameter list*, if specified)

ROUTINE

Functions or procedures with the specified *routine name* (and parameter types that match *routine parameter list*, if specified)

Language-Level Privileges (IDS)

A user must have the Usage privilege on SPL to register a UDR written in SPL.

Language-Level Privileges:

```
|—USAGE ON LANGUAGE—SPL—|
```

When a user registers a UDR that is written in SPL, Dynamic Server verifies that the user has the Usage privilege on the SPL language. If the user does not, the CREATE FUNCTION or CREATE PROCEDURE statement fails with an error. (The C language and the Java language do not require the Usage privilege.)

To revoke the Usage privilege on the SPL language from a user or role, issue a REVOKE statement that includes the USAGE ON LANGUAGE SPL keywords. If this statement succeeds, any user or role that you specify in the FROM clause can no longer register UDRs that are written in the specified language. For example, if you revoke the default Usage privilege on SPL from PUBLIC, the ability to create SPL routines is taken away from all users:

```
REVOKE USAGE ON LANGUAGE SPL FROM PUBLIC;
```

You can issue a GRANT USAGE ON LANGUAGE statement to restore Usage privilege on SPL to a restricted group, such as to the role named **developers**:

```
GRANT USAGE ON LANGUAGE SPL TO developers;
```

Sequence-Level Privileges (IDS)

Although Dynamic Server implements sequence objects as tables, only the following subset of the table privileges (as described in “Table-Level Privileges” on page 2-399) can be granted or revoked on a sequence:

- Select privilege
- Alter privilege

Use the following syntax to specify privileges to revoke on a sequence object:

Sequence-Level Privileges:

```
|—ALL—ON—'|  
|—ALTER—|  
|—SELECT—|  
|—'owner.'—|  
|—sequence—|  
|—synonym—|
```

Element	Description	Restrictions	Syntax
<i>owner</i>	Owner of the sequence or of its synonym	Must be the owner	" Owner Name " on page 5-45
<i>sequence</i>	Sequence on which to revoke privileges	Must exist	"Identifier " on page 5-21
<i>synonym</i>	Synonym for a sequence object	Must point to a sequence	"Identifier " on page 5-21

The sequence must reside in the current database. (You can qualify the *sequence* or *synonym* identifier with a valid *owner* name, but the name of a remote *database* (or *database@server*) is not valid as a qualifier.) Syntax to revoke sequence-level privileges is an extension to the ANSI/ISO standard for SQL.

Alter Privilege

You can revoke the Alter privilege on a sequence from another user, from PUBLIC, or from a role. The Alter privilege enables a specified user or role to modify the definition of a sequence with the ALTER SEQUENCE statement or to rename the sequence with the RENAME SEQUENCE statement.

The following REVOKE statement cancels any Alter privilege that was granted individually to user **mark** on the **cust_seq** sequence object:

```
REVOKE ALTER ON cust_seq FROM mark;
```

Select Privilege

You can revoke the Select privilege on a sequence from another user, from PUBLIC, or from a role. Select privilege enables a user or role to use the *sequence.CURRVAL* and *sequence.NEXTVAL* in SQL statements to access and to increment the value of a sequence.

The following REVOKE statement cancels any Select privilege that was granted individually to user **mark** on the **cust_seq** sequence object:

```
REVOKE SELECT ON cust_seq FROM mark;
```

ALL Keyword

You can use the ALL keyword to revoke both Alter and Select privileges from another user, from PUBLIC, or from a role.

The following example cancels any Alter and Select privileges that user **mark** holds on the **cust_seq** sequence object:

```
REVOKE ALL ON cust_seq FROM mark;
```

Whether **mark** can still access **cust_seq** after this statement executes depends on whether the user still holds Alter or Select privileges on **cust_seq** that were granted to PUBLIC, or if he holds a role to which unrevoked privileges on **cust_seq** have been granted.

User List

The authorization identifiers (or the PUBLIC keyword) that follow the FROM keyword of REVOKE specify who loses the revoked privileges or revoked roles. If you use the PUBLIC keyword as the user list, the REVOKE statement revokes the specified privileges or roles from PUBLIC, thereby revoking them from all users to whom the privileges or roles have not been explicitly granted, or who do not hold some other role through which they have received the role or privilege.

The *user list* can consist of the authorization identifier of a single user or of multiple users, separated by commas. If you use the PUBLIC keyword as the user list, the REVOKE statement revokes the specified privileges from all users.

User List:



Element	Description	Restrictions	Syntax
<i>user</i>	Login name of a user whose privilege or role you are revoking	Must be a valid authorization identifier	“ Owner Name ” on page 5-45

Spell the user names in the list exactly as they were spelled in the GRANT statement. You can optionally use quotation marks around each user name in the list to preserve the lettercase. In an ANSI-compliant database, if you do not use quotation marks to delimit *user*, the name of the user is stored in uppercase letters unless the ANSIOWNER environment variable was set to 1 before the database server was initialized.

When you specify login names, you can use the REVOKE statement and the GRANT statement to secure various types of database objects selectively. For examples, see “When to Use REVOKE Before GRANT” on page 2-499.

Role Name

Only the DBA or a user who was granted a role WITH GRANT OPTION can revoke a role or its privileges. Users cannot revoke roles from themselves.

Role Name:



Element	Description	Restrictions	Syntax
<i>role</i>	A role with one of these attributes: <ul style="list-style-type: none"> • Loses an existing privilege or role • Is lost by a user or by another role 	Must exist. If enclosed between quotation marks, <i>role</i> is case sensitive.	“ Owner Name ” on page 5-45

Immediately after the REVOKE keyword, the name of a *role* specifies a role to be revoked from the user list. After the FROM keyword, however, the name of a *role* specifies a role from which access privilege (or another role) is to be revoked. The same FROM clause can include both *user* and *role* names if no other REVOKE options conflict with the *user* or *role* specifications. Syntax to revoke privileges on a role or from a role are extensions to the ANSI/ISO standard for SQL.

When you include a *role* after the FROM keyword of the REVOKE statement, the specified privilege (or another role) is revoked from that role, but users who have that role retain any privileges or roles that were granted to them individually.

If you enclose *role* between quotation marks, the name is case sensitive and is stored exactly as you typed it. In an ANSI-compliant database, if you do not use quotation marks as delimiters, the *role* is stored in uppercase letters.

When you revoke a role that was granted to a user with the WITH GRANT OPTION keywords, you revoke both the role and the option to grant it.

The following examples show the effects of REVOKE *role*:

- Remove users or remove another role from inclusion in the specified role:

```
REVOKE accounting FROM mary;  
REVOKE payroll FROM accounting;
```
- Remove one or more access privileges from a role:

```
REVOKE UPDATE ON employee FROM accounting;
```

When you revoke table-level privileges from a role, you cannot include the RESTRICT or CASCADE keywords.

Revoking a Default Role

The DBA or the owner of the database can define a *default role* for one or more users or for PUBLIC with the GRANT DEFAULT ROLE statement. Unlike a non-default role, which does not take effect until the SET ROLE statement activates the role, a default role takes effect automatically when the user connects to the database. The default role can specify a set of access privileges for all the users who are granted that default role. Conversely, the REVOKE DEFAULT ROLE statement cancels the specified *role* as the default role for the specified *user-list*, as in the following example:

```
REVOKE DEFAULT ROLE accounting FROM mary;
```

This statement removes from user **mary** whatever privileges the default **accounting** role had conferred. If **mary** issues the SET ROLE DEFAULT statement, it has no effect until she is granted some new default role.

After you execute REVOKE DEFAULT ROLE specifying one or more users or PUBLIC, any privileges that those users held only through the default role are cancelled. (But this statement does not revoke any privileges that were granted to a user individually, or privileges that were granted to a user through another role, or privileges that PUBLIC holds.)

After REVOKE DEFAULT ROLE successfully cancels the default role of *user*, the default role of *user* becomes NULL, and the default role information is removed from the system catalog. (In this context, NULL and NONE are synonyms.)

No warning is issued if REVOKE DEFAULT ROLE specifies a user who has not been granted a default role.

No options besides the *user-list* are valid after the FROM keyword in the REVOKE DEFAULT ROLE statement.

Revoking the EXTEND Role (IDS)

Only the Database Server Administrator (DBSA), by default user **informix**, can grant the built-in EXTEND role to one or more users or to PUBLIC with the GRANT EXTEND TO *user-list* statement. Only users who have the EXTEND role can create or drop external UDRs that are written in the C or Java languages, both of which support shared libraries. (It is sufficient to hold the EXTEND role; it is not necessary to activate it with the SET ROLE statement for a user to be able to

create and drop external UDRs.) Conversely, the REVOKE EXTEND FROM *user-list* statement cancels the EXTEND role of the specified users, preventing them from creating or dropping any external UDRs, as in the following example:

```
REVOKE EXTEND FROM 'max';
```

This statement prevents user **max** from creating or dropping external UDRs, even if **max** is the owner of a UDR that he subsequently attempts to drop.

In databases for which this security feature is not needed, the DBSA can disable this restriction on who can create or drop external UDRs by setting the IFX_EXTEND_ROLE configuration parameter to OFF in the ONCONFIG file. When IFX_EXTEND_ROLE is set to OFF, users with at least the Resource privilege on the database can create or drop external UDRs. See “Database-Level Privileges” on page 2-496 for information about the Resource privilege.

Revoking Privileges Granted WITH GRANT OPTION

If you revoke from *user* privileges or a role that you granted using the WITH GRANT OPTION keywords, you sever the chain of privileges granted by that *user*. Thus, when you revoke privileges from users or from a role, you also revoke the same privilege resulting from GRANT statements in the following contexts:

- Issued by your grantee
- Allowed because your grantee specified the WITH GRANT OPTION clause
- Allowed because subsequent grantees granted the same privilege or role using the WITH GRANT OPTION clause

The following examples illustrate this situation. You, as the owner of the table **items**, issue the following statements to grant access privileges to user **mary**:

```
REVOKE ALL ON items FROM PUBLIC;  
GRANT SELECT, UPDATE ON items TO mary WITH GRANT OPTION;
```

User **mary** then uses her new privilege to grant users **cathy** and **paul** access to the **items** table:

```
GRANT SELECT, UPDATE ON items TO cathy;  
GRANT SELECT ON items TO paul;
```

Later you revoke privileges on the **items** table from user **mary**:

```
REVOKE SELECT, UPDATE ON items FROM mary;
```

This single statement effectively revokes all privileges on the **items** table from users **mary**, **cathy**, and **paul**.

The CASCADE keyword has the same effect as this default condition.

The AS Clause

Without the AS clause, the user who executes the REVOKE statement must be the grantor of the privilege that is being revoked. The DBA or the owner of the object can use the AS clause to specify another user (who must be the grantor of the privilege) as the revoker of the privileges. The AS clause provides the only mechanism by which privileges can be revoked on a database object whose *owner* is an authorization identifier, such as **informix**, that is not also a valid user account known to the operating system.

In Extended Parallel Server, the *AS revoker* clause is not valid in a REVOKE statement that cancels a *role*.

Effect of CASCADE Keyword on UNDER Privileges (IDS)

If you revoke the Under privilege on a typed table with the CASCADE option, the Under privilege is removed from the specified user, and any subtables created under the typed table by that user are dropped from the database.

If you revoke the Under privilege on a named ROW type with the CASCADE option when that data type is in use, the REVOKE fails. This exception to the default behavior of the CASCADE option occurs because the database server supports the DROP ROW TYPE statement with the RESTRICT keyword only.

For example, assume that user **jeff** creates a ROW type named **rtype1** and grants the Under privilege on that ROW type to user **mary**. User **mary** now creates a ROW type named **rtype2** under ROW type **rtype1** and grants the Under privilege on ROW type **rtype2** to user **andy**. Then user **andy** creates a ROW type named **rtype3** under ROW type **rtype2**.

If user **jeff** now tries to revoke the Under privilege on ROW type **rtype1** from user **mary** with the CASCADE option, the REVOKE statement fails, because ROW type **rtype2** is still in use by ROW type **rtype3**.

Controlling the Scope of REVOKE with the RESTRICT Option

The RESTRICT keyword causes the REVOKE statement to fail when any of the following dependencies exist:

- A view depends on a Select privilege that you are attempting to revoke.
- A foreign-key constraint depends on a References privilege that you attempt to revoke.
- You attempt to revoke a privilege from a user who subsequently granted this privilege to another user or to a role.

REVOKE does not fail if it specifies a user who has the right to grant the privilege to others but has not exercised that right. For example, assume that user **clara** specifies WITH GRANT OPTION when she grants the Select privilege on the **customer** table to user **ted**. Further assume that user **ted**, in turn, grants the Select privilege on the **customer** table to user **rania**. The following statement that **clara** issued has no effect, because **ted** has used his authority to grant the Select privilege:

```
REVOKE SELECT ON customer FROM ted RESTRICT;
```

In contrast, if user **ted** does not grant the Select privilege to **rania** or to any other user, the same REVOKE statement succeeds. Even if **ted** does grant the Select privilege to another user, either of the following statements succeeds:

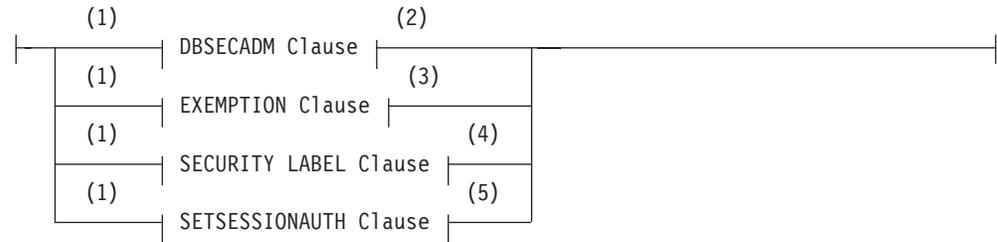
```
REVOKE SELECT ON customer FROM ted CASCADE;  
REVOKE SELECT ON customer FROM ted;
```

Security Administration Options

In conjunction with the GRANT statement, the REVOKE statement supports the discretionary access control (DAC) data security feature of Dynamic Server by specifying which users or roles hold privileges that are required to access the database or objects within the database.

The Security Administration Options of the REVOKE statement, like their counterparts for the GRANT statement, support an additional set of data security features, called label-based access control (LBAC). These features enable Dynamic Server to allow or withhold access to protected data on the basis of a comparing a row security label or column security label that is contained in the data object to the user security label and other credentials that have been granted to the user who is seeking access.

Security Administration Options:



Notes:

- 1 Dynamic Server only
- 2 See “DBSECADM Clause”
- 3 See “EXEMPTION Clause” on page 2-511
- 4 See “SECURITY LABEL Clause” on page 2-513
- 5 See “SETSESSIONAUTH Clause” on page 2-515

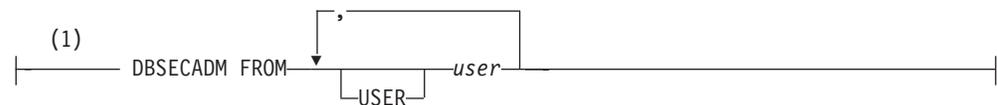
Use of these REVOKE statement security administration options is restricted:

- Only the Database Server Administrator (DBSA), by default user **informix**, can use the REVOKE DBSECADM statement to revoke the DBSECADM role.
- Only a user who holds the DBSECADM role can issue the REVOKE EXEMPTION, REVOKE SECURITY LABEL, or REVOKE SETSESSIONAUTH statements.

DBSECADM Clause

The REVOKE DBSECADM statement prevents the user to whom the DBSECADM role was granted from issuing DDL statements that can create, alter, rename, or drop security objects, including security policies, security labels, and security components.

DBSECADM Clause:



Notes:

- 1 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>user</i>	User from whom the role is to be revoked	Must be the authorization identifier of a user	" Owner Name " on page 5-45

The DBSECADM role is a built-in role that only the DBSA can revoke. Unlike user-defined roles, whose scope is the database in which the role is created, the scope of the DBSECADM role is all of the databases of the Dynamic Server instance. It is not necessary for DBSA to reissue the REVOKE DBSECADM statement in other databases of the same server.

Only a user who holds the DBSECADM role can issue the following SQL statements that create or modify security objects:

- ALTER SECURITY LABEL COMPONENT
- CREATE SECURITY LABEL
- CREATE SECURITY LABEL COMPONENT
- CREATE SECURITY POLICY
- DROP SECURITY LABEL
- DROP SECURITY LABEL COMPONENT
- DROP SECURITY POLICY
- RENAME SECURITY LABEL
- RENAME SECURITY LABEL COMPONENT
- RENAME SECURITY POLICY

Only a user who holds the DBSECADM role can use the following SQL statements to reference tables that are protected by a security policy:

- ALTER TABLE ... ADD SECURITY POLICY
- ALTER TABLE ... ADD ... IDSSECURITYLABEL [DEFAULT *label*]
- ALTER TABLE ... ADD ... [COLUMN] SECURED WITH
- ALTER TABLE ... DROP SECURITY POLICY
- ALTER TABLE ... MODIFY ... [COLUMN] SECURED WITH
- ALTER TABLE ... MODIFY ... DROP COLUMN SECURITY
- CREATE TABLE ... COLUMN SECURED WITH
- CREATE TABLE ... IDSSECURITYLABEL [DEFAULT *label*]
- CREATE TABLE ... SECURITY POLICY

The following GRANT and REVOKE statements also cannot be issued by a user who does not hold the DBSECADM role:

- GRANT EXEMPTION
- GRANT SECURITY LABEL
- GRANT SETSESSIONAUTH
- REVOKE EXEMPTION
- REVOKE SECURITY LABEL
- REVOKE SETSESSIONAUTH

The USER keyword that can follow the FROM keyword is optional, and has no effect, but any authorization identifier that the DBSA specifies in the REVOKE

DBSECADM statement must be the identifier of an individual user, rather than the identifier of a role. The *user* cannot be the DBSA who issues this REVOKE DBSECADM statement.

In the following example, the DBSA cancels the DBSECADM role of user **niccolo**:
 REVOKE DBSECADM FROM niccolo;

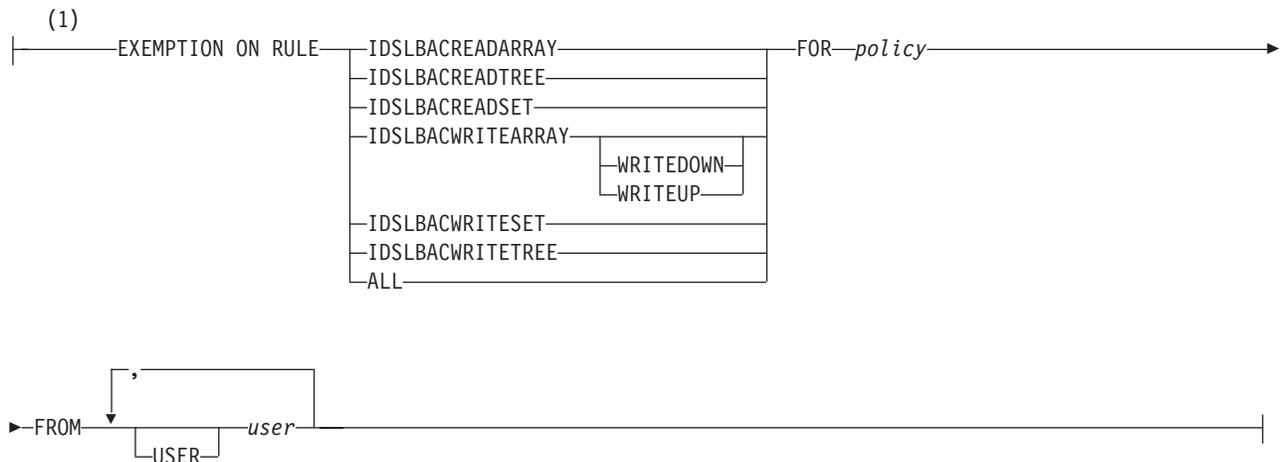
If this statement executes successfully, user **niccolo** can no longer perform the operations listed above.

After the DBSECADM role is revoked, only the DBSA can grant it again to the user from whom it was revoked.

EXEMPTION Clause

The REVOKE EXEMPTION statement modifies the security credentials of the specified user (or list of users) by enabling one or all of the rules of a specified security policy from which the user had been exempt.

EXEMPTION Clause:



Notes:

- 1 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>policy</i>	Security policy for which the exemption is revoked	Must exist in the database	"Identifier " on page 5-21
<i>user</i>	User to whom the exemption is to be revoked	Must be the authorization identifier of a user	" Owner Name " on page 5-45

Only a user who holds the DBSECADM role can issue the REVOKE EXEMPTION statement.

Rules on Which Exemptions Are Revoked: The keyword that follows the ON keyword specifies the predefined access rule of the security policy (whose identifier follows the FOR keyword) for which an exemption is cancelled. The access rule for which exemption is revoked applies when a table that is protected by the specified policy is accessed by a user from whom the exemption is revoked.

For descriptions of the predefined rules for read access and for write access that are associated with a security policy, see the section “Rules Associated with a Security Policy” on page 2-182.

The following keywords of the REVOKE EXEMPTION statement identify specific **IDSLBACRULES** rules that this statement can apply to formerly exempt users:

- **IDSLBACREADARRAY** applies to the user the **IDSLBACREADARRAY** rule for the specified security policy. For a user with no exemption, this rule requires that each array component of the user security label must be greater than or equal to the corresponding array component of the data row security label.
- **IDSLBACREADSET** applies to the user the **IDSLBACREADSET** rule for the specified security policy. For a user with no exemption, this rule requires that each set component of the user security label must include the set component of the data row security label
- **IDSLBACREADTREE** applies to the user the **IDSLBACREADTREE** rule for the specified security policy. For a user with no exemption, this rule requires that each tree component of the user security label must include at least one of the elements in the tree component of the data row security label, or else an ancestor of one such element.
- **IDSLBACWRITEARRAY WRITEDOWN** exempts the user from one aspect of the **IDSLBACWRITEARRAY** rule for the specified security policy. The user who loses this exemption cannot write to a row protected by a label that includes an array component level below the level in the label of the user.
- **IDSLBACWRITEARRAY WRITEUP** exempts the user from one aspect of the **IDSLBACWRITEARRAY** rule for the specified security policy. The user who loses this exemption cannot write to a row protected by a label that includes an array component level above the level in the label of the user.
- **IDSLBACWRITEARRAY** (with no **WRITEDOWN** or **WRITEUP** keyword) applies to the user the **IDSLBACWRITEARRAY** rule for the specified security policy. The user who loses this exemption cannot write to a row whose array component level is above or below the level in the label of the user. .
- **IDSLBACWRITASET** applies to the user the **IDSLBACWRITASET** rule for the specified security policy. For a user with no exemption, that rule requires that each set component of the user security label must include the set component of the data row security label
- **IDSLBACWRITETREE** applies to the user the **IDSLBACWRITETREE** rule for the specified security policy. For a user with no exemption, that rule requires that each tree component of the user security label must include at least one of the elements in the tree component of the data row security label, or the ancestor of one such element.
- **ALL** revokes an exemption from all **IDSLBACRULES** rules for the specified security policy.

In the following example, **DBSECADM** revokes an exemption from all of the rules of the **MegaCorp** security policy from users **manoj** and **sam**:

```
REVOKE EXEMPTION ON RULE ALL FOR MegaCorp FROM manoj, sam;
```

Security Policies and Grantees of Exemptions: An exemption applies only to the rules of a single security policy. whose name follows the **FOR** keyword. Because a protected table can have multiple security labels, but no more than one security policy, revocation of an exemption can prevent a user with insufficient security credentials from accessing data in tables that are protected by the specified security policy.

The REVOKE EXEMPTION statement fails with an error if the specified policy does not exist in the database.

The USER keyword that can follow the FROM keyword is optional, and has no effect, but any authorization identifier specified in the REVOKE EXEMPTION statement must be the identifier of an individual user, rather than the identifier of a role. This *user* cannot be the DBSECADM who issues the same REVOKE EXEMPTION statement.

In the following example, DBSECADM revokes an exemption from user **lynette** for rule **IDSLBACREADARRAY** of the **MegaCorp** security policy:

```
REVOKE EXEMPTION ON RULE IDSLBACREADARRAY FOR MegaCorp FROM lynette;
```

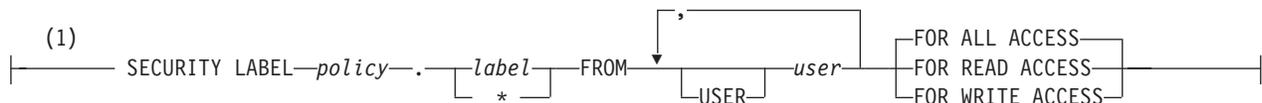
This exemption restores the read access rules for all array components for subsequent read operations that user **lynette** attempts on tables protected by security labels of the specified policy.

When the REVOKE EXEMPTION statement successfully cancels an exemption of a user, the database server updates the **syssecpolicyexemptions** table of the system catalog to unregister the revoked exemption (or multiple exemptions, if several users are listed after the FROM keyword).

SECURITY LABEL Clause

The REVOKE SECURITY LABEL statement cancels a security label (or all the security labels of a specified security policy) held by one or more users.

SECURITY LABEL Clause:



Notes:

- 1 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>label</i>	Name of an existing security label	Must exist as a label for the specified security <i>policy</i>	"Identifier " on page 5-21
<i>policy</i>	The security policy of this <i>label</i>	Must already exist in the database	"Identifier " on page 5-21
<i>user</i>	User from whom the label is revoked	Must be the authorization identifier of a user	" Owner Name " on page 5-45

Only a user who holds the DBSECADM role can issue the REVOKE SECURITY LABEL statement.

A security label is a database object that is always associated with a security policy. That policy defines the set of valid security components that make up the security label. The label stores a set of one or more values for each component of the security policy.

The DBSECADM can associate a security label with the following entities:

- A column of a database table, which a *column security label* can protect
- A row of a database table, which a *row security label* can protect
- A user, whose *user security label* (and any exemptions from rules of the security policy that have been granted to the user) are called the *security credentials* of the user.

When a user who holds a security label for a specific security policy attempts to access a row that is protected by a row security label of the same security policy, the database server compares the sets of values in the user security label and in the row security label in determining whether or not the user should be allowed to access the data. Similarly, LBAC takes into account the user security label and the column security label in determining whether or not the credentials of the user should be allowed to access a protected column.

The GRANT SECURITY LABEL and REVOKE SECURITY LABEL statements enable DBSECADM to control the association of a user with a label. (Data values in a protected table are associated with a row security label or with a column security label by options to the CREATE TABLE or ALTER TABLE statements that only DBSECADM can execute, rather than by the GRANT SECURITY LABEL statement.)

Immediately following the LABEL keyword, the asterisk (*) symbol in the *policy.** specification instructs the database server to revoke every security label of the *policy*. If instead of an asterisk you specify *policy.label*, that *label* must be the name of a security label of the specified *policy*. In this case, if the statement is successful, only that security label is revoked from the user list.

The USER keyword that can follow the FROM keyword is optional, but any authorization identifier specified in the REVOKE SECURITY LABEL statement must be the identifier of an individual user, rather than the identifier of a role.

Access Specifications: The list of users from whom the security label is revoked can optionally be followed by keywords that specify the type of access to data that the security policy of the label protects

- FOR WRITE ACCESS

These keywords restrict the label to the write access rules of **IDSLBACRULES**, namely **IDLSBACWRITEARRAY**, **IDLSBACWRITESET**, and **IDLSBACWRITETREE**.

- FOR READ ACCESS

These keywords restrict the label to the read access rules of **IDSLBACRULES**, namely **IDLSBACWREADARRAY**, **IDLSBACREADSET**, and **IDLSBACREADTREE**.

- FOR ALL ACCESS

These keywords apply the label to all of the read and write access rules that are listed above. If the REVOKE SECURITY LABEL statement includes no FOR ... ACCESS specification, this option takes effect as the default.

For more information about these **IDSLBACRULES** rules for label-based read and write access, see “Rules Associated with a Security Policy” on page 2-182. For information about exemptions to these rules that can be granted for a specific security policy, see “Rules on Which Exemptions Are Revoked” on page 2-511.

Examples of Revoking User Security Labels:

The following three statements create three security label components called **level**, **compartments**, and **groups** respectively:

```
CREATE SECURITY LABEL COMPONENT
  level ARRAY ['TS','S','C','U'];

CREATE SECURITY LABEL COMPONENT
  compartments SET {'A','B','C','D'};

CREATE SECURITY LABEL COMPONENT
  groups TREE ('G1' ROOT,
              'G2' UNDER ROOT,
              'G3' UNDER ROOT);
```

The following statement creates a security policy called **secPolicy** based on the three components above:

```
CREATE SECURITY POLICY secPolicy COMPONENTS
  level, compartments, groups;
```

The following statement creates a security label called **secLabel1**:

```
CREATE SECURITY LABEL secPolicy.secLabel1
  COMPONENT level 'S',
  COMPONENT compartments 'A', 'B',
  COMPONENT groups 'G2';
```

The following statement grants this security label for read access to user **sam**:

```
GRANT SECURITY LABEL secPolicy.secLabel1
  TO sam FOR READ ACCESS;
```

The following statement revokes the security label for read access from user **sam**.

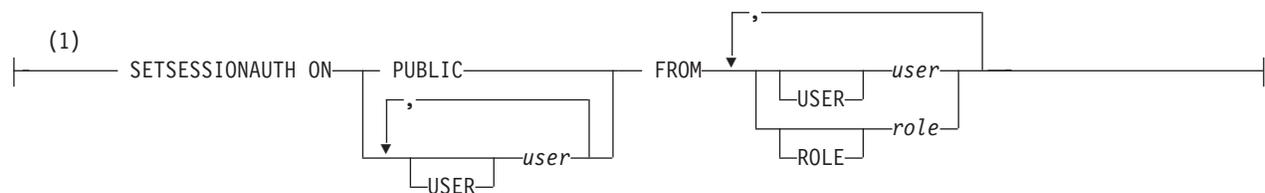
```
REVOKE SECURITY LABEL secPolicy.secLabel1
  FROM sam FOR READ ACCESS;
```

When the REVOKE SECURITY LABEL statement successfully cancels a security label that was held by a user, the database server updates the **sysseclabelauth** table of the system catalog to remove the user from the list of those who hold that security label.

SETSESSIONAUTH Clause

The REVOKE SETSESSIONAUTH statement revokes the SETSESSIONAUTH privilege from one or more users or roles. The SETSESSIONAUTH privilege allows users who also hold the DBA privilege to use the SET SESSION AUTHORIZATION statement to set the session authorization to one of a set of specified users.

SETSESSIONAUTH Clause:



Notes:

- 1 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>role</i>	Role from which the privilege is to be revoked	Must be the authorization identifier of a role	“ Owner Name ” on page 5-45
<i>user</i>	After the FROM keyword, a user from whom the privilege is to be revoked. After the ON keyword, a user whose identity the grantee can specify in the SET AUTHORIZATION statement.	Must be the authorization identifier of a user	“ Owner Name ” on page 5-45

Only a user who holds the DBSECADM role can revoke the SETSESSIONAUTH privilege.

The user or PUBLIC specification that follows the ON keyword specifies whose identity the grantee of the SETSESSIONAUTH privilege is no longer able to assume while using the SET SESSION AUTHORIZATION statement. This can be a user or PUBLIC, but not a role. If PUBLIC is specified, then the grantee of the privilege no longer has the ability to assume the identity of an arbitrary database user.

The USER and ROLE keywords that can follow the FROM keyword are optional. Neither the *user* nor the *role* can be the holder of the DBSECADM role who issues the REVOKE SETSESSIONAUTH statement. The FROM clause cannot specify PUBLIC.

The following example grants to user **sam** the ability to set the session authorization to users **lynette** and **manor**:

```
REVOKE SETSESSIONAUTH ON lynette, manoj TO sam;
```

The next example revokes from user **lynette** the ability to set the session authorization to PUBLIC:

```
REVOKE SETSESSIONAUTH ON PUBLIC FROM lynette;
```

The PUBLIC scope of the privilege that this statement revokes had enabled user **lynette** to assume the access privileges and the security credentials of any user whose name she specified in the SET SESSION AUTHORIZATION statement.

Related Information

Related Statements: “GRANT ” on page 2-395, “ GRANT FRAGMENT” on page 2-423, and “REVOKE FRAGMENT ” on page 2-517

For information about roles, see the following statements: “CREATE ROLE ” on page 2-164, “DROP ROLE ” on page 2-334, and “SET ROLE ” on page 2-647.

For a discussion of discretionary access privileges, see the *IBM Informix Database Design and Implementation Guide*.

For a discussion of LBAC security objects, see your IBM Informix Dynamic Server Security Guide.

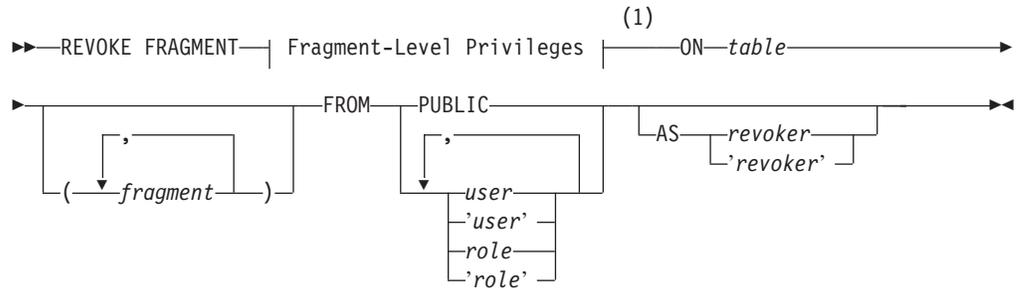
For a discussion of how to embed GRANT and REVOKE statements in programs, see the *IBM Informix Guide to SQL: Tutorial*.

REVOKE FRAGMENT

Use the REVOKE FRAGMENT statement to revoke from one or more users or roles the Insert, Update, or Delete fragment-level privileges that were granted on individual fragments of a table that has been fragmented by expression.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 See "Fragment-Level Privileges" on page 2-518

Element	Description	Restrictions	Syntax
<i>fragment</i>	Name of partition or dbspace that stores one fragment. Default is all fragments of <i>table</i> .	Must exist and must store a fragment of <i>table</i>	"Identifier " on page 5-21
<i>revoker</i>	User (who is not executing this statement) who was grantor of privileges to be revoked	Must be grantor of the fragment-level privileges	" Owner Name " on page 5-45
<i>role</i>	Role from which privileges are to be revoked	Must exist in the database	" Owner Name " on page 5-45
<i>table</i>	Fragmented table whose fragment-level privileges are to be revoked	Must exist and must be fragmented by expression	"Database Object Name " on page 5-16
<i>user</i>	User from whom privileges are to be revoked	Must be a valid authorization identifier	" Owner Name " on page 5-45

Usage

The REVOKE FRAGMENT statement is a special case of the REVOKE statement for assigning privileges on table fragments. Use the REVOKE FRAGMENT statement to revoke the Insert, Update, or Delete privilege on one or more table fragments from one or more users or roles. The DBA can use this statement to revoke privileges on a fragment whose owner is another user.

The REVOKE FRAGMENT statement is valid only for tables that are fragmented by an expression-based distribution scheme. For an explanation of this fragmentation strategy, see "Expression Distribution Scheme" on page 2-15.

Specifying Fragments

If you specify no *fragment*, the privileges are revoked for all fragments of *table*. You can specify one fragment or a comma-separated list of fragments enclosed between parentheses that immediately follow the ON *table* specification.

Each *fragment* must be referenced by the name of the partition where it is stored. If you did not declare an explicit identifier when you created the partition, its name defaults to the name of the dbspace that contains the partition.

After a dbspace is renamed successfully by the onspaces utility, only the new name is valid. Dynamic Server automatically updates existing fragmentation strategies in the system catalog to substitute the new dbspace name, but you must specify the new name in REVOKE FRAGMENT statement to reference a fragment whose default name is the name of a renamed dbspace.

The FROM Clause

You can specify the PUBLIC keyword to revoke the specified fragment-level privileges from PUBLIC, thereby revoking the privileges from all users to whom the privileges have not been explicitly granted, or who do not hold a role through which they have received the privileges.

If you enclose *user* or *role* in quotation marks, the name is case sensitive and is stored exactly as you typed it. In an ANSI-compliant database, if you do not use quotation marks around *user* or around *role*, the name is stored in uppercase letters by default, although you can set the ANSIOWNER environment variable to preserve lowercase characters in *owner* specifications.

When you include a *role* in the FROM clause of REVOKE FRAGMENT, the specified fragment privilege is revoked from that role. Users who have that role, however, retain any fragment privileges they hold that were granted to them individually or to PUBLIC.

Fragment-Level Privileges

The keyword or keywords that follow the FRAGMENT keyword specify *fragment-level privileges*, which are a logical subset of table-level privileges:

Fragment-Level Privileges:



You can revoke fragment-level privileges individually or in combination. The following keywords specify the fragment-level privileges that you can revoke.

Keyword

Effect

INSERT

Prevents the user from inserting rows in the fragment

DELETE

Prevents the user from deleting rows in the fragment

UPDATE

Prevents the user from updating rows in the fragment

ALL Cancels Insert, Delete, and Update privileges on a fragment

If you specify the **ALL** keyword in a **REVOKE FRAGMENT** statement, the specified users and roles lose all fragment-level privileges that they currently possess on the specified fragments. For example, assume that a user currently has the Update privilege on one fragment of a table. If you use the **ALL** keyword to revoke all current privileges on this fragment from this user, the user loses the Update privilege that he or she had on this fragment.

For the distinction between fragment-level and table-level privileges, see the sections “Definition of Fragment-Level Authorization” on page 2-424 and “Effect of Fragment-Level Authorization in Statement Validation” on page 2-425.

The AS Clause

Without the **AS** clause, the user who executes the **REVOKE** statement must be a grantor of the privilege that is being revoked. The DBA or the owner of the fragment can use the **AS** clause to specify another user (who must be the grantor of the privilege) as the revoker of privileges on a fragment.

The **AS** clause provides the only mechanism by which privileges can be revoked on a fragment whose *owner* is an authorization identifier that is not a valid user account known to the operating system.

Examples of the REVOKE FRAGMENT Statement

Examples that follow are based on the **customer** table. They all assume that the **customer** table is fragmented by expression into three fragments that reside in partitions that are named **part1**, **part2**, and **part3**.

Revoking Privileges on One Fragment

The following statement revokes the Update privilege on the fragment of the **customer** table in **part1** from user **ed**:

```
REVOKE FRAGMENT UPDATE ON customer (part1) FROM ed;
```

The following statement revokes the Update and Insert privileges on the fragment of the **customer** table in **part1** from user **susan**:

```
REVOKE FRAGMENT UPDATE, INSERT ON customer (part1) FROM susan;
```

The following statement revokes all privileges currently granted to user **harry** on the fragment of the **customer** table in **part1**:

```
REVOKE FRAGMENT ALL ON customer (part1) FROM harry;
```

Revoking Privileges on More Than One Fragment

The following statement revokes all privileges currently granted to user **millie** on the fragments of the **customer** table in **part1** and **part2**:

```
REVOKE FRAGMENT ALL ON customer (part1, part2) FROM millie;
```

Revoking Privileges from More Than One User

The following statement revokes all privileges currently granted to users **jerome** and **hilda** on the fragment of the **customer** table in **part3**:

```
REVOKE FRAGMENT ALL ON customer (part3) FROM jerome, hilda;
```

Revoking Privileges Without Specifying Fragments

The following statement revokes all current privileges from user **mel** on all fragments for which this user currently has privileges:

```
REVOKE FRAGMENT ALL ON customer FROM mel;
```

Related Statements

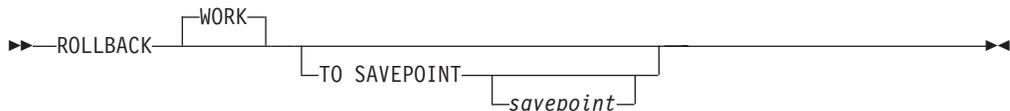
Related statements: “GRANT FRAGMENT” on page 2-423 and “REVOKE” on page 2-494

For a discussion of fragment-level and table-level privileges, see the section “Fragment-Level Privileges” on page 2-424. See also the *IBM Informix Database Design and Implementation Guide*.

ROLLBACK WORK

Use the ROLLBACK WORK statement to cancel all or part of the current transaction intentionally, undoing any changes that occurred since the beginning of the transaction, or between the ROLLBACK WORK statement and a specified or default savepoint object. The ROLLBACK WORK statement restores the database to its state that existed before the cancelled portion of the transaction began.

Syntax



Element	Description	Restrictions	Syntax
<i>savepoint</i>	Name of the savepoint that delimits the scope of the rollback	Must exist in the current transaction.	“Identifier ” on page 5-21

Usage

The ROLLBACK WORK statement is valid only in databases that support transaction logging. Only logged operations can be rolled back. Use ROLLBACK WORK only at the end of a multistatement operation.

In a database that is not ANSI-compliant, the BEGIN WORK statement starts a transaction. You can end a transaction with the COMMIT WORK statement or cancel all or part of the transaction with the ROLLBACK WORK statement. If you issue the ROLLBACK WORK statement when no transaction is pending in a database that is not ANSI-compliant, Dynamic Server issues an error.

In an ANSI-compliant database, multistatement transactions are implicit. You do not need to mark the beginning of a transaction with the `BEGIN WORK` statement. You only need to mark the end of each transaction with a `COMMIT WORK` statement or cancel the transaction with a `ROLLBACK WORK` statement. If you issue the `ROLLBACK WORK` statement when no transaction is pending, the statement is accepted but has no effect.

The `ROLLBACK WORK` statement restores the database to the state that existed before the cancelled portion of the transaction began. Unless you include the `TO SAVEPOINT` keywords, `ROLLBACK WORK` cancels the entire transaction.

The `ROLLBACK WORK` statement releases all row and table locks that the cancelled transaction holds.

In Informix ESQL/C and SPL, the `ROLLBACK WORK` statement closes all open cursors except those that are declared as *hold cursors* by including the `WITH HOLD` keywords. Hold cursors remain open after a transaction is committed or rolled back.

If you use the `ROLLBACK WORK` statement within an SPL routine that the `WHENEVER` statement calls, specify `WHENEVER SQLERROR CONTINUE` and `WHENEVER SQLWARNING CONTINUE` before the `ROLLBACK WORK` statement. This step prevents the program from looping if the `ROLLBACK WORK` statement encounters an error or a warning.

If a program terminates abnormally, the current transaction is implicitly rolled back.

WORK Keyword

The `WORK` keyword is optional in a `ROLLBACK WORK` statement. The following two statements are equivalent:

```
ROLLBACK;
```

```
ROLLBACK WORK;
```

TO SAVEPOINT Clause

The optional `TO SAVEPOINT` clause specifies a partial rollback. This clause can restrict the scope of the rollback to the operations of the current savepoint level between the `ROLLBACK` statement and the specified or default savepoint. If no *savepoint* is specified after the `SAVEPOINT` keyword, the rollback ends at the most recently set savepoint within the current savepoint level.

When the `ROLLBACK WORK TO SAVEPOINT` statement executes successfully, any effects of DDL and DML statements that preceded the savepoint persist, but changes to the schema of the database or to its data values by statements that follow the savepoint are cancelled. Any locks acquired by these cancelled statements persist, but are released at the end of the transaction. Any savepoints between the specified savepoint and the `ROLLBACK` statement are destroyed, but the savepoint referenced by the `ROLLBACK` statement (and any savepoints that precede the referenced savepoint) continue to exist. Program control passes to the statement that immediately follows the `ROLLBACK` statement.

If the TO SAVEPOINT clause is omitted, the ROLLBACK statement rolls back the entire transaction, and all savepoints within the transaction are released.

If the specified *savepoint* does not exist in the current transaction, the database server issues an exception.

The TO SAVEPOINT clause is not valid in a ROLLBACK statement that immediately follows the TRUNCATE statement. In this case, the attempted partial rollback fails with an error. To cancel uncommitted changes that the TRUNCATE statement has made to a table, issue ROLLBACK WORK as the next statement, but with no TO SAVEPOINT clause.

The following program fragment rolls back part of the current transaction to a savepoint called **pt109**:

```
BEGIN WORK;
DROP TABLE tab03;
CREATE TABLE tab03 (col1 CHAR(24), col2 DATE);
SAVEPOINT 108;
...
INSERT INTO tab03 VALUES ('First day of autumn', '09/23/2012');
SAVEPOINT pt109;
...
DELETE FROM tab03 WHERE col2 < '12/09/2009';
SAVEPOINT pt110;
...
ROLLBACK TO SAVEPOINT pt109;
```

The ROLLBACK statement in this example has these effects:

- Cancels the DML operation that deleted any rows with **col2** date values earlier than December 9, 2009.
- Releases savepoint **pt110**, and any other savepoints between **pt109** and the ROLLBACK statement.
- Cancels any other changes to the database by operations that follow savepoint **pt109** in the lexical order of SQL statements within the current transaction.

Savepoint **pt108**, however, is not released, because it was set earlier than **pt109** in the transaction. Not cancelled by this partial rollback are the effects of any uncommitted DDL or DML operations of the transaction before savepoint **pt109** was set, including the creation of table **tab03** and the INSERT operation that added a row to that table. These persist after the partial rollback, pending the possibility of another partial rollback to a savepoint, and the eventual commitment or rollback of the entire transaction.

Related Statements

Related statements: “BEGIN WORK ” on page 2-74, “COMMIT WORK ” on page 2-81, “RELEASE SAVEPOINT” on page 2-485, and “SAVEPOINT ” on page 2-526.

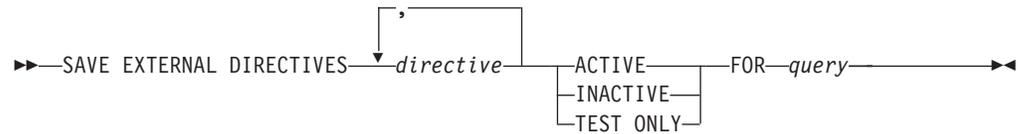
For a discussion of transactions and ROLLBACK WORK, see the *IBM Informix Guide to SQL: Tutorial*.

SAVE EXTERNAL DIRECTIVES

Use the SAVE EXTERNAL DIRECTIVES statement to create external optimizer directives for a specified query, and save the directives in the database. These directives are applied automatically to subsequent instances of the same query.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Element	Description	Restrictions	Syntax
<i>directive</i>	Optimizer directive valid for <i>query</i>	Must be valid for <i>query</i>	"Optimizer Directives" on page 5-35
<i>query</i>	Text of a valid SELECT statement	NULL string is not valid	"SELECT" on page 2-528

Usage

SAVE EXTERNAL DIRECTIVES associates one or more optimizer directives with a query, and stores a record of this association in the **sysdirectives** system catalog table, for subsequent use with queries that match the specified *query* string. This statement establishes an association between the list of optimizer directives and the text of a query, but it does not execute the specified *query*.

Only the DBA or user **informix** can execute SAVE EXTERNAL DIRECTIVES. Optimizer directives that it stores in the database are called *external directives*.

External Optimizer Directives

Like inline optimizer directives that are embedded within a query, external directives that SAVE EXTERNAL DIRECTIVES associates with the text of a query can improve performance in some queries for which the default behavior of the query optimizer is not satisfactory. Unlike inline directives, external directives can be applied without revising or recompiling existing applications.

Enabling or Disabling External Directives for a Session

Dynamic Server ignores external directives if the EXT_DIRECTIVES parameter is set to 0 in the ONCONFIG file or the EXTDIRECTIVES keyword in the SET ENVIRONMENT statement is set to 0, OFF, or off during a session.

In addition, the client system can disable external directives for its current session when the **IFX_EXTDIRECTIVES** environment variable is set to 0.

The following table shows whether external directives are disabled (OFF) or enabled (ON) for various combinations of valid **IFX_EXTDIRECTIVES** settings on the client system and valid EXT_DIRECTIVES configuration parameter settings on Dynamic Server:

Table 2-2. Combinations of IFX_DIRECTIVES Settings and EXT_DIRECTIVES Configuration Parameter Settings

IFX_EXTDIRECTIVES Environment Variable Setting on Client System	EXT_DIRECTIVES Configuration Parameter = 0	EXT_DIRECTIVES Configuration Parameter = 1	EXT_DIRECTIVES Configuration Parameter = 2
IFX_EXTDIRECTIVES not set	OFF	OFF	ON

Table 2-2. Combinations of IFX_DIRECTIVES Settings and EXT_DIRECTIVES Configuration Parameter Settings (continued)

IFX_EXTDIRECTIVES Environment Variable Setting on Client System	EXT_DIRECTIVES Configuration Parameter = 0	EXT_DIRECTIVES Configuration Parameter = 1	EXT_DIRECTIVES Configuration Parameter = 2
IFX_EXTDIRECTIVES = 1	OFF	ON	ON
IFX_EXTDIRECTIVES = 0	OFF	OFF	OFF

If EXT_DIRECTIVES is set to 1 or 2 when the database server is initialized, then the server is enabled for external directives. Individual sessions can enable or disable external directives by setting IFX_EXTDIRECTIVES, as the table shows. Any settings other than 1 or 2 are interpreted as zero, disabling this feature.

When external directives are enabled, the status of individual external directives is specified by the ACTIVE, INACTIVE, or TEST ONLY keywords. (But only queries on which directives are effective can benefit from external directives.)

You can also use the EXTDIRECTIVES option of the SET ENVIRONMENT statement to enable or disable external directives during a session. What you specify using the EXTDIRECTIVES option overwrites the external directive value that is specified in the EXT_DIRECTIVES configuration parameter in the ONCONFIG file.

To overwrite the value for enabling or disabling the external directive in the ONCONFIG file and:

- To enable the external directives during a session, specify 1, on, or ON as the value for SET ENVIRONMENT EXTDIRECTIVES.
- To disable the external directives during a session, specify 0, off, or OFF as the value for SET ENVIRONMENT EXTDIRECTIVES.

To enable the default values specified in the EXT_DIRECTIVES configuration parameter and in the client-side IFX_EXTDIRECTIVES environment variable during a session, specify DEFAULT as the value for the EXTDIRECTIVES option of the SET ENVIRONMENT statement.

For more information on using the EXTDIRECTIVES option of the SET ENVIRONMENT statement, see “SET ENVIRONMENT ” on page 2-615.

The directive Specification

Each *directive* specification in the SAVE EXTERNAL DIRECTIVES statement must follow the syntax of the Optimizer Directives segment, as described in “Optimizer Directives ” on page 5-35. If you specify more than one directive, separate them by a comma (,) symbol, as in the following example:

```
SAVE EXTERNAL DIRECTIVES /** AVOID_INDEX (table1 index1)*/, /** FULL(table1) */
ACTIVE FOR
    SELECT /** INDEX( table1 index1 ) */ col1, col2
    FROM table1, table2 WHERE table1.col1 = table2.col1;
```

This example associates AVOID_INDEX and FULL directives with the specified query. The inline INDEX directive is ignored by the query optimizer when the external directives are applied to a query that matches the SELECT statement.

The ACTIVE, INACTIVE, and TEST ONLY Keywords

You must include one of the ACTIVE, INACTIVE, or TEST ONLY keyword options to enable, disable, or restrict the scope of external directives:

- If external directives are enabled, the ACTIVE keyword applies the list of directives to any subsequent query that matches the *query* string.
- The INACTIVE keyword causes Dynamic Server to ignore the directive. (It is associated with the query in **sysdirectives**, but it is dormant, with no effect.)
- If external directives are enabled, the TEST ONLY keywords apply the directives only to matching queries that the DBA or user **informix** executes. Queries by any other users cannot use TEST ONLY external directives.

An INACTIVE directive has no effect unless the DBA or user **informix** changes the **sysdirectives.active** system catalog column value from zero (INACTIVE) to one (ACTIVE) or two (TEST ONLY) for that directive. External directives do not have SQL identifiers, but the DBA can reference the **sysdirectives.id** column in an UPDATE statement to specify which external directive to update.

Alternatively, the DBA or user **informix** can delete an INACTIVE or TEST ONLY row from **sysdirectives** and use the SET EXTERNAL DIRECTIVES statement to redefine the deleted directive, but now specifying the ACTIVE keyword. This can give other users access to TEST ONLY directives that the DBA has validated.

The query Specification

The *query* specification that follows the FOR keyword in SAVE EXTERNAL DIRECTIVES must specify the syntax of a valid SELECT statement, as described in “SELECT” on page 2-528. If the *query* text also includes any inline optimizer directives, the inline directives are ignored when external directives are applied to the query.

When external directives are enabled and the **sysdirectives** system catalog table is not empty, the database server compares every query with the *query* text of every ACTIVE external directive, and for queries executed by the DBA or user **informix**, with every TEST ONLY external directive. If an external directive has been applied to a query, output from the SET EXPLAIN statement indicates “EXTERNAL DIRECTIVES IN EFFECT” for that query.

The purpose of external directives is to improve the performance of queries that match the *query* string, but the use of such directives can potentially slow other queries, if the query optimizer must compare the *query* strings of a large number of active external directives with the text of every SELECT statement. For this reason, IBM recommends that the DBA not allow the **sysdirectives** table to accumulate more than a few ACTIVE rows. (Another way to avoid unintended performance impact on other queries is to disable this feature.)

If more than one SET EXTERNAL DIRECTIVES statements associate active external directives with the same query, the effect is unpredictable, because the optimizer uses the first **sysdirectives** row whose *query* string matches the query.

Related Statements

For information about optimizer directives and their syntax, see the segment “Optimizer Directives” in “Optimizer Directives ” on page 5-35.

For information about the **sysdirectives** table and the **IFX_EXTDIRECTIVES** environment variable, see the *IBM Informix Guide to SQL: Reference*.

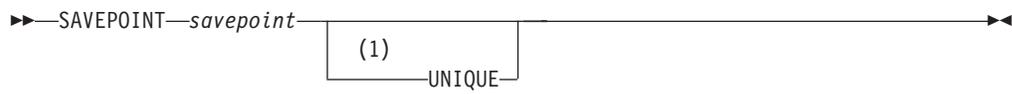
SAVEPOINT

Use the SAVEPOINT statement to declare the name of a new savepoint within the current SQL transaction, and to set the position of the new savepoint within the lexical order of SQL statements within the transaction.

You can use the SAVEPOINT statement in SQL transactions to support error handling with DB-Access and in SPL, C, and Java routines.

This statement is compliant with the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 Informix extension and Dynamic Server only

Element	Description	Restrictions	Syntax
<i>savepoint</i>	Name declared here for the new savepoint	Cannot be the name of an existing unique savepoint in the same savepoint level	“Identifier ” on page 5-21

Usage

The SAVEPOINT statement sets the specified savepoint at the current position in the lexical order of statements within the current transaction. After the SAVEPOINT statement executes successfully, subsequent ROLLBACK TO SAVEPOINT statements that reference this savepoint can cancel any uncommitted changes to the database from logged DML or DDL operations in the current transaction that follow the new savepoint but precede the ROLLBACK TO SAVEPOINT statement.

If an existing savepoint in the same transaction has the same name that the SAVEPOINT statement specifies, the existing savepoint is destroyed, unless one of the following conditions is true:

- The existing savepoint was set in a different savepoint level.
- The existing savepoint name was declared with the UNIQUE keyword option. In this case, the SAVEPOINT statement fails with an error, unless the existing UNIQUE savepoint was set in a different savepoint level.

Destroying a savepoint to reuse its name for another savepoint is not the same as releasing the savepoint. Reusing a savepoint name destroys only one savepoint. Releasing a savepoint with the RELEASE SAVEPOINT statement releases the specified savepoint and all savepoints that have been subsequently set.

The UNIQUE Option

This optional keyword specifies that the application does not intend to reuse the name of this savepoint in another SAVEPOINT statement while this savepoint is active within the current savepoint level.

If a savepoint already exists that was set with the same and with the UNIQUE keyword within the current savepoint level, the SAVEPOINT statement fails with an error, and the existing savepoint is not destroyed.

Savepoint Levels

Dynamic server supports the construct of nested savepoint levels. A single SQL transaction can have multiple savepoint levels. New savepoint levels are automatically created for the duration of execution of an SPL routine or an external UDR. Recursive calls to the same SPL routine or UDR also increment the savepoint level of the current transaction.

A savepoint level ends when the UDR in which it was created finishes execution. When a savepoint level ends, all savepoints within it are automatically released. Any DDL or DML modifications are inherited by the parent savepoint level (that is, by the savepoint level within which the one that just ended was created), and are subject to any savepoint-related statements that are issued against the parent savepoint level.

The following rules apply to actions within a savepoint level:

- Savepoints can only be referenced within the savepoint level in which they are established. You cannot release, destroy, or roll back to a savepoint established outside of the current savepoint level.
- The uniqueness of savepoint names is only enforced within the current savepoint level. The names of savepoints that are active in other savepoint levels can be reused in the current savepoint level without affecting those savepoints in other savepoint levels.

Savepoints in Distributed SQL Transactions

Savepoints are valid in cross-database distributed SQL transactions of a single Dynamic Server instance that supports transactions if all of the participating databases support transaction logging. Savepoints are also supported in cross-server SQL transactions, including operations in high-availability clusters, if all of the participating Dynamic Server instances support savepoints, and all of the databases that are accessed in the transaction use logging.

If any of the participating database servers in a cross-server transaction does not support savepoints, however, and a connection is established between a coordinator that can support savepoints and a subordinate server that cannot, any ROLLBACK TO SAVEPOINT statement within the distributed transaction fails with an error.

Persistence of Savepoints

Savepoints are position markers within SQL transactions, not database objects. An existing savepoint **S** is destroyed by any of the following events within the same transaction:

- A COMMIT WORK or ROLLBACK WORK (without the TO SAVEPOINT clause) statement is executed.
- A RELEASE SAVEPOINT statement is executed that specifies S in the same savepoint level.
- A ROLLBACK TO SAVEPOINT or RELEASE SAVEPOINT statement is executed that specifies a savepoint that was established earlier than S in the same savepoint level.
- A SAVEPOINT statement specifies the same name as S in the same savepoint level, and S was not created with the UNIQUE keyword.

Restrictions on Savepoints

Savepoints and savepoint levels are not supported in the following contexts:

- in databases that do not support transaction logging
- in triggered actions
- in XA global transactions
- in applications or UDRs where the AUTOCOMMIT connection attribute is enabled.

In addition, the SAVEPOINT statement (like the RELEASE SAVEPOINT and ROLLBACK WORK TO SAVEPOINT statements) is not valid in UDRs that are invoked within DML statements, as in the following example:

```
SELECT first_1 foo() FROM systables;
```

Here the `foo()` routine cannot set a savepoint.

Related Statements

Related statements: “RELEASE SAVEPOINT” on page 2-485 and “ROLLBACK WORK” on page 2-520

SELECT

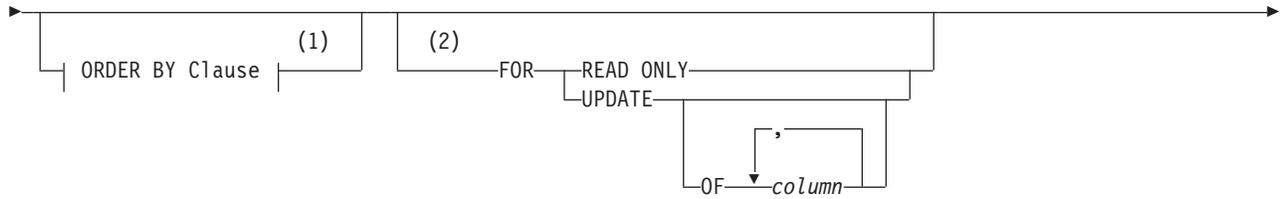
Use the SELECT statement to retrieve values from a database or from an SPL or Informix ESQL/C collection variable. A SELECT operation is called a *query*.

Rows or values that satisfy the specified search criteria of the query are called *qualifying* rows or values. What the query retrieves to its calling context, after applying any additional logical conditions, is called the *result set* of the query. This result set can be empty.

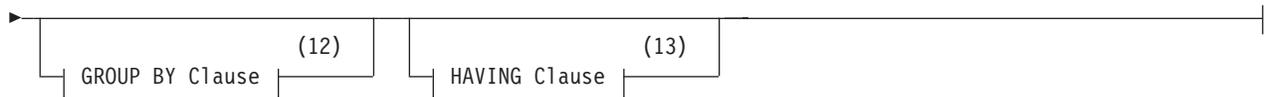
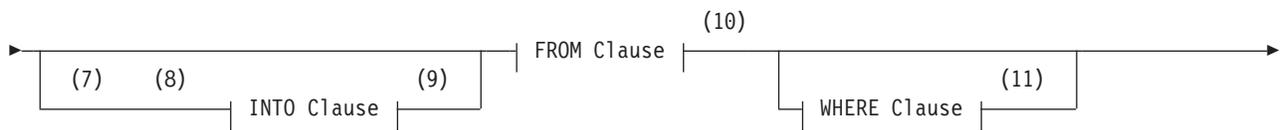
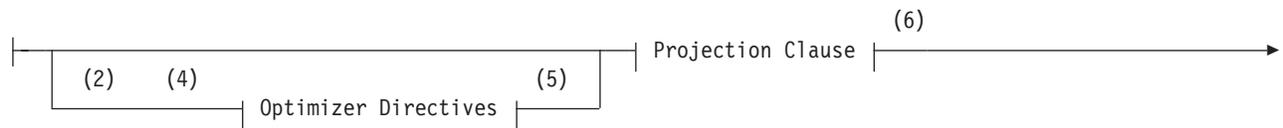
You need the CONNECT privilege to execute a query, as well as the SELECT privilege on the tables from which the query retrieves rows.

Syntax





Select Options:



Notes:

- 1 See "ORDER BY Clause" on page 2-568
- 2 Informix extension
- 3 See "INTO Table Clauses" on page 2-574
- 4 Dynamic Server only
- 5 See "Optimizer Directives " on page 5-35
- 6 See "Projection Clause" on page 2-531
- 7 ESQ/C only
- 8 Stored Procedure Language only
- 9 See "INTO Clause" on page 2-541
- 10 See "FROM Clause" on page 2-544
- 11 See "WHERE Clause of SELECT" on page 2-560
- 12 See "GROUP BY Clause" on page 2-566
- 13 See "HAVING Clause" on page 2-567

Element	Description	Restrictions	Syntax
<i>column</i>	Name of a column that can be updated after a FETCH	Must be in a FROM clause <i>table</i> , but does not need to be in the select list of the Projection clause	"Identifier " on page 5-21

Usage

The SELECT statement can return data from tables in the current database, or in another database of the current database server, or in a database of another database server. Only the SELECT keyword, the Projection clause, and the FROM clause are required specifications.

The SELECT statement can include various basic clauses, which are identified in the following list.

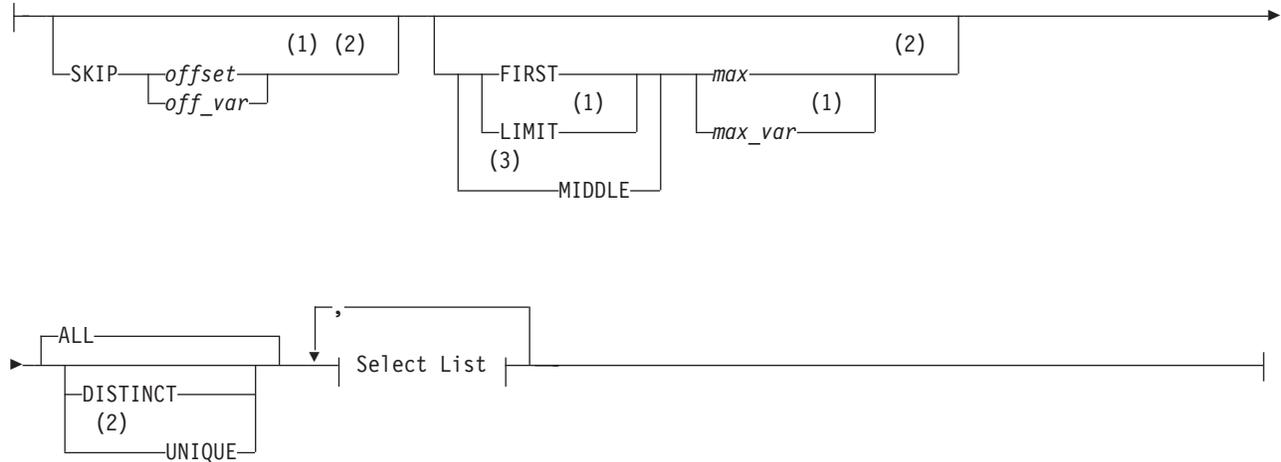
Clause	Page	Effect
" Optimizer Directives " on page 5-35		Specifies how the query should be implemented
"Projection Clause" on page 2-531		Specifies a list of items to be read from the database
"INTO Clause" on page 2-541		Specifies variables to receive the result set
"FROM Clause" on page 2-544		Specifies the data sources of Projection clause items
"Using the ON Clause" on page 2-557		Specifies join conditions as pre-join filters
"WHERE Clause of SELECT" on page 2-560		Sets conditions on qualifying rows and post-join filters
"GROUP BY Clause" on page 2-566		Combines groups of rows into summary results
"HAVING Clause" on page 2-567		Sets conditions on the summary results
"ORDER BY Clause" on page 2-568		Sorts qualifying rows according to column values
"FOR UPDATE Clause" on page 2-572		Enables updating of the result set after a FETCH
"Using the FOR READ ONLY Clause in Read-Only Mode" on page 2-573		Disables updating of the result set after a FETCH
" INTO TEMP Clause" on page 2-575		Puts the result set into a temporary table
" INTO EXTERNAL Clause (XPS)" on page 2-577		Stores the result set in an external table
" INTO SCRATCH Clause (XPS)" on page 2-578		Stores the result set in an unlogged temporary table
"UNION Operator" on page 2-578		Combines the result sets of two SELECT statements
"UNION Operator" on page 2-578		Same as UNION ALL, but discards duplicate rows

Sections that follow describe these clauses of the SELECT statement.

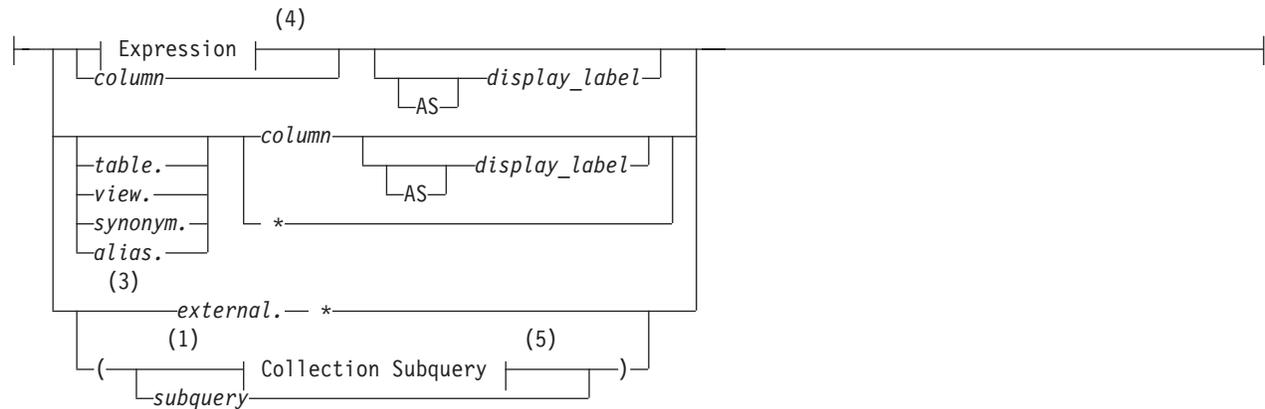
Projection Clause

The Projection clause (sometimes called the *Select clause*) specifies a list of database objects or expressions to retrieve, and can set restrictions on qualifying rows. (The *select list* is sometimes also called the *projection list*.)

Projection Clause:



Select List:



Notes:

- 1 Dynamic Server only
- 2 Informix extension
- 3 Extended Parallel Server only
- 4 See "Expression" on page 4-40
- 5 See "Collection Subquery" on page 4-3

Element	Description	Restrictions	Syntax
<i>alias</i>	Temporary <i>table</i> or <i>view</i> name. See "FROM Clause" on page 2-544.	Valid only if the FROM clause declares the <i>alias</i> for <i>table</i> or <i>view</i>	"Identifier " on page 5-21
<i>column</i>	Column from which to retrieve data	Must exist in a data source that the FROM clause references	"Identifier " on page 5-21
<i>display_label</i>	Temporary name declared here for a <i>column</i> or for an expression	See "Declaring a Display Label" on page 2-541	"Identifier " on page 5-21
<i>external</i>	External table from which to retrieve data	Must exist	"Database Object Name " on page 5-16
<i>max</i>	Integer (> 0) specifying maximum number of rows to return	If <i>max</i> > number of qualifying rows then all matching rows are returned	"Literal Number " on page 4-185
<i>max_var</i>	Host variable or local SPL variable storing the value of max	Same as <i>max</i> ; valid in prepared objects and in SPL routines	Language dependent
<i>offset</i>	Integer (> 0) specifying how many qualifying rows to exclude before the first row of the result set	Cannot be negative. If <i>offset</i> > (number of qualifying rows), then no rows are returned	"Literal Number " on page 4-185
<i>off_var</i>	Host variable or local SPL variable storing the value of offset	Same as <i>offset</i> ; valid in prepared objects and in user-defined routines	Language dependent
<i>subquery</i>	Embedded query	A subquery within the Projection clause cannot include the SKIP, FIRST, INTO TEMP, or the ORDER BY clause.	"SELECT" on page 2-528
<i>table, view, synonym</i>	Name of a table, view, or synonym from which to retrieve data	The synonym and the table or view to which it points must exist	"Database Object Name " on page 5-16

The asterisk (*) specifies all columns in the *table* or *view* in their defined order. To retrieve all columns in another order, or a subset of columns, you must specify individual *column* names explicitly. A solitary asterisk (*) can be a valid Projection clause if the FROM clause specifies only a single data source.

The SKIP, FIRST, LIMIT, MIDDLE, DISTINCT, and UNIQUE specifications can restrict results to a subset of the qualifying rows, as sections that follow explain.

The Order of Qualifying Rows

To execute a query, the database server constructs a query plan and retrieves all qualifying rows that match the WHERE clause conditions. (Here a *row* refers to one set of values, as specified in the select list, from a single record in the table or joined tables that the FROM clause specifies.) If the query has no ORDER BY clause, the qualifying rows are sequenced in the order of their retrieval, which might vary with each execution; otherwise, their sequence follows the ORDER BY specification, as described in "ORDER BY Clause" on page 2-568.

Whether or not the query specifies ORDER BY can affect which qualifying rows are in the result set if the Projection clause includes any of the following options:

- the FIRST option
- the SKIP and LIMIT options (Dynamic Server only)
- the MIDDLE option (Extended Parallel Server only)

Using the SKIP Option (IDS)

The SKIP *offset* option specifies how many of the qualifying rows to exclude, for *offset* an integer in the SERIAL8 range, counting from the first qualifying row. The following example retrieves the values from all rows except the first 10 rows:

```
SELECT SKIP 10 a, b FROM tab1;
```

You can also use a host variable to specify how many rows to exclude. In an SPL routine, you can use an input parameter or a local variable to provide this value.

When you use the SKIP option in a query with an ORDER BY clause, you can exclude the first *offset* rows that have the lowest values according to the ORDER BY criteria. You can also use SKIP to exclude rows with the highest values, if the ORDER BY clause includes the DESC keyword. For example, the following query returns all rows of the **orders** table, except for the fifty oldest orders:

```
SELECT SKIP 50 * FROM orders ORDER BY order_date;
```

Here the result set is empty if there are fewer than 50 rows in the **orders** table. An *offset* = 0 is not invalid, but in that case the SKIP option does nothing.

You can also use the SKIP option to restrict the result sets of prepared SELECT statements, of UNION queries, in queries whose result set defines a collection-derived table, and in the events and actions of triggers.

You can use the SKIP and the FIRST options together to specify which and how many qualifying rows are in the result set, as illustrated by examples in the section “Using the SKIP Option with the FIRST Option (IDS)” on page 2-534.

The SKIP option is not valid in the following contexts:

- In the definition of a view
- In nested SELECT statements
- In subqueries.

Using the FIRST Option

The FIRST *max* option specifies that the result set includes no more than *max* rows (or exactly *max*, if *max* is not greater than the number of qualifying rows). Any additional rows that satisfy the selection criteria are not returned. The following example retrieves at most 10 rows from table **tab1**:

```
SELECT FIRST 10 a, b FROM tab1;
```

Dynamic Server can use a host variable or the value of an SPL input parameter in a local variable to specify *max*.

With an ORDER BY clause, you can retrieve the first *max* qualifying rows. For example, the following query finds the ten highest-paid employees:

```
SELECT FIRST 10 name, salary FROM emp ORDER BY salary DESC;
```

You can use the FIRST option in a query whose result set defines collection-derived table (CDT) within the FROM clause of another SELECT statement. The following query specifies a CDT that has no more than ten rows:

```
SELECT *
  FROM TABLE(MULTISET(SELECT FIRST 10 * FROM employees
                       ORDER BY employee_id)) vt(x,y), tab2
 WHERE tab2.id = vt.x;
```

The FIRST and SKIP keywords are also valid in queries that include table expressions in the FROM clause:

```
SELECT * FROM (SELECT SKIP 2 FIRST 8 col1 FROM tab1 WHERE col1 > 50 );
```

The next example applies the FIRST option to the result of a UNION expression:

```
SELECT FIRST 10 a, b FROM tab1 UNION SELECT a, b FROM tab2;
```

The FIRST option is not valid in any of the following contexts:

- In the definition of a view
- In nested SELECT statements
- In subqueries, except for subqueries that specify table expressions in the FROM clause
- In a singleton SELECT (where *max* = 1) within an SPL routine
- Where embedded SELECT statements are used as expressions

The LIMIT Keyword (IDS)

For Dynamic Server, LIMIT is a keyword synonym for the FIRST keyword in the Projection clause. You cannot, however, substitute LIMIT for FIRST in other syntactic contexts where FIRST is valid, such as in the FETCH statement.

Extended Parallel Server does not support the LIMIT keyword, and it returns an error if FIRST is not immediately followed by a non-negative literal integer.

Using SKIP, FIRST, LIMIT, or MIDDLE as a Column Name

If no integer follows the FIRST keyword, the database server interprets FIRST as a column identifier. For example, if table **T** has columns **first**, **second**, and **third**, the following query would return data from the column named **first**:

```
SELECT first FROM T
```

The same considerations apply to the SKIP and LIMIT keywords of Dynamic Server, and to the MIDDLE keyword of Extended Parallel Server. If no literal integer nor integer variable follows the LIMIT keyword in the Projection clause, Dynamic Server interprets LIMIT as a column name. If no data source in the FROM clause has a column with that name, the query fails with an error.

Using the SKIP Option with the FIRST Option (IDS)

If a Projection clause with the SKIP *offset* option also includes FIRST or LIMIT, the result set begins with the row whose ordinal position is (*offset* + 1) in the set of qualifying rows, rather than with the first row. The row in position (*offset* + *max*) is the last row in the result set, unless there are fewer than (*offset* + *max*) qualifying rows. The following example ignores the first 50 rows from table **tab1**, but returns a result set of at most 10 rows, beginning with the fifty-first row:

```
SELECT SKIP 50 FIRST 10 a, b FROM tab1;
```

The next example uses in a query with SKIP and FIRST to insert no more than five rows from table **tab1** into table **tab2**, beginning with the eleventh row:

```
INSERT INTO tab2 SELECT SKIP 10 FIRST 5 * FROM tab1;
```

The following collection subquery returns only the eleventh through fifteenth qualifying rows as a collection-derived table, orders these five rows by the value in column **a**, and stores this result set in a temporary table.

```
SELECT * FROM TABLE (MULTISET (SELECT SKIP 10 FIRST 5 a FROM tab3  
ORDER BY a)) INTO TEMP;
```

The following INSERT statement includes a collection subquery whose results define a collection-derived table. The rows are ordered by the value in column **a**, and are inserted into table **tab1**.

```
INSERT INTO tab1 (a) SELECT * FROM TABLE (MULTISET (SELECT SKIP 10 FIRST 5 a  
FROM tab3 ORDER BY a));
```

Queries that combine the FIRST or LIMIT and SKIP options with the ORDER BY clause can impose a unique order on the qualifying rows, so successive queries that increment the *offset* value by the value of *max* can partition the qualifying rows into disjunct subsets of *max* rows. This can support web applications that require a fixed page size, without requiring cursor management.

You can use these features in distributed queries only if all of the participating database servers support the SKIP and FIRST options.

Using the MIDDLE Option (XPS)

The MIDDLE keyword, like the FIRST, can specify a maximum number of rows to retrieve that match conditions specified in the SELECT statement. The FIRST option returns the first *max* rows (for *max* substitute a number that you specify) that satisfy the selection criteria, but the MIDDLE option returns *max* rows from the middle of the set of qualifying rows.

The syntax and restrictions for this option are the same as those for the FIRST option. For more information, see “Using the FIRST Option” on page 2-533.

Allowing Duplicates

You can apply the ALL, UNIQUE, or DISTINCT keywords to indicate whether duplicate values are returned, if any exist. If you do not specify any of these keywords in the Projection clause, all qualifying rows are returned by default.

Keyword

Effect

ALL Specifies that all qualifying rows are returned, regardless of whether duplicates exist. (This is the default specification.)

DISTINCT

Excludes duplicates of qualifying rows from the result set

UNIQUE

Excludes duplicate. (Here UNIQUE is a synonym for DISTINCT. This is an extension to the ANSI/ISO standard.)

For example, the next query returns all the unique ordered pairs of values from the **stock_num** and **manu_code** columns in rows of the **items** table. If several rows have the same pair of values, that pair appears only once in the result set:

```
SELECT DISTINCT stock_num, manu_code FROM items;
```

You can specify DISTINCT or UNIQUE no more than once in each level of a query or subquery. The following example uses DISTINCT in both the query and in the subquery:

```
SELECT DISTINCT stock_num, manu_code FROM items
  WHERE order_num = (SELECT DISTINCT order_num FROM orders
                    WHERE customer_num = 120);
```

Data Types in Distributed Queries

Queries whose only data sources are tables and views in the local database to which the session is connected can return values from columns or expressions of any built-in or user-defined data type that is registered in the local database. Queries that reference tables or views in other databases are called *distributed queries*, and the data types that they can access are a subset of the data types that Dynamic Server supports in local queries.

Among distributed queries, the restrictions on data types depend on the number of participating database servers.

- If all the databases that the query accesses are databases of the same Dynamic Server instance, the query is called a *cross-database* distributed query.
- If the query accesses databases of multiple Dynamic Server instances, the query is called a *cross-server* distributed query.

In both types of distributed queries, all participating databases must have the same ANSI/ISO-compliance status. A cross-server distributed query can use both the SKIP and FIRST options if all participating servers support the SKIP option; otherwise the query fails with an error. More generally, all cross-server operations require that the participating database server instances support the SQL syntax that specifies the operation.

For additional information about distributed queries, see the *IBM Informix Database Design and Implementation Guide*.

Data Types in Cross-Database Transactions: Distributed queries (and other distributed DML operations or function calls) that access only databases of the local Dynamic Server instance can access data types of the following categories:

- The *built-in data types* that are not opaque, including these:
 - BIGINT
 - BIGSERIAL
 - BYTE
 - CHAR
 - DATE
 - DATETIME
 - DECIMAL
 - FLOAT
 - INT
 - INTERVAL
 - INT8
 - MONEY
 - NCHAR
 - NVARCHAR
 - SERIAL
 - SERIAL8
 - SMALLFLOAT
 - SMALLINT
 - TEXT
 - VARCHAR
- Most *built-in opaque data types*, including these:
 - BLOB
 - BOOLEAN
 - CLIENTBINVAL
 - CLOB
 - IFX_LO_SPEC
 - IFX_LO_STAT
 - INDEXKEYARRAY

- LVARCHAR
- POINTER
- RTNPARAMTYPES,
- SELFUNCARGS
- STAT
- XID
- User-defined types (UDTs) that are cast explicitly to any of the built-in types that are listed above
- DISTINCT of any of the built-in types in the preceding list.

Distributed operations across databases of the local Dynamic Server instance can return UDTs and DISTINCT types based on built-in data types only if all the UDTs and DISTINCT types are cast explicitly to built-in data types.

All the opaque UDTs, DISTINCT types, data type hierarchies, and casts must have exactly the same definitions in each database that participates in the distributed query. For queries or other DML operations in cross-server UDRs that use the data types listed above as parameters or as returned data types, the UDR must also have the same definition in each participating database.

A cross-database distributed query (or any other cross-database DML operation) fails with an error if it references a table, view, or synonym in another database of the local Dynamic Server instance that includes a column of any of the following data types:

- LOLIST
- IMPEX
- IMPEXBIN
- SENDRECV
- DISTINCT of any of the opaque data types that are listed above.
- Complex types (named or unnamed ROW, COLLECTION, LIST, MULTISSET, or SET)

Data Types in Cross-Server Transactions: A distributed query (or any other distributed DML operation or function call) across databases of two or more Dynamic Server instances cannot return complex or large-object data types, nor most user-defined data types (UDTs) or opaque data types. Cross-server distributed queries, DML operations, and function calls can return only the following data types:

- Any non-opaque built-in data type
- BOOLEAN
- LVARCHAR
- DISTINCT of non-opaque built-in types
- DISTINCT of BOOLEAN
- DISTINCT of LVARCHAR
- DISTINCT of any of the DISTINCT types that appear above in this list.

A cross-server distributed query can support DISTINCT data types only if they are cast explicitly to built-in types, and all of the DISTINCT types, their data type hierarchies, and their casts are defined exactly the same way in each database that participates in the distributed query. For queries or other DML operations in

cross-server UDRs that use the data types in the preceding list as parameters or as returned data types, the UDR must also have the same definition in every participating database.

The built-in DISTINCT data type IDSSecurityLabel, which stores security label objects, can be accessed in cross-server and cross-database operations on protected data by users who hold sufficient security credentials. Like local operations on protected data, distributed queries that access remote tables protected by a security policy can return only the qualifying rows that IDSLBACRULES allow, after the database server has compared the security label that secures the data with the security credentials of the user who issues the query.

For additional information about the data types that Dynamic Server supports in cross-server DML operations, see “Data Types in Cross-Server Transactions” on page 2-537. For information about the table hierarchies of the DISTINCT data types that are valid in cross-server operations, see “DISTINCT Types in Distributed Operations” on page 4-33.

A cross-server query (or any other cross-server DML operation) fails with an error if it references a table, view, or synonym in a database of another Dynamic Server instance that includes a column of any of the following data types:

- BLOB
- CLOB
- INDEXKEYARRAY
- POINTER
- RTNPARAMTYPES
- SELFUNCARGS
- IFX_LO_SPEC
- IFX_LO_STAT
- STAT
- CLIENTBINVAL
- User-defined OPAQUE types
- Complex types (named or unnamed ROW, COLLECTION, LIST, MULTISSET, or SET)
- DISTINCT of any of the opaque or complex data types that are listed above.

Queries cannot access a database of another database server unless both servers define TCP/IP connections in DBSERVERNAME or DBSERVERALIASES configuration parameters. This applies to any communication between Dynamic Server instances, even if both database servers reside on the same computer.

Expressions in the Select List

You can use any basic type of expression (column, constant, built-in function, aggregate function, and user-defined routine), or combination thereof, in the select list. The expression types are described in “Expression” on page 4-40. Sections that follow present examples of simple expression in the select list.

You can combine simple numeric expressions by connecting them with arithmetic operators for addition, subtraction, multiplication, and division. If you combine a column expression and an aggregate function, however, you must include the column expression in the GROUP BY clause. (See also “Relationship of GROUP BY and Projection Clauses” on page 2-567.)

In general, you cannot use variables (for example, host variables in an Informix ESQL/C application) in the select list by themselves. A variable is valid in the select list, however, if an arithmetic or concatenation operator connects it to a constant.

In a FOREACH SELECT statement, you cannot use SPL variables in the select list, by themselves or with column names, when the tables in the FROM clause are remote tables. You can use SPL variables by themselves or with a constant in the select list only when the tables in the FROM clause are local tables.

In distributed queries of Dynamic Server, values in expressions (and returned by expressions) are restricted, as “Data Types in Cross-Server Transactions” on page 2-537 describes. Any UDRs whose return values are used as expressions in other databases of the same Dynamic Server instance must be defined in each participating database.

The Boolean operator NOT is not valid in the Projection clause.

Selecting Columns: Column expressions are the most commonly used expressions in a SELECT statement. For a complete description of the syntax and use of column expressions, see “Column Expressions” on page 4-59. The following examples use column expressions in the Projection clause:

```
SELECT orders.order_num, items.price FROM orders, items;
SELECT customer.customer_num ccnum, company FROM customer;
SELECT catalog_num, stock_num, cat_advert [1,15] FROM catalog;
SELECT lead_time - 2 UNITS DAY FROM manufact;
```

Selecting Constants: If you include a constant expression in the projection list, the same value is returned for each row that the query returns (except when the constant expression is NEXTVAL). For a complete description of the syntax and use of constant expressions, see “Constant Expressions” on page 4-70. Examples that follow show constant expressions within a select list:

```
SELECT 'The first name is', fname FROM customer;
SELECT TODAY FROM cust_calls;
SELECT SITENAME FROM systables WHERE tabid = ;1
SELECT lead_time - 2 UNITS DAY FROM manufact;
SELECT customer_num + LENGTH('string') from customer;
```

Selecting Built-In Function Expressions: A built-in function expression uses a function that is evaluated for each row in the query. All built-in function expressions require arguments. This set of expressions contains the time functions and the length function when they are used with a column name as an argument. The following examples show built-in function expressions within the select list of the Projection clause:

```
SELECT EXTEND(res_dtime, YEAR TO SECOND) FROM cust_calls;
SELECT LENGTH(fname) + LENGTH(lname) FROM customer;
SELECT HEX(order_num) FROM orders;
SELECT MONTH(order_date) FROM orders;
```

Selecting Aggregate Function Expressions: An aggregate function returns one value for a set of queried rows. This value depends on the set of rows that the WHERE clause of the SELECT statement qualifies. In the absence of a WHERE clause, the aggregate functions take on values that depend on all the rows that the FROM clause forms. Examples that follow show aggregate functions in a select list:

```
SELECT SUM(total_price) FROM items WHERE order_num = 1013;
SELECT COUNT(*) FROM orders WHERE order_num = 1001;
SELECT MAX(LENGTH(fname) + LENGTH(lname)) FROM customer;
```

Selecting User-Defined Function Expressions: User-defined functions extend the range of functions that are available to you and allow you to perform a subquery on each row that you select.

The following example calls the **get_orders()** user-defined function for each **customer_num** and displays the returned value under the **n_orders** label:

```
SELECT customer_num, lname, get_orders(customer_num) n_orders
      FROM customer;
```

If an SPL routine in a SELECT statement contains certain SQL statements, the database server returns an error. For information on which SQL statements cannot be used in an SPL routine that is called within a query, see “Restrictions on SPL Routines in Data-Manipulation Statements” on page 5-78.

For the complete syntax of user-defined function expressions, see “User-Defined Functions” on page 4-161.

Selecting Expressions That Use Arithmetic Operators: You can combine numeric expressions with arithmetic operators to make complex expressions. You cannot combine expressions that contain aggregate functions with column expressions. These examples show expressions that use arithmetic operators within a select list in the Projection clause:

```
SELECT stock_num, quantity*total_price FROM customer;
SELECT price*2 doubleprice FROM items;
SELECT count(*)+2 FROM customer;
SELECT count(*)+LENGTH('ab') FROM customer;
```

Selecting ROW Fields (IDS): You can select a specific field of a named or unnamed ROW type column with *row.field* notation, using a period (.) as a separator between the *row* and *field* names. For example, suppose you have the following table structure:

```
CREATE ROW TYPE one (a INTEGER, b FLOAT);
CREATE ROW TYPE two (c one, d CHAR(10));
CREATE ROW TYPE three (e CHAR(10), f two);

CREATE TABLE new_tab OF TYPE two;
CREATE TABLE three_tab OF TYPE three;
```

The following examples show expressions that are valid in the select list:

```
SELECT t.c FROM new_tab t;
SELECT f.c.a FROM three_tab;
SELECT f.d FROM three_tab;
```

You can also enter an asterisk (*) in place of a field name to signify that all fields of the ROW-type column are to be selected.

For example, if the **my_tab** table has a ROW-type column named **rowcol** that contains four fields, the following SELECT statement retrieves all four fields of the **rowcol** column:

```
SELECT rowcol.* FROM my_tab;
```

You can also retrieve all fields from a row-type column by specifying only the column name. This example has the same effect as the previous query:

```
SELECT rowcol FROM my_tab;
```

You can use *row.field* notation not only with ROW-type columns but with expressions that evaluate to ROW-type values. For more information, see “Column Expressions” on page 4-59 in the Expression segment.

Declaring a Display Label

You can declare a display label for any column or column expression in the select list of the Projection clause. This temporary name is in scope only while the SELECT statement is executing.

In DB–Access, a display label appears as the heading for that column in the output of the SELECT statement.

In Informix ESQL/C, the value of *display_label* is stored in the **sqlname** field of the **sqlda** structure. For more information on the **sqlda** structure, see the *IBM Informix ESQL/C Programmer’s Manual*.

If your display label is an SQL keyword, use the AS keyword to clarify the syntax. For example, to use UNITS, YEAR, MONTH, DAY, HOUR, MINUTE, SECOND, or FRACTION as display labels, use the AS keyword with the display label. The following statement uses AS with **minute** as a display label:

```
SELECT call_dtime AS minute FROM cust_calls;
```

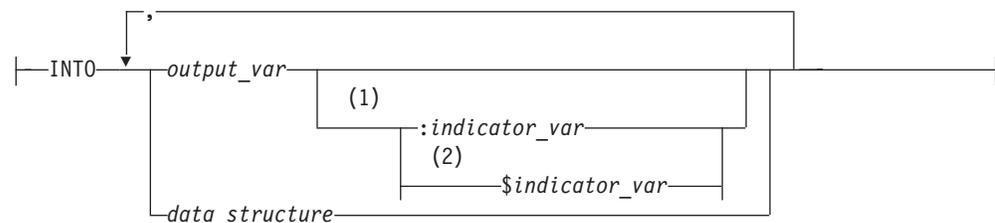
For the keywords of SQL, see Appendix A, “Reserved Words for IBM Informix Dynamic Server,” on page A-1, or Appendix B, “Reserved Words for IBM Informix Extended Parallel Server,” on page B-1.

If you are creating a temporary table, you must supply a display label for any columns that are not simple column expressions. The display label is used as the name of the column in the temporary table. If you are using the SELECT statement to define a view, do not use display labels. Specify the desired label names in the CREATE VIEW column list instead.

INTO Clause

Use the INTO clause in an SPL routine or an Informix ESQL/C program to specify the program variables or host variables to receive data that SELECT retrieves.

INTO Clause:



Notes:

- 1 ESQL/C only
- 2 Informix extension

Element	Description	Restrictions	Syntax
<i>data_structure</i>	Structure that was declared as a host variable	Data types of elements must be able to store the values that are being selected	Language specific
<i>indicator_var</i>	Program variable to receive a return code if corresponding <i>output_var</i> receives a NULL value	Optional; use an indicator variable if the possibility exists that the value of the corresponding <i>output_var</i> is NULL	Language specific
<i>output_var</i>	Program or host variable to receive value of the corresponding select list item. Can be a collection variable	Order of receiving variables must match the order of corresponding items in the select list of Projection clause	Language specific

The INTO clause specifies one or more variables that receive the values that the query returns. If it returns multiple values, they are assigned to the list of variables in the order in which you specify the variables.

If the SELECT statement stands alone (that is, it is not part of a DECLARE statement and does not use the INTO clause), it must be a singleton SELECT statement. A *singleton* SELECT statement returns only one row.

The number of receiving variables must be equal to the number of items in the select list of the Projection clause. The data type of each receiving variable should be compatible with the data type of the corresponding column or expression in the select list.

For the actions that the database server takes when the data type of the receiving variable does not match that of the selected item, see “Warnings in ESQL/C” on page 2-544.

The following example shows a singleton SELECT statement in Informix ESQL/C:

```
EXEC SQL select fname, lname, company_name
      into :p_fname, :p_lname, :p_coname
      where customer_num = 101;
```

In an SPL routine, if a SELECT returns more than one row, you must use the FOREACH statement to access the rows individually. The INTO clause of the SELECT statement holds the fetched values. For more information, see “FOREACH” on page 3-22.

INTO Clause with Indicator Variables

If the possibility exists that a data value returned from the query is NULL, use an ESQL/C indicator variable in the INTO clause. For more information, see the *IBM Informix ESQL/C Programmer's Manual*.

INTO Clause with Cursors

If the SELECT statement returns more than one row, you must use a cursor in a FETCH statement to fetch the rows individually. You can put the INTO clause in the FETCH statement rather than in the SELECT statement, but you should not put it in both.

The following Informix ESQL/C code examples show different ways you can use the INTO clause. As both examples show, first you must use the DECLARE statement to declare a cursor.

Using the INTO clause in the SELECT statement:

```
EXEC SQL declare q_curs cursor for
    select lname, company
        into :p_lname, :p_company
        from customer;
EXEC SQL open q_curs;
while (SQLCODE == 0)
    EXEC SQL fetch q_curs;
EXEC SQL close q_curs;
```

Using the INTO clause in the FETCH statement:

```
EXEC SQL declare q_curs cursor for
    select lname, company from customer;
EXEC SQL open q_curs;
while (SQLCODE == 0)
    EXEC SQL fetch q_curs into :p_lname, :p_company;
EXEC SQL close q_curs;
```

Preparing a SELECT ... INTO Query

In Informix ESQL/C, you cannot prepare a query that has an INTO clause. You can prepare the query without the INTO clause, declare a cursor for the prepared query, open the cursor, and then use the FETCH statement with an INTO clause to fetch the cursor into the program variable.

Alternatively, you can declare a cursor for the query without first preparing the query and include the INTO clause in the query when you declare the cursor. Then open the cursor and fetch the cursor without using the INTO clause of the FETCH statement.

Using Array Variables with the INTO Clause

In Informix ESQL/C, if you use a DECLARE statement with a SELECT statement that contains an INTO clause, and the variable is an array element, you can identify individual elements of the array with integer literals or variables. The value of the variable that is used as a subscript is determined when the cursor is declared; the subscript variable subsequently acts as a constant.

The following Informix ESQL/C code example declares a cursor for a SELECT ... INTO statement using the variables *i* and *j* as subscripts for the array *a*. After you declare the cursor, the INTO clause of the SELECT statement is equivalent to INTO *a*[5], *a*[2].

```
i = 5
j = 2
EXEC SQL declare c cursor for
    select order_num, po_num into :a[i], :a[j] from orders
        where order_num =1005 and po_num =2865;
```

You can also use program variables in the FETCH statement to specify an element of a program array in the INTO clause. The program variables are evaluated at each fetch, rather than when you declare the cursor.

Error Checking

If the data type of the receiving variable does not match that of the selected item, the data type of the selected item is converted, if possible, to the data type of the variable. If the conversion is impossible, an error occurs, and a negative value is returned in the **status** variable, **sqlca.sqlcode**, or **SQLCODE**. In this case, the value in the program variable is unpredictable.

In an ANSI-compliant database, if the number of variables that are listed in the INTO clause differs from the number of items in the select list of the Projection clause, you receive an error.

Warnings in ESQL/C: In Informix ESQL/C, if the number of variables listed in the INTO clause differs from the number of items in the Projection clause, a warning is returned in the `sqlwarn` structure: `sqlca.sqlwarn.sqlwarn3`. The actual number of variables that are transferred is the lesser of the two numbers. For information about the `sqlwarn` structure, see the *IBM Informix ESQL/C Programmer's Manual*.

FROM Clause

The FROM clause lists the table objects from which you are selecting the data.

FROM Clause:

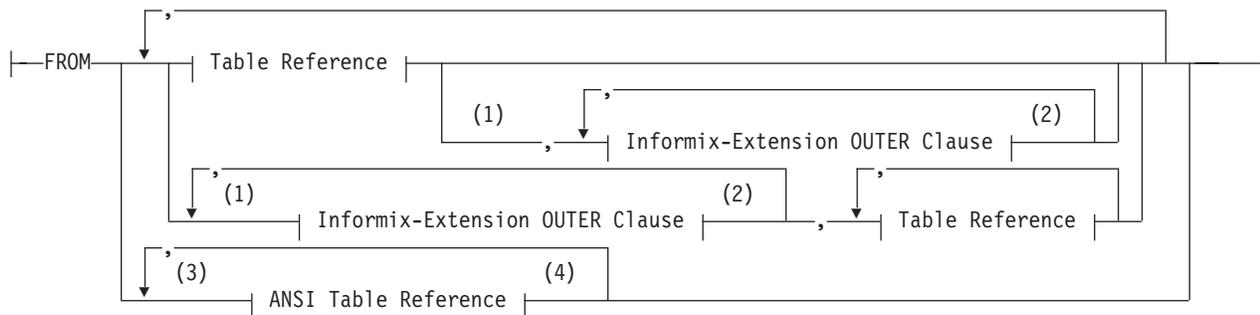
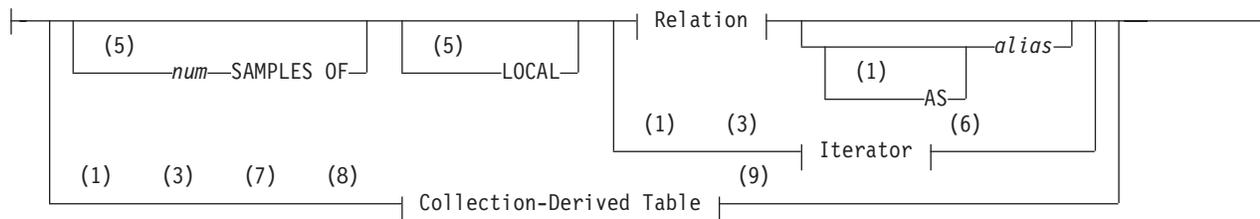
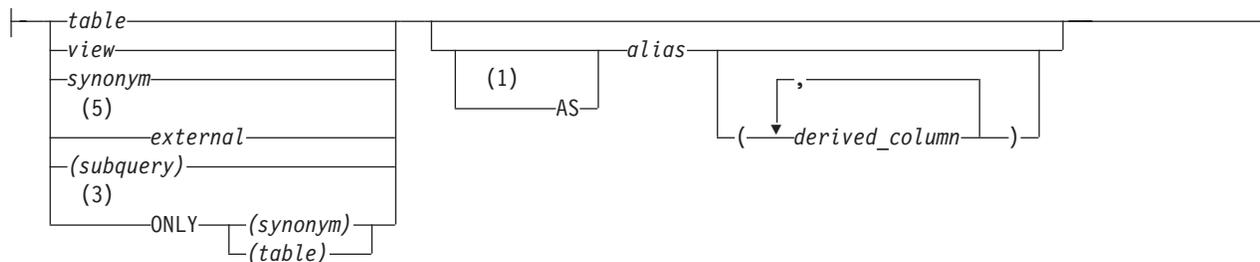


Table Reference:



Relation:



Notes:

- 1 Informix extension
- 2 See "Informix-Extension Outer Joins" on page 2-559
- 3 Dynamic Server only
- 4 See "ANSI Table Reference" on page 2-554
- 5 Extended Parallel Server only

- 6 See “Iterator Functions (IDS)” on page 2-551
- 7 Stored Procedure Language only
- 8 ESQL/C only
- 9 See “Collection-Derived Table” on page 5-4

Element	Description	Restrictions	Syntax
<i>alias</i>	Temporary name for a table, view, or derived table in this query	See “The AS Keyword.”	“Identifier ” on page 5-21
<i>derived_column</i>	Temporary name for a derived column in a table expression	Unless the underlying collection is a ROW type, you can declare no more than one <i>derived_column</i> name	“Identifier ” on page 5-21
<i>external</i>	External table from which to retrieve data	Must exist but cannot be the outer table in an outer join	“Database Object Name ” on page 5-16
<i>num</i>	Number of rows to sample	Unsigned integer > 0	“Literal Number ” on page 4-185
<i>subquery</i>	Specifies rows to be retrieved	Cannot be a correlated subquery	“SELECT” on page 2-528
<i>synonym, table, view</i>	Synonym for a table from which to retrieve data	Synonym and table or view to which it points must exist	“Database Object Name ” on page 5-16

If the FROM clause specifies more than one data source, the query is called a *join*, because its result set can join rows from several table references. For more information about joins, see “Queries that Join Tables” on page 2-552.

Aliases for Tables or Views

You can declare an alias for a table or view in the FROM clause. If you do so, you must use the alias to refer to the table or view in other clauses of the SELECT statement. You can also use aliases to make the query shorter.

The following examples show typical uses of the FROM clause. The first query selects all the columns and rows from the **customer** table. The second query uses a join between the **customer** and **orders** table to select all the customers who have placed orders.

```
SELECT * FROM customer;
SELECT fname, lname, order_num FROM customer, orders
WHERE customer.customer_num = orders.customer_num;
```

The next example is equivalent to the second query in the preceding example, but it declares aliases in the FROM clause and uses them in the WHERE clause:

```
SELECT fname, lname, order_num FROM customer c, orders o
WHERE c.customer_num = o.customer_num;
```

Aliases (sometimes called *correlation names*) are especially useful with a self-join. For more information about self-joins, see “Self-Joins” on page 2-565. In a self-join, you must list the table name twice in the FROM clause and declare a different alias for each of the two instances of the table name.

The AS Keyword: If you use a potentially ambiguous word as an alias (or as a display label), you must begin its declaration with the keyword AS. This keyword is required if you use any of the keywords ORDER, FOR, AT, GROUP, HAVING, INTO, NOT, UNION, WHERE, WITH, CREATE, or GRANT as an alias for a table or view.

The database server would issue an error if the next example did not include the AS keyword to indicate that **not** is a display label, rather than an operator:

```
CREATE TABLE t1(a INT);  
SELECT a AS not FROM t1;
```

If you do not declare an alias for a collection-derived table, the database server assigns an implementation-dependent name to it.

Table Expressions

A *table expression* (sometimes called a *derived table*) is a view name, a table name, or an uncorrelated subquery in the FROM clause. Table expressions can be simple or complex:

- Simple table expressions

A *simple* table expression is one whose underlying query can be folded into the main query while preserving the correctness of the query result.

- Complex table expressions

A *complex* table expression is one whose underlying query cannot be folded into the main query while preserving the correctness of the query result. The database server materializes such table expressions into a temporary table that is used in the main query. Subqueries in the FROM clause that specify aggregates, unions, or the ORDER BY clause are implemented as complex table expressions, which typically require more resources of the database server than simple table expressions.

In either case, the table expression is evaluated as a general SQL query and its results can be thought of as a logical table. This logical table and its columns can be used just like an ordinary base table, but it is not persistent. It exists only during the execution of the query that references it.

Table expressions have the same syntax as general SELECT statements, but with most of the restrictions that apply to subqueries in other contexts. A table expression cannot include the SELECT INTO clause.

Unlike subqueries in the Projection clause or in the WHERE clause, subqueries in the FROM clause cannot be correlated subqueries. A *correlated subquery* is a subquery that refers to a column of a table that is not listed in its FROM clause. Conversely, any subquery that references only columns in tables that are listed in its FROM clause is an *uncorrelated subquery*.

Although a derived table in the FROM clause cannot be defined by a correlated subquery, the uncorrelated subquery that defines a derived table can contain a correlated subquery outside its FROM clause, as in the following example:

```
SELECT * FROM (SELECT * FROM t1  
WHERE a IN (SELECT b FROM t2 WHERE t1.a = t2.b));
```

Here the subquery in the first WHERE clause is a correlated subquery, because it references column **a** of table **t1**, but its FROM clause specifies only table **t2**. No error is returned, because the correlated subquery is specified in the WHERE clause of the subquery in which it is nested, not the FROM clause.

In FROM clause table expressions, Dynamic Server also supports the ORDER BY clause, which is not valid in subqueries outside the FROM clause. Columns or expressions that are specified by the ORDER BY clause in a table expression need not be included in the Projection clause.

In Extended Parallel Server, table expressions in the FROM clause have these additional restrictions:

- Table expressions are not valid in CREATE GK INDEX statements
- Table expressions are not valid in the triggered actions of CREATE TRIGGER statements
- Table expressions cannot include the ORDER BY clause

Dynamic Server does not support GK indexes, but is not subject to the last two restrictions. It supports table expressions in the triggered actions of CREATE TRIGGER statements, and as the triggering event of a Select trigger. Dynamic Server also supports the ORDER BY clause in table expressions, and iterator functions as FROM clause table expressions.

The CALL statement of SPL cannot invoke a TABLE function within a subquery in the FROM clause of the SELECT statement.

Apart from these restrictions, any valid SQL query can be a table expression. A table expression can be nested within another table expression (within the 64 kilobyte size limit on the text of a single SQL statement), and can include tables and views in its definition. You can use table expressions in CREATE VIEW statements to define views.

Usability and Performance Considerations: Although equivalent functionality is available through views, subqueries as table expressions simplify the formulation of queries, make the syntax more flexible and intuitive, and support the ANSI/ISO standard for SQL.

The query optimizer does not materialize simple table expressions that the FROM clause specifies. The performance of a query that uses the ANSI/ISO syntax for a table expression in the FROM clause is at least as good as that of a query that uses the Informix-extension TABLE (MULTISET (SELECT ...)) syntax to specify an equivalent derived table in the FROM clause. Subqueries in the FROM clause that include aggregate functions, the UNION operator, or ORDER BY specifications are implemented as complex table expressions that can impose greater costs than simple table expressions. Use the SET EXPLAIN statement to examine the query plan and the estimated cost of a table expression.

The following are examples of valid table expressions:

```
SELECT * FROM (SELECT * FROM t);

SELECT * FROM (SELECT * FROM t) AS s;

SELECT * FROM (SELECT * FROM t) AS s WHERE t.a = s.b;

SELECT * FROM (SELECT * FROM t) AS s, (SELECT * FROM u) AS v WHERE s.a = v.b;

SELECT * FROM (SELECT SKIP 2 col1 FROM tab1 WHERE col1 > 50 ORDER BY col1 DESC);

SELECT * FROM (SELECT col1,col3 FROM tab1
WHERE col1 < 50 GROUP BY col1,col3 ORDER BY col3 ) vtab(vcol0,vcol1);

SELECT * FROM (SELECT * FROM t WHERE t.a = 1) AS s,
OUTER
(SELECT * FROM u WHERE u.b = 2 GROUP BY 1) AS v WHERE s.a = v.b;

SELECT * FROM (SELECT a AS colA FROM t WHERE t.a = 1) AS s,
OUTER
(SELECT b AS colB FROM u WHERE u.b = 2 GROUP BY 1) AS v
```

```

WHERE s.colA = v.colB;

CREATE VIEW vu AS SELECT * FROM (SELECT * FROM t);

SELECT * FROM ((SELECT * FROM t) AS r) AS s;

```

Restrictions on External Tables in Joins and Subqueries

In Extended Parallel Server, when you use external tables in joins or subqueries, the following restrictions apply:

- No more than one external table is valid in a query.
- The external table cannot be the outer table in an outer join.
- For subqueries that cannot be converted to joins, you can use an external table in the main query, but not the subquery.
- You cannot do a self-join on an external table.

For more information on subqueries, see your *IBM Informix Performance Guide*.

Application Partitioning: The LOCAL Keyword (XPS)

In Extended Parallel Server, the LOCAL table feature allows client applications to read data only from the *local* fragments of a table. In other words, it allows the application to read only the fragments that reside on the coserver to which the client is connected. This feature supports application partitioning. An application can connect to multiple coservers, execute a LOCAL read on each coserver, and assemble the final result on the client computer.

You qualify the name of a table with the LOCAL keyword to indicate that you want to retrieve rows from fragments only on the local coserver. The LOCAL keyword has no effect on data retrieved from nonfragmented tables.

When a query involves a join, you must plan carefully if you want to extract data that the client can aggregate. The simplest way to ensure that a join will retrieve data suitable for aggregation is to limit the number of LOCAL tables to one. The client can then aggregate data with respect to that table. The following example shows a query that returns data suitable for aggregation by the client:

```

SELECT x.col1, y.col2 FROM LOCAL tab1 x, tab2 y      -- tab1 is local
      INTO TEMP t1 WHERE x.col1 = y.col1;          -- can be aggregated by client

```

The following example shows data that the client cannot aggregate:

```

SELECT x.col1, y.col2 FROM LOCAL tab1 x, LOCAL tab2 -- tab1 and tab2 are local
      INTO SCRATCH s4 WHERE x.col1 = y.col1;        -- client cannot aggregate

```

The client must submit exactly the same query to each coserver to retrieve data that can be aggregated.

Sampled Queries: The SAMPLES OF Keywords (XPS)

In Extended Parallel Server, *sampled queries* are supported. Sampled queries are queries that are based on *sampled tables*. A sampled table is the result of randomly selecting a specified number of rows from the table, rather than all rows that match the selection criteria.

You can use a sampled query to gather quickly an approximate profile of data within a large table. If you use a sufficiently large sample size, you can examine trends in the data by sampling the data instead of scanning all the data. In such cases, sampled queries can provide better performance than scanning the data. If the value specified for the sample size is greater than the number of rows in the table, the whole table is scanned.

To indicate that a table is to be sampled, specify the number of samples to return before the SAMPLES OF keywords of the FROM clause. You can run sampled queries against tables and synonyms, but not against views.

Sampled queries are not supported in the INSERT, DELETE, or UPDATE statements, nor in other SQL statements that can include a query.

A sampled query has at least one sampled table. You do not need to sample all tables in a sampled query. You can specify the SAMPLES OF option for some tables in the FROM clause but not specify it for other tables.

The sampling method is known as *sampling without replacement*. This means that a sampled row is not sampled again. The database server applies selection criteria *after* samples are selected. Thus, the selection criteria restrict the sample set, rather than the rows from which the sample is taken.

If a table is fragmented, the database server divides the specified number of samples among the fragments. The number of samples from a fragment is proportional to the ratio of the size of a fragment to the size of the table. In other words, the database server takes more samples from larger fragments.

Important: You must run UPDATE STATISTICS MEDIUM before you run the query with the SAMPLES OF option. If you do not run UPDATE STATISTICS, the SAMPLES OF clause is ignored, and all data values are returned. For better results, run UPDATE STATISTICS MEDIUM before you run the query with the SAMPLES OF option.

The results of a sampled query contain a certain amount of deviation from a complete scan of all rows. You can reduce this expected error to an acceptable level by increasing the proportion of sampled rows to actual rows. When you use sampled queries in joins, the expected error increases dramatically; you must use larger samples in each table to retain an acceptable level of accuracy.

For example, you might want to generate a list of how many of each part is sold from the **parts_sold** table, which is known to contain approximately 100,000,000 rows.

The following query provides a sampling ratio of one percent and returns an approximate result:

```
SELECT part_number, COUNT(*) * 100 AS how_many
FROM 1000000 SAMPLES OF parts_sold
GROUP BY part_number;
```

The ONLY Keyword (IDS)

If the SELECT statement queries a supertable, rows from both the supertable and its subtables are returned. To query rows from the supertable only, you must include the ONLY keyword in the FROM clause, as in this example:

```
SELECT * FROM ONLY(super_tab);
```

Selecting from a Collection Variable (IDS)

The SELECT statement in conjunction with the Collection-Derived Table segment allows you to select elements from a collection variable.

The Collection-Derived Table segment identifies the collection variable from which to select the elements. (See “Collection-Derived Table” on page 5-4.)

Using Collection Variables with SELECT: To modify the contents of a column of a collection data type, you can use the SELECT statement with a collection variable in various ways:

- You can select the contents (if any) of a collection column into a collection variable.

You can assign the data type of the column to a collection variable of type COLLECTION (that is, an untyped collection variable).

- You can select the contents from a collection variable to determine the data that you might want to update.
- You can select the contents from a collection variable INTO another variable in order to update certain collection elements.

The INTO clause identifies the variable for the element value that is selected from the collection variable. The data type of the host variable in the INTO clause must be compatible with that of the corresponding collection element.

- You can use a Collection cursor to select one or more elements from an Informix ESQL/C collection variable.

For more information, including restrictions on the SELECT statement, see “Associating a Cursor with a Prepared Statement” on page 2-297.

- You can use a Collection cursor to select one or more elements from an SPL collection variable.

For more information, including restrictions on the SELECT statement, see “Using a SELECT ... INTO Statement” on page 3-24.

When one of the tables to be joined is a collection, the FROM clause cannot specify a join. This restriction applies when the collection variable holds your collection-derived table. See also “Collection-Derived Table” on page 5-4. and the INSERT, UPDATE, and DELETE statement descriptions in this chapter.

Selecting from a Row Variable (IDS, ESQL/C)

The SELECT statement can include the Collection-Derived Table segment to select one or more fields from a **row** variable. The Collection-Derived Table segment identifies the **row** variable from which to select the fields. For more information, see “Collection-Derived Table” on page 5-4.

To select fields

1. Create a **row** variable in your Informix ESQL/C program.
2. Optionally, fill the **row** variable with field values. You can select a ROW-type column into the **row** variable with the SELECT statement (without the Collection-Derived Table segment). Alternatively, you can insert field values into the **row** variable with the UPDATE statement and the Collection-Derived Table segment.
3. Select row fields from the **row** variable with the SELECT statement and the Collection-Derived Table segment.
4. Once the **row** variable contains the correct field values, you can use the INSERT or UPDATE statement on a table or view name to save the contents of the **row** variable in a named or unnamed row column.

The INTO clause can specify a host variable to hold a field value selected from the **row** variable.

The type of the host variable must be compatible with that of the field. For example, this code fragment puts the **width** field value into the **rect_width** host variable.

```
EXEC SQL BEGIN DECLARE SECTION;
    ROW (x INT, y INT, length FLOAT, width FLOAT) myrect;
    double rect_width;
EXEC SQL END DECLARE SECTION;
...
EXEC SQL SELECT rect INTO :myrect FROM rectangles
    WHERE area = 200;
EXEC SQL SELECT width INTO :rect_width FROM table(:myrect);
```

The **SELECT** statement on a **row** variable has the following restrictions:

- No expressions are allowed in the select list of the Projection clause.
- ROW columns cannot be in a WHERE clause comparison condition.
- The Projection clause must be an asterisk (*) if the row-type contains fields of opaque, distinct, or built-in data types.
- Columns listed in the Projection clause can have only unqualified names. They cannot use the syntax *database@server:table.column*.
- The following clauses are not allowed: GROUP BY, HAVING, INTO TEMP, ORDER BY, and WHERE.
- The FROM clause has no provisions to do a join.

You can modify the **row** variable with the Collection-Derived Table segment of the UPDATE statements. (The INSERT and DELETE statements do not support a **row** variable in the Collection-Derived Table segment.)

The **row** variable stores the fields of the row. It has no intrinsic connection, however, with a database column. Once the **row** variable contains the correct field values, you must then save the variable into the ROW column with one of the following SQL statements:

- To update the ROW column in the table with the **row** variable, use an UPDATE statement on a table or view name and specify the **row** variable in the SET clause. For more information, see “Updating ROW-Type Columns (IDS)” on page 2-696.
- To insert a row into a ROW column, use the INSERT statement on a table or view and specify the **row** variable in the VALUES clause. See “Inserting Values into ROW-Type Columns (IDS)” on page 2-437.

For examples of how to use SPL row variables, see the *IBM Informix Guide to SQL: Tutorial*. For information on using Informix ESQL/C **row** variables, see the discussion of complex data types in the *IBM Informix ESQL/C Programmer's Manual*.

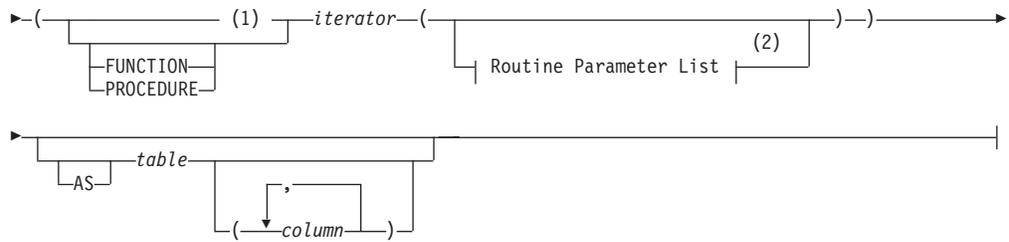
Iterator Functions (IDS)

The FROM clause can include a call to an iterator function to specify the source for a query. An *iterator function* is a user-defined function that returns to its calling SQL statement several times, each time returning a value.

You can query the returned result set of an iterator UDR using a virtual table interface. Use this syntax to invoke an iterator function in the FROM clause:

Iterator:

|—TABLE—————→



Notes:

- 1 Informix extension
- 2 See "Routine Parameter List" on page 5-67

Element	Description	Restrictions	Syntax
<i>column</i>	Name declared here for a virtual column in <i>table</i>	Must be unique among <i>column</i> names in <i>table</i> . Cannot include qualifiers	"Identifier " on page 5-21
<i>iterator</i>	Name of the iterator function	Must be registered in the database	"Identifier " on page 5-21
<i>table</i>	Name declared here for virtual table holding <i>iterator</i> result set	Cannot include qualifiers	"Identifier " on page 5-21

The keyword FUNCTION (or PROCEDURE) was required in releases earlier than Dynamic Server 10.5. These keyword extensions to the ANSI/ISO standard for SQL are optional in this release, and have no effect. The following two query specifications, which specify **fibGen()** as an iterator function, are equivalent:

```
SELECT * FROM TABLE FUNCTION ( fibGen(10));
SELECT * FROM TABLE ( fibGen(10));
```

The *table* can only be referenced within the context of this query. After the SELECT statement terminates, the virtual table no longer exists.

The number of columns must match the number of values returned by the iterator. An external function can return no more than one value (but that can be of a collection data type). An SPL routine can return multiple values.

To reference the virtual *table* columns in other parts of the SELECT statement, for example, in the WHERE clause or HAVING clause, you must declare its name and the virtual column names in the FROM clause. You do not need to declare the *table* name or *column* names in the FROM clause if you use the asterisk notation in the Select list of the Projection clause:

```
SELECT * FROM ...
```

For more information and examples of using iterator functions in queries, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

Queries that Join Tables

If the FROM clause specifies more than one table reference, the query can join rows from several tables. A *join condition* specifies a relationship between at least one column from each table to be joined. Because the columns in a join condition are being compared, they must have compatible data types.

The FROM clause of the SELECT statement can specify several types of joins.

FROM Clause Keywords	Corresponding Result Set
CROSS JOIN	Cartesian product (all possible pairs of rows)
INNER JOIN	Only rows from CROSS that satisfy the join condition
LEFT OUTER JOIN	Qualifying rows of one table, and all rows of another
RIGHT OUTER JOIN	Same as LEFT, but roles of the two tables are reversed
FULL OUTER JOIN	The union of all rows from an INNER join of the two tables, and of all rows of each table that have no match in the other table (using NULL values in the selected columns of the other table)

The last four categories are collectively called “Join Types” in the literature of the relational model; a CROSS JOIN ignores the specific data values in joined tables, returning every possible pair of rows, where one row is from each table.

In an inner (or simple) join, the result contains only the combination of rows that satisfy the join conditions. Outer joins preserve rows that otherwise would be discarded by inner joins. In an outer join, the result contains the combination of rows that satisfy the join conditions *and* the rows from the dominant table that would otherwise be discarded. The rows from the dominant table that do not have matching rows in the subordinate table contain NULL values in the columns selected from the subordinate table.

Dynamic Server supports the two different syntaxes for left outer joins:

- Informix-extension syntax
- ANSI-compliant syntax

Earlier versions of the database server supported only Informix-extension syntax for outer joins. Dynamic Server continues to support this legacy syntax, but using the ANSI-compliant syntax for joins provides greater flexibility in creating queries. In view definitions, however, the legacy syntax does not require materialized views, so it might offer performance advantages.

If you use ANSI-compliant syntax to specify a join in the FROM clause, you must also use ANSI-compliant syntax for all outer joins in the same query block. Thus, you cannot begin another outer join with only the OUTER keyword. The following query, for example, is not valid:

```
SELECT * FROM customer, OUTER orders RIGHT JOIN cust_calls
  ON (customer.customer_num = orders.customer_num)
  WHERE customer.customer_num = 104);
```

This returns an error because it attempts to combine the Informix-extension OUTER syntax with the ANSI-compliant RIGHT JOIN syntax for outer joins.

See the section “Informix-Extension Outer Joins” on page 2-559 for the Informix-extension syntax for LEFT OUTER joins.

ANSI-Compliant Joins (IDS)

The ANSI-compliant syntax for joins supports these join specifications:

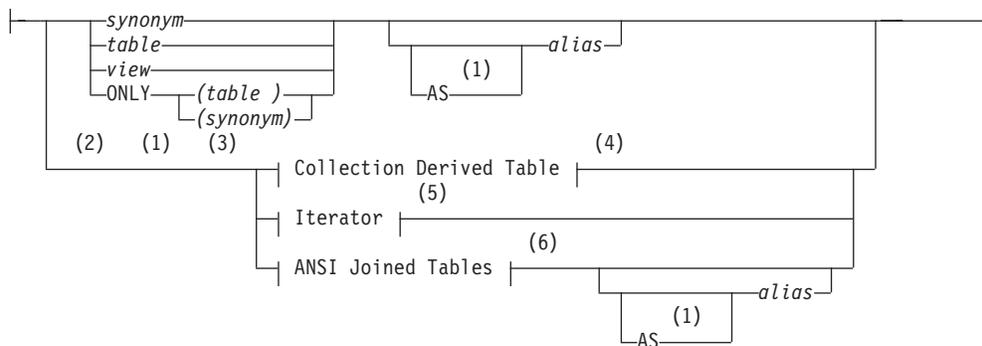
- To use a CROSS join, a LEFT OUTER, RIGHT OUTER, or FULL OUTER join, or an INNER (or simple) join, see “Creating an ANSI Join” on page 2-555 and “ANSI INNER Joins” on page 2-557.
- To use pre-join filters, see “Using the ON Clause” on page 2-557.
- To use one or more post-join filters in the WHERE clause, see “Specifying a Post-Join Filter” on page 2-558.
- To have the dominant or subordinate part of an outer join be the result set of another join, see “Using a Join as the Dominant or Subordinate Part of an Outer Join” on page 2-559.

Important: Use the ANSI-compliant syntax for joins when you create new queries in Dynamic Server. In Dynamic Server releases earlier than IDS 10.00.xC3, cross-server distributed queries with ANSI-compliant joins use query plans that are inefficient for this release. For any UDR older than IDS 10.00.xC3 that performs a cross-server ANSI-compliant join, use the UPDATE STATISTICS statement to replace the original query plan with a reoptimized plan.

ANSI Table Reference:

This diagram shows the ANSI-compliant syntax for a table reference.

ANSI Table Reference:



Notes:

- 1 Informix extension
- 2 Stored Procedure Language only
- 3 ESQL/C only
- 4 See “Collection-Derived Table” on page 5-4
- 5 See “Iterator Functions (IDS)” on page 2-551
- 6 See “ANSI Joined Tables” on page 2-555

Element	Description	Restrictions	Syntax
<i>alias</i>	Temporary name for a table or view within the scope of the SELECT	See “The AS Keyword” on page 2-545.	“Identifier ” on page 5-21
<i>synonym, table, view</i>	Source from which to retrieve data	The synonym and the table or view to which it points must exist	“Database Object Name ” on page 5-16

Here the ONLY keyword is the same semantically as in the Informix-extension Table Reference segment, as described in "The ONLY Keyword (IDS)" on page 2-549.

The AS keyword is optional when you declare an alias (also called a *correlation name*) for a table reference, as described in "The AS Keyword" on page 2-545, unless the alias conflicts with an SQL keyword.

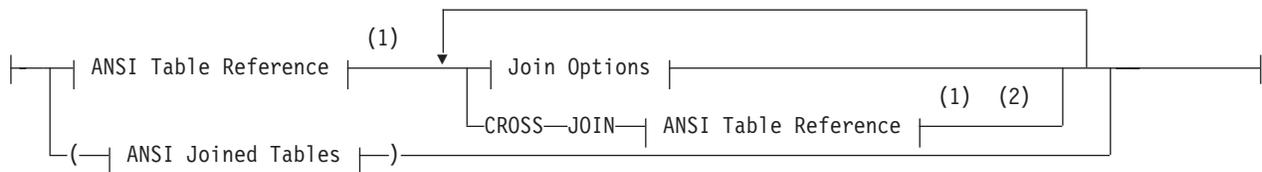
Creating an ANSI Join: With ANSI-compliant joined table syntax, as shown in the following diagram, you can specify the INNER JOIN, CROSS JOIN, NATURAL JOIN, LEFT JOIN (or LEFT OUTER JOIN), RIGHT JOIN and FULL OUTER JOIN keywords. The OUTER keyword is optional in ANSI-compliant outer joins.

You must use the same form of join syntax (either Informix extension or ANSI-compliant) for all of the outer joins in the same query block. When you use the ANSI-compliant syntax, you must also specify the join condition in the ON clause, as described in "Using the ON Clause" on page 2-557.

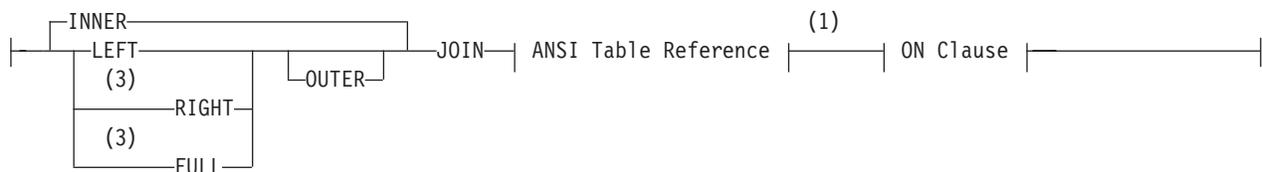
ANSI Joined Tables:

This is the ANSI-compliant syntax for specifying inner and outer joins.

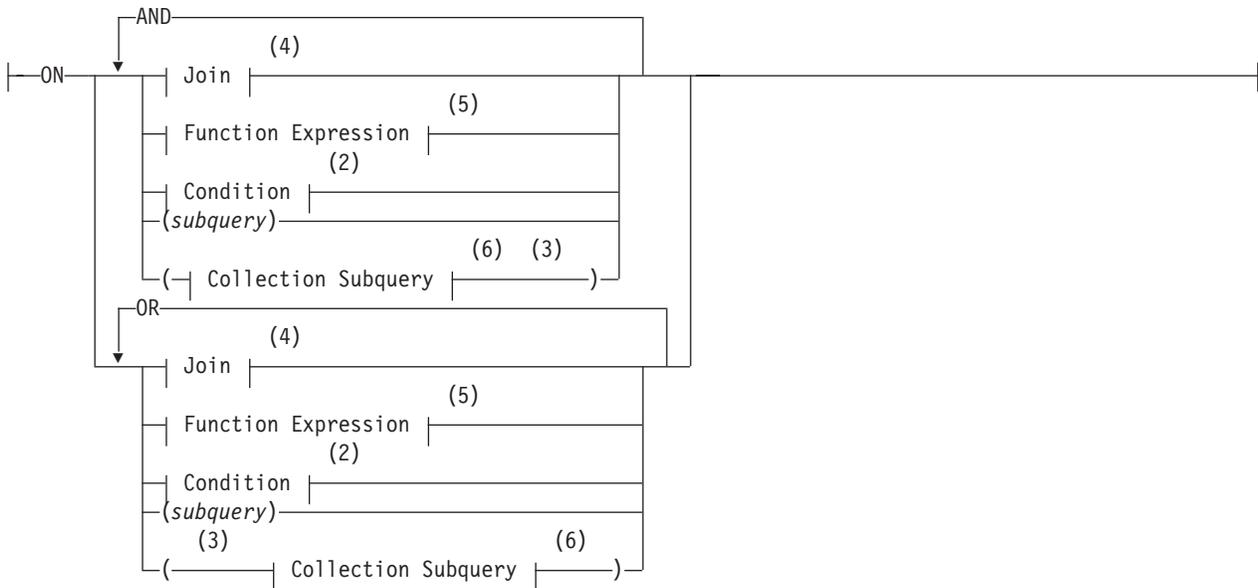
ANSI Joined Tables:



Join Options:



ON Clause:



Notes:

- 1 See "ANSI Table Reference" on page 2-554
- 2 See "Condition " on page 4-5
- 3 Dynamic Server only
- 4 See "Specifying a Join in the WHERE Clause" on page 2-564
- 5 See "Function Expressions" on page 4-87
- 6 See "Collection Subquery " on page 4-3

Element	Description	Restrictions	Syntax
<i>subquery</i>	Embedded query	Cannot contain the FIRST or the ORDER BY clause	"SELECT" on page 2-528

The ANSI-Joined Table segment must be enclosed between parentheses if it is immediately followed by another join specification. For example, the first of the following two queries returns an error; the second query is valid:

```
SELECT * FROM (T1 LEFT JOIN T2) CROSS JOIN T3 ON (T1.c1 = T2.c5)
WHERE (T1.c1 < 100); -- Ambiguous order of operations;
```

```
SELECT * FROM (T1 LEFT JOIN T2 ON (T1.c1 = T2.c5)) CROSS JOIN T3
WHERE (T1.c1 < 100); -- Unambiguous order of operations;
```

The following valid query specifies nested LEFT OUTER joins of table expressions within the FROM clause of the outer SELECT statement:

```
SELECT * FROM
( (SELECT C1,C2 FROM T3) AS VT3(V31,V32)
LEFT OUTER JOIN
( (SELECT C1,C2 FROM T1) AS VT1(VC1,VC2)
LEFT OUTER JOIN
(SELECT C1,C2 FROM T2) AS VT2(VC3,VC4)
ON VT1.VC1 = VT2.VC3)
ON VT3.V31 = VT2.VC3);
```

ANSI CROSS Joins: The CROSS keyword specifies the Cartesian product, returning all possible paired combinations that include one row from each of the joined tables.

ANSI INNER Joins: To create an inner (or simple) join using the ANSI-compliant syntax, specify the join with the JOIN or INNER JOIN keywords. If you specify only the JOIN keyword, the database server creates an implicit inner join by default. An inner join returns all the rows in a table that have one or more matching rows in the other table (or tables). The unmatched rows are discarded.

ANSI LEFT OUTER Joins: The LEFT keyword specifies a join that treats the first table reference as the dominant table in the join. In a left outer join, the subordinate part of the outer join appears to the right of the keyword that begins the outer join specification. The result set includes all the rows that an INNER join returns, plus all rows that would otherwise have been discarded from the dominant table.

ANSI RIGHT OUTER Joins: The RIGHT keyword specifies a join that treats the second table reference as the dominant table in the join. In a right outer join, the subordinate part of the outer join appears to the left of the keyword that begins the outer join specification. The result set includes all the rows that an INNER join returns, plus all rows that would otherwise have been discarded from the dominant table.

ANSI FULL OUTER Joins: The FULL keyword specifies a join in which the result set includes all the rows from the Cartesian product for which the join condition is true, plus all the rows from each table that do not match the join condition.

In an ANSI-compliant join that specifies the LEFT, RIGHT, or FULL keywords in the FROM clause, the OUTER keyword is optional.

Join-method optimizer directives that you specify for an ANSI-compliant joined query are ignored, but are listed under *Directives Not Followed* in the **sqexplain.out** file.

Using the ON Clause

Use the ON clause to specify the join condition and any expressions as optional join filters.

The following example from the **stores_demo** database illustrates how the join condition in the ON clause combines the **customer** and **orders** tables:

```
SELECT c.customer_num, c.company, c.phone, o.order_date
FROM customer c LEFT JOIN orders o
ON c.customer_num = o.customer_num;
```

The following table shows part of the joined **customer** and **orders** tables.

customer_num	company	phone	order_date
101	All Sports Supplies	408-789-8075	05/21/2008
102	Sports Spot	415-822-1289	NULL
103	Phil's Sports	415-328-4543	NULL
104	Play Ball!	415-368-1100	05/20/2008
—	—	—	—

In an outer join, the join filters (expressions) that you specify in the ON clause determine which rows of the subordinate table join to the dominant (or outer) table. The dominant table, by definition, returns all its rows in the joined table. That is, a join filter in the ON clause has no effect on the dominant table.

If the ON clause specifies a join filter on the dominant table, the database server joins only those dominant table rows that meet the criterion of the join filter to rows in the subordinate table. The joined result contains all rows from the dominant table. Rows in the dominant table that do not meet the criterion of the join filter are extended with NULL values for the subordinate columns.

The following example from the **stores_demo** database illustrates the effect of a join filter in the ON clause:

```
SELECT c.customer_num, c.company, c.phone, o.order_date
FROM customer c LEFT JOIN orders o
ON c.customer_num = o.customer_num
AND c.company <> "All Sports Supplies";
```

The row that contains All Sports Supplies remains in the joined result.

customer_num	company	phone	order_date
101	All Sports Supplies	408-789-8075	NULL
102	Sports Spot	415-822-1289	NULL
103	Phil's Sports	415-328-4543	NULL
104	Play Ball!	415-368-1100	05/20/2008
—	—	—	—

Even though the order date for customer number 101 is 05/21/2008 in the **orders** table, the effect of placing the join filter (`c.company <> "All Sports Supplies"`) prevents this row in the dominant **customer** table from being joined to the subordinate **orders** table. Instead, a NULL value for **order_date** is extended to the row of All Sports Supplies.

Applying a join filter to a base table in the subordinate part of an outer join can improve performance. For more information, see your *IBM Informix Performance Guide*.

Specifying a Post-Join Filter: When you use the ON clause to specify the join, you can use the WHERE clause as a post-join filter. The database server applies the post-join filter of the WHERE clause to the results of the outer join.

The following example illustrates the use of a post-join filter. This query returns data from the **stores_demo** database. Suppose you want to determine which items in the catalog are not being ordered. The next query creates an outer join of the data from the **catalog** and **items** tables and then determines which catalog items from a specific manufacturer (HRO) have not sold:

```
SELECT c.catalog_num, c.stock_num, c.manu_code, i.quantity
FROM catalog c LEFT JOIN items i
ON c.stock_num = i.stock_num AND c.manu_code = i.manu_code
WHERE i.quantity IS NULL AND c.manu_code = "HRO";
```

The WHERE clause contains the post-join filter that locates the rows of HRO items in the catalog for which nothing has been sold.

When you apply a post-join filter to a base table in the dominant or subordinate part of an outer join, you might improve performance. For more information, see your *IBM Informix Performance Guide*.

Using a Join as the Dominant or Subordinate Part of an Outer Join: With the ANSI join syntax, you can nest joins. You can use a join as the dominant or subordinate part of an outer or inner join.

Suppose you want to modify the previous query (the post-join filter example) to get more information that will help you determine whether to continue carrying each unsold item in the catalog. You can modify the query to include information from the **stock** table so that you can see a short description of each unsold item with its cost:

```
SELECT c.catalog_num, c.stock_num, s.description, s.unit_price,
       s.unit_descr, c.manu_code, i.quantity
FROM (catalog c INNER JOIN stock s
      ON c.stock_num = s.stock_num
      AND c.manu_code = s.manu_code)
LEFT JOIN items i
      ON c.stock_num = i.stock_num
      AND c.manu_code = i.manu_code
WHERE i.quantity IS NULL
      AND c.manu_code = "HRO";
```

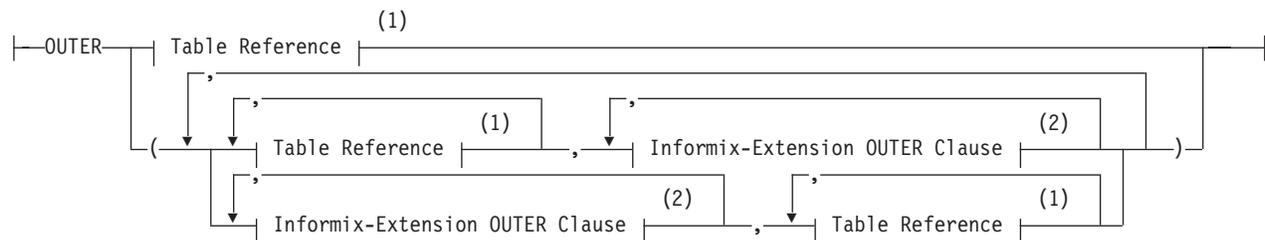
In this example, an inner join between the **catalog** and **stock** tables forms the dominant part of an outer join with the **items** table.

For additional examples of outer joins, see the *IBM Informix Guide to SQL: Tutorial*.

Informix-Extension Outer Joins

The Informix-extension syntax for outer joins begins with an implicit left outer join. That is, you begin an outer join with the **OUTER** keyword. This is the syntax of the Informix-extension **OUTER** clause.

Informix-Extension OUTER Clause:



Notes:

- 1 See "FROM Clause" on page 2-544
- 2 Informix extension

If you use this syntax for an outer join, you must use Informix-extension syntax for all outer joins in a single query block, and you must include the join condition in the **WHERE** clause. You cannot begin another outer join with the **LEFT JOIN** or the **LEFT OUTER JOIN** keywords.

This example uses the **OUTER** keyword to create an outer join that lists all customers and their orders, regardless of whether they have placed orders:

```
SELECT c.customer_num, c.lname, o.order_num FROM customer c,
       OUTER orders o WHERE c.customer_num = o.customer_num;
```

This example returns all the rows from the **customer** table with the rows that match in the **orders** table. If no record for a customer appears in the **orders** table, the returned **order_num** column for that customer has a NULL value.

If you have a complex outer join, that is, the query has more than one outer join, you must either embed the additional outer join or joins in parentheses, as the syntax diagram shows, or establish join conditions, or relationships, between the dominant table and each subordinate table in the WHERE clause.

When an expression or a condition in the WHERE clause relates two subordinate tables, you must use parentheses around the joined tables in the FROM clause to enforce dominant-subordinate relationships, as in this example:

```
SELECT c.company, o.order_date, i.total_price, m.manu_name
       FROM customer c,
          OUTER (orders o, OUTER (items i, OUTER manufact m))
       WHERE c.customer_num = o.customer_num
          AND o.order_num = i.order_num
          AND i.manu_code = m.manu_code;
```

When you omit parentheses around the subordinate tables in the FROM clause, you must establish join conditions between the dominant table and each subordinate table in the WHERE clause. If a join condition is between two subordinate tables, the query fails, but the following example successfully returns a result:

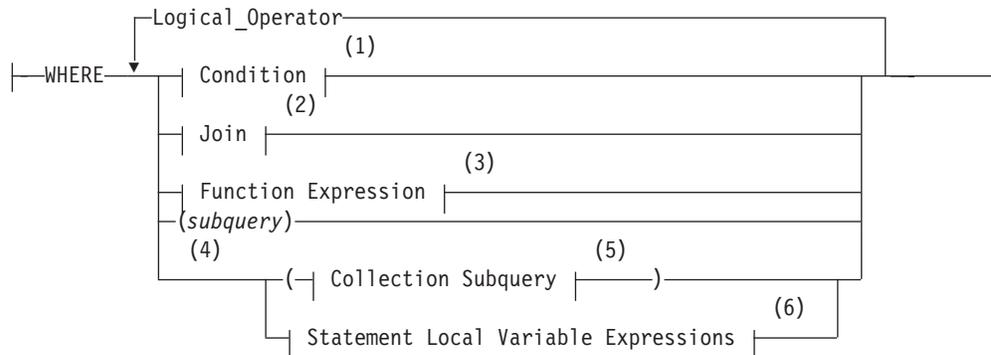
```
SELECT c.company, o.order_date, c2.call_descr
       FROM customer c, OUTER orders o, OUTER cust_calls c2
       WHERE c.customer_num = o.customer_num
          AND c.customer_num = c2.customer_num;
```

The *IBM Informix Guide to SQL: Tutorial* has examples of complex outer joins.

WHERE Clause of SELECT

The WHERE clause can specify join conditions for Informix-extension joins, post-join filters for ANSI-compliant joins, and for search criteria on data values.

WHERE Clause:



Notes:

- 1 See "Condition " on page 4-5

- 2 See "Specifying a Join in the WHERE Clause" on page 2-564
- 3 See "Function Expressions" on page 4-87
- 4 Dynamic Server only
- 5 See "Collection Subquery " on page 4-3
- 6 See "Statement-Local Variable Expressions (IDS)" on page 4-164

Element	Description	Restrictions	Syntax
<i>Logical_ Operator</i>	Combines two conditions	Valid options are <i>logical union</i> (= OR or OR NOT) or <i>logical intersection</i> (= AND or AND NOT)	"Conditions with AND or OR" on page 4-20
<i>subquery</i>	Embedded query	Cannot include the FIRST or ORDER BY keywords	"SELECT" on page 2-528

Using a Condition in the WHERE Clause

You can use these simple conditions or comparisons in the WHERE clause:

- Relational-operator condition
- IN or BETWEEN . . . AND
- IS NULL or IS NOT NULL
- LIKE or MATCHES

You also can use a SELECT statement within the WHERE clause; this is called a *subquery*. The following WHERE clause operators are valid in a subquery:

- IN or EXISTS
- ALL, ANY, or SOME

For more information, see "Condition " on page 4-5.

In the WHERE clause, an aggregate function is not valid unless it is part of a subquery or is on a correlated column originating from a parent query, and the WHERE clause is in a subquery within a HAVING clause.

Relational-Operator Condition: A relational-operator condition is satisfied if the expressions on each side of the operator fulfill the relation that the operator specifies. The following statements use the greater than (>) and equal (=) relational operators:

```
SELECT order_num FROM orders
   WHERE order_date > '6/04/08';
SELECT fname, lname, company
   FROM customer
   WHERE city[1,3] = 'San';
```

Single quotation marks are required around 'San' because the substring is from a character column. See the "Relational-Operator Condition" on page 4-9.

IN Condition: The IN condition is satisfied when the expression to the left of the IN keyword is included in the list of values to the right of the keyword.

The following examples show the IN condition:

```
SELECT lname, fname, company FROM customer
   WHERE state IN ('CA', 'WA', 'NJ');
SELECT * FROM cust_calls
   WHERE user_id NOT IN (USER );
```

For more information, see the “IN Subquery” on page 4-17.

BETWEEN Condition: The BETWEEN condition is satisfied when the value to the left of BETWEEN is in the inclusive range of the two values on the right of BETWEEN. The first two queries in the following example use literal values after the BETWEEN keyword. The third query uses the built-in CURRENT function and a literal interval to search for dates between the current day and seven days earlier.

```
SELECT stock_num, manu_code FROM stock
  WHERE unit_price BETWEEN 125.00 AND 200.00;
SELECT DISTINCT customer_num, stock_num, manu_code
  FROM orders, items
  WHERE order_date BETWEEN '6/1/07' AND '9/1/07';
SELECT * FROM cust_calls WHERE call_dtime
  BETWEEN (CURRENT - INTERVAL(7) DAY TO DAY) AND CURRENT;
```

For more information, see the “BETWEEN Condition” on page 4-9.

IS NULL and IS NOT NULL Conditions: The IS NULL condition is satisfied if the specified *column* contains a NULL value, or if the specified *expression* evaluates to NULL.

If you use the IS NOT NULL predicate, the condition is satisfied when the *column* contains a value that is not NULL, or when the *expression* does not evaluate to NULL. The following example selects the order numbers and customer numbers for which the order has not been paid:

```
SELECT order_num, customer_num FROM orders
  WHERE paid_date IS NULL;
```

For a complete description of the IS NULL and IS NOT NULL operators, see the “IS NULL and IS NOT NULL Conditions” on page 4-12.

LIKE or MATCHES Condition: The LIKE or MATCHES condition is satisfied if either of the following is true:

- The value of the column that precedes the LIKE or MATCHES keyword matches the pattern that the quoted string specifies. You can use wildcard characters in the string.
- The value of the column that precedes the LIKE or MATCHES keyword matches the pattern that is specified by the column that follows the LIKE or MATCHES keyword. The value of the column on the right serves as the matching pattern in the condition.

The following SELECT statement returns all rows in the **customer** table in which the **lname** column begins with the literal string 'Baxter'. Because the string is a literal string, the condition is case sensitive.

```
SELECT * FROM customer WHERE lname LIKE 'Baxter'
```

The next SELECT statement returns all rows in the **customer** table in which the value of the **lname** column matches the value of the **fname** column:

```
SELECT * FROM customer WHERE lname LIKE fname;
```

The following examples use the LIKE condition with a wildcard. The first SELECT statement finds all stock items that are some kind of ball. The second SELECT statement finds all company names that contain a percent (%) sign. Backslash (\) is used as the default escape character for the percent (%) sign wildcard. The third SELECT statement uses the ESCAPE option with the LIKE condition to

retrieve rows from the **customer** table in which the **company** column includes a percent (%) sign. The z is used as an escape character for the percent (%) sign:

```
SELECT stock_num, manu_code FROM stock
WHERE description LIKE '
SELECT * FROM customer WHERE company LIKE '
SELECT * FROM customer WHERE company LIKE '%z%' ESCAPE 'z';
```

The following examples use MATCHES with a wildcard in SELECT statements. The first SELECT statement finds all stock items that are some kind of ball. The second SELECT statement finds all company names that contain an asterisk (*). The backslash (\) is used as the default escape character for a literal asterisk (*) character. The third statement uses the ESCAPE option with the MATCHES condition to retrieve rows from the **customer** table where the **company** column includes an asterisk (*). The z character is specified as an escape character for the asterisk (*) character:

```
SELECT stock_num, manu_code FROM stock
WHERE description MATCHES '*ball';

SELECT * FROM customer WHERE company MATCHES '*\**';

SELECT * FROM customer WHERE company MATCHES '*z**' ESCAPE 'z';
```

See also the “LIKE and MATCHES Condition” on page 4-13.

IN Subquery: With the IN subquery, more than one row can be returned, but only one column can be returned.

This example shows the use of an IN subquery in a SELECT statement:

```
SELECT DISTINCT customer_num FROM orders
WHERE order_num NOT IN
  (SELECT order_num FROM items
   WHERE stock_num = 1);
```

For additional information, see the “ IN Condition” on page 4-10.

EXISTS Subquery: With the EXISTS subquery, one or more columns can be returned.

The following example of a SELECT statement with an EXISTS subquery returns the stock number and manufacturer code for every item that has never been ordered (and is therefore not listed in the **items** table).

It is appropriate to use an EXISTS subquery in this SELECT statement because you need the correlated subquery to test both **stock_num** and **manu_code** in the **items** table.

```
SELECT stock_num, manu_code FROM stock
WHERE NOT EXISTS
  (SELECT stock_num, manu_code FROM items
   WHERE stock.stock_num = items.stock_num AND
   stock.manu_code = items.manu_code);
```

The preceding example would work equally well if you use a SELECT * in the subquery in place of the column names, because you are testing for the existence of a row or rows.

For additional information, see the “ EXISTS Subquery” on page 4-18.

ALL, ANY, SOME Subqueries: The following examples return the order number of all orders that contain an item whose total price is greater than the total price of every item in order number 1023. The first SELECT uses the ALL subquery, and the second SELECT produces the same result by using the MAX aggregate function.

```
SELECT DISTINCT order_num FROM items
  WHERE total_price > ALL (SELECT total_price FROM items
    WHERE order_num = 1023);
```

```
SELECT DISTINCT order_num FROM items
  WHERE total_price > SELECT MAX(total_price) FROM items
    WHERE order_num = 1023);
```

The following SELECT statements return the order number of all orders that contain an item whose total price is greater than the total price of at least one of the items in order number 1023. The first SELECT statement uses the ANY keyword, and the second SELECT statement uses the MIN aggregate function:

```
SELECT DISTINCT order_num FROM items
  WHERE total_price > ANY (SELECT total_price FROM items
    WHERE order_num = 1023);
```

```
SELECT DISTINCT order_num FROM items
  WHERE total_price > (SELECT MIN(total_price) FROM items
    WHERE order_num = 1023);
```

You can omit the keywords ANY, ALL, or SOME in a subquery if the subquery returns exactly one value. If you omit ANY, ALL, or SOME, and the subquery returns more than one value, you receive an error. The subquery in the next example returns only one row, because it uses an aggregate function:

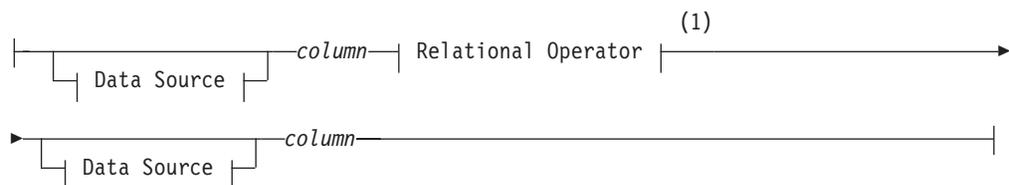
```
SELECT order_num FROM items
  WHERE stock_num = 9 AND quantity =
    (SELECT MAX(quantity) FROM items WHERE stock_num = 9);
```

See also “ ALL, ANY, and SOME Subqueries” on page 4-18.

Specifying a Join in the WHERE Clause

You join two tables by creating a relationship in the WHERE clause between at least one column from one table and at least one column from another. The join creates a temporary composite table where each pair of rows (one from each table) that satisfies the join condition is linked to form a single row.

Join:



Data Source:



Notes:

- 1 See "Relational Operator" on page 4-193
- 2 Extended Parallel Server only

Element	Description	Restrictions	Syntax
<i>alias</i>	Temporary alternative name declared in the FROM clause for a table or view	See " Self-Joins"; "FROM Clause" on page 2-544	"Identifier " on page 5-21
<i>column</i>	Column of a table or view to be joined	Must exist in the table or view	"Identifier " on page 5-21
<i>external</i>	External table from which to retrieve data	External table must exist	"Database Object Name " on page 5-16
<i>synonym, table, view</i>	Name of a synonym, table, or view to be joined in the query	Synonym and the table or view to which it points must exist	"Database Object Name " on page 5-16

Rows from the tables or views are *joined* when there is a match between the values of specified columns. When the columns to be joined have the same name, you must qualify each column name with its data source.

Two-Table Joins:

You can create two-table joins, multiple-table joins, self-joins, and outer joins (Informix-extension syntax). The following example shows a two-table join:

```
SELECT order_num, lname, fname FROM customer, orders
WHERE customer.customer_num = orders.customer_num;
```

Multiple-Table Joins:

A multiple-table join is a join of more than two tables. Its structure is similar to the structure of a two-table join, except that you have a join condition for more than one pair of tables in the WHERE clause. When columns from different tables have the same name, you must qualify the column name with its associated table or table alias, as in *table.column*. For the full syntax of a table name, see "Database Object Name " on page 5-16.

The following multiple-table join yields the company name of the customer who ordered an item as well as its stock number and manufacturer code:

```
SELECT DISTINCT company, stock_num, manu_code
FROM customer c, orders o, items i
WHERE c.customer_num = o.customer_num
AND o.order_num = i.order_num;
```

Self-Joins:

You can join a table to itself. To do so, you must list the table name twice in the FROM clause and assign it two different table aliases. Use the aliases to refer to each of the *two* tables in the WHERE clause. The next example is a self-join on the **stock** table. It finds pairs of stock items whose unit prices differ by a factor greater than 2.5. The letters x and y are each aliases for the **stock** table.

```
SELECT x.stock_num, x.manu_code, y.stock_num, y.manu_code
       FROM stock x, stock y WHERE x.unit_price > 2.5 * y.unit_price;
```

Extended Parallel Server does not support self-joins with an external table.

Informix-Extension Outer Joins: The next outer join lists the company name of the customer and all associated order numbers, if the customer has placed an order. If not, the company name is still listed, and a NULL value is returned for the order number.

```
SELECT company, order_num FROM customer c, OUTER orders o
       WHERE c.customer_num = o.customer_num;
```

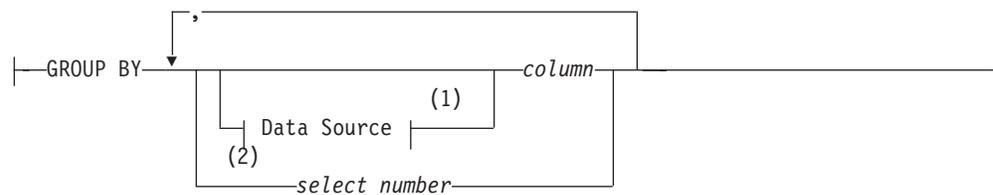
In Extended Parallel Server, you cannot use an external table as the outer table in an outer join.

For more information about outer joins, see the *IBM Informix Guide to SQL: Tutorial*.

GROUP BY Clause

Use the GROUP BY clause to produce a single row of results for each group. A *group* is a set of rows that have the same values for each column listed.

GROUP BY Clause:



Notes:

- 1 See "WHERE Clause of SELECT" on page 2-560
- 2 Informix extension

Element	Description	Restrictions	Syntax
<i>column</i>	A column (or an expression) from the select list of the Projection clause	See "Relationship of GROUP BY and Projection Clauses" on page 2-567	"Identifier " on page 5-21
<i>select_number</i>	Integer specifying ordinal position of a column or expression in select list	See "Using Select Numbers" on page 2-567.	"Literal Number " on page 4-185

The SELECT statement with a GROUP BY clause returns a single row of results for each group of rows that have the same value in *column*, or that have the same value in the column or expression that the *select_number* specifies.

Relationship of GROUP BY and Projection Clauses

A GROUP BY clause restricts what the Projection clause can specify. If you use a GROUP BY clause, each *column* specified in the select list of the Projection clause must also be included in the GROUP BY clause.

If you use an aggregate function and one or more column expressions in the select list, you must put all the column names that are not used as part of an aggregate or time expression in the GROUP BY clause.

Constant expressions and BYTE or TEXT column expressions are not valid in the GROUP BY list.

If the select list includes a BYTE or TEXT column, you cannot use the GROUP BY clause. In addition, you cannot include a ROWID in a GROUP BY clause.

In Dynamic Server, if your select list includes a column of a user-defined data type, the column cannot be used in a GROUP BY clause unless the UDT can use the built-in bit-hashing function. Any UDT that cannot use the built-in bit-hashing function must be created with the CANNOTHASH modifier, which tells the database server that the UDT cannot be used in a GROUP BY clause.

The following example specifies one column that is not in an aggregate expression. The **total_price** column should not be in the GROUP BY list because it appears as the argument of an aggregate function. The COUNT and SUM aggregates are applied to each group, not to the whole query set.

```
SELECT order_num, COUNT(*), SUM(total_price)
FROM items GROUP BY order_num;
```

If a column stands alone in a column expression in the select list, you must use it in the GROUP BY clause. If a column is combined with another column by an arithmetic operator, you can choose to group by the individual columns or by the combined expression using the number.

NULL Values in the GROUP BY Clause

In a column listed in a GROUP BY clause, each row that contains a NULL value belongs to a single group. That is, all NULL values are grouped together.

Using Select Numbers

You can use one or more integers in the GROUP BY clause to stand for column expressions. In the next example, the first SELECT statement uses select numbers for **order_date** and **paid_date - order_date** in the GROUP BY clause. You can group only by a combined expression using the select numbers.

In the second SELECT statement, you cannot replace the 2 with the arithmetic expression **paid_date - order_date**:

```
SELECT order_date, COUNT(*), paid_date - order_date
FROM orders GROUP BY 1, 3;
SELECT order_date, paid_date - order_date
FROM orders GROUP BY order_date, 2;
```

HAVING Clause

Use the HAVING clause to apply one or more qualifying conditions to groups.

HAVING Clause:

|—HAVING—| Condition | (1) |

Notes:

1 See “Condition ” on page 4-5

In the following examples, each condition compares one calculated property of the group with another calculated property of the group or with a constant. The first SELECT statement uses a HAVING clause that compares the calculated expression COUNT(*) with the constant 2. The query returns the average total price per item on all orders that have more than two items.

The second SELECT statement lists customers and the call months for customers who have made two or more calls in the same month:

```
SELECT order_num, AVG(total_price) FROM items
  GROUP BY order_num HAVING COUNT(*) > 2;
SELECT customer_num, EXTEND (call_dtime, MONTH TO MONTH)
  FROM cust_calls GROUP BY 1, 2 HAVING COUNT(*) > 1;
```

You can use the HAVING clause to place conditions on the GROUP BY column values as well as on calculated values. This example returns **cust_code** and **customer_num**, **call_dtime**, and groups them by **call_code** for all calls that have been received from customers with **customer_num** less than 120:

```
SELECT customer_num, EXTEND (call_dtime), call_code
  FROM cust_calls GROUP BY call_code, 2, 1
  HAVING customer_num < 120;
```

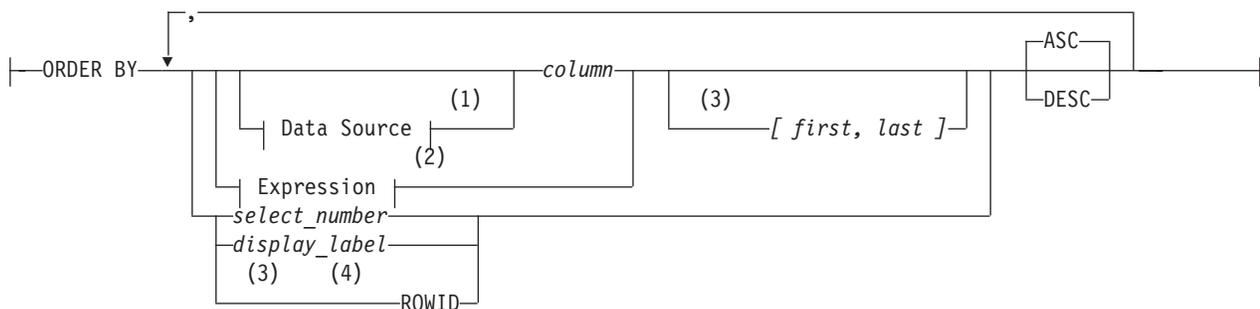
The HAVING clause generally complements a GROUP BY clause. If you omit the GROUP BY clause, the HAVING clause applies to all rows that satisfy the query, and all rows in the table make up a single group. The following example returns the average price of all the values in the table, as long as more than ten rows are in the table:

```
SELECT AVG(total_price) FROM items HAVING COUNT(*) > 10;
```

ORDER BY Clause

The ORDER BY clause sorts query results by specified columns or expressions.

ORDER BY Clause:



Notes:

- 1 See "WHERE Clause of SELECT" on page 2-560
- 2 See "Expression" on page 4-40
- 3 Informix extension
- 4 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>column</i>	Sort rows by value in this column	Must also be in select list (Extended Parallel Server only).	"Identifier " on page 5-21
<i>display_label</i>	Temporary name for a column or for a column expression	Must be unique among labels declared in the Projection clause	"Identifier " on page 5-21
<i>first, last</i>	First and last byte in column substring to sort the result set	Integers; for BYTE, TEXT, and character data types only	"Literal Number " on page 4-185
<i>select_number</i>	Ordinal position of a column in select list of the Projection clause	See "Using Select Numbers" on page 2-567.	"Literal Number " on page 4-185

The ORDER BY clause implies that the query returns more than one row. In SPL, the database server issues an error if you specify the ORDER BY clause without a FOREACH loop to process the returned rows individually within the SPL routine.

The following query specifies a derived table in the FROM clause whose rows are ordered by the **col1** value, and declares **vtab** as the name of the derived table, and **vcol** as the name of its only column:

```
SELECT vcol FROM (SELECT FIRST 5 col1 FROM tab1 ORDER BY col1) vtab(vcol);
```

Ordering by a Column or by an Expression

For Extended Parallel Server, any *column* in the ORDER BY clause must also appear explicitly or by * notation in the select list of the Projection clause. To order query results by an expression, you must also declare a display label for the expression in the Projection clause, as in the following example, which declares the display label **span** for the difference between two columns:

```
SELECT paid_date - ship_date span, customer_num FROM orders  
ORDER BY span;
```

Dynamic Server supports columns and expressions in the ORDER BY clause that do not appear in the select list of the Projection clause. You can omit a display label for the derived column in the select list and specify the derived column by means of a select number in the ORDER BY clause.

The select list of the Projection clause must include any column or expression that the ORDER BY clause specifies, however, if any of the following is true:

- The query includes the DISTINCT, UNIQUE, or UNION operator.
- The query includes the INTO TEMP *table* clause.
- The distributed query accesses a remote database whose server requires every column or expression in the ORDER BY clause to also appear in the select list of the Projection clause.
- An expression in the ORDER BY clause includes a display label for a column substring. (See the next section, "Ordering by a Substring" on page 2-570.)

The next query selects one column from the **orders** table and sorts the results by the value of another column. (With Extended Parallel Server, **order_date** must also appear in the Projection clause.) By default, the rows are listed in ascending order.

```
SELECT ship_date FROM orders ORDER BY order_date;
```

You can order by an aggregate expression only if the query also has a GROUP BY clause. This query declares the display label **maxwgt** for an aggregate in the ORDER BY clause:

```
SELECT ship_charge, MAX(ship_weight) maxwgt
   FROM orders GROUP BY ship_charge ORDER BY maxwgt;
```

If the current processing locale defines a localized collation, then NCHAR and NVARCHAR column values are sorted in that localized order.

In Dynamic Server, no *column* in the ORDER BY clause can be a collection type, but a query whose result set defines a collection-derived table can include the ORDER BY clause. For an example, see “Collection-Derived Table” on page 5-4.

You might improve the performance of some non-PDQ queries that use the ORDER BY clause to sort a large set of rows if you increase the setting of the DS_NONPDQ_QUERY_MEM configuration parameter.

Ordering by a Substring

You can order by a substring instead of by the entire length of a character, BYTE, or TEXT column, or of an expression returning a character string. The database server uses the substring to sort the result set. Define the substring by specifying integer subscripts (the *first* and *last* parameters), representing the starting and ending byte positions of the substring within the column value.

The following SELECT statement queries the **customer** table and specifies a column substring in the ORDER BY column. This instructs the database server to sort the query results by the portion of the **lname** column contained in the sixth through ninth bytes of the column value.

```
SELECT * from customer ORDER BY lname[6,9];
```

Assume that the value of **lname** in one row of the **customer** table is Greenburg. Because of the column substring in the ORDER BY clause, the database server determines the sort position of this row by using the value *burg*, rather than the complete column value Greenburg.

When ordering by an expression, you can specify substrings only for expressions that return a character data type. If you specify a column substring in the ORDER BY clause, the column must have one of the following data types: BYTE, CHAR, NCHAR, NVARCHAR, TEXT, or VARCHAR.

Dynamic Server can also support LVARCHAR column substrings in the ORDER BY clause, if the column is in a database of the local database server.

For information on the GLS aspects of using column substrings in the ORDER BY clause, see the *IBM Informix GLS User's Guide*.

Ascending and Descending Orders

You can use the ASC and DESC keywords to specify ascending (smallest value first) or descending (largest value first) order. The default order is ascending. For DATE and DATETIME data types, *smallest* means earliest in time and *largest* means

latest in time. For character data types in the default locale, the order is the ASCII collating sequence, as listed in “Collating Order for U.S. English Data” on page 4-195.

For NCHAR or NVARCHAR data types, the localized collating order of the current session is used, if that is different from the code set order. For more information about collation, see “SET COLLATION” on page 2-584.

If you specify the ORDER BY clause, NULL values are ordered as less than values that are not NULL. Using the ASC order, a NULL value comes before any non-NULL value; using DESC order, the NULL comes last.

Nested Ordering

If you list more than one column in the ORDER BY clause, your query is ordered by a nested sort. The first level of sort is based on the first column; the second column determines the second level of sort. The following example of a nested sort selects all the rows in the **cust_calls** table and orders them by **call_code** and by **call_dtime** within **call_code**:

```
SELECT * FROM cust_calls ORDER BY call_code, call_dtime;
```

Using Select Numbers

In place of column names, you can enter in the ORDER BY clause one or more integers that refer to the position of items listed in the select list of the Projection clause. You can also use a select number to sort by an expression.

The following example orders by the expression **paid_date - order_date** and **customer_num**, using select numbers in a nested sort:

```
SELECT order_num, customer_num, paid_date - order_date
FROM orders
ORDER BY 3, 2;
```

Select numbers are required in the ORDER BY clause when SELECT statements are joined by the UNION or UNION ALL keywords, or when compatible columns in the same position have different names.

Ordering by Rowids

You can specify the ROWID keyword in the ORDER BY clause. This specifies the **rowid** column, a hidden column in nonfragmented tables and in fragmented tables that were created with the WITH ROWIDS clause. The **rowid** column contains a unique internal record number that is associated with a row in a table. (It is recommended, however, that you utilize primary keys as your access method, rather than exploiting the **rowid** column.)

The ORDER BY clause cannot specify the **rowid** column if the table from which you are selecting is a fragmented table that has no **rowid** column.

In Extended Parallel Server, you cannot specify the ROWID keyword in the ORDER BY clause unless you also included ROWID in the Projection clause.

In Dynamic Server, you do not need to include the ROWID keyword in the Projection clause when you specify ROWID in the ORDER BY clause.

For further information about **rowid** values and how to use the **rowid** column in column expressions, see “WITH ROWIDS Option” on page 2-19 and “Using Rowids (IDS)” on page 4-64.


```

{
EXEC SQL fetch xc into $fname, $lname;
if (strncmp(SQLSTATE, '00', 2) != 0) break;
printf("%d %s %s\n",cnum, fname, lname );
if (cnum == 999)          --update rows with 999 customer_num
    EXEC SQL update customer set fname = 'rosey' where current of xc;
}

EXEC SQL close xc;
EXEC SQL disconnect current;

```

A SELECT...FOR UPDATE statement, like an update cursor, allows you to perform updates that are not possible with the UPDATE statement alone, because both the decision to update and the values of the new data items can be based on the original contents of the row. The UPDATE statement cannot query the table that is being updated.

Syntax That is Incompatible with the FOR UPDATE Clause

A SELECT statement that uses a FOR UPDATE clause must conform to the following restrictions:

- The statement can select data from only one table.
- The statement cannot include any aggregate functions.
- The statement cannot include any of the following clauses or keywords: DISTINCT, FOR READ ONLY, GROUP BY, INTO SCRATCH, INTO TEMP, INTO EXTERNAL, ORDER BY, UNION, or UNIQUE.

For information on how to declare an update cursor for a SELECT statement that does not include a FOR UPDATE clause, see “Using the FOR UPDATE Option” on page 2-290.

FOR READ ONLY Clause

Use the FOR READ ONLY keywords to specify that the Select cursor declared for the SELECT statement is a read-only cursor. A read-only cursor is a cursor that cannot modify data. This section provides the following information about the FOR READ ONLY clause:

- When you must use the FOR READ ONLY clause
- Syntax restrictions on a SELECT statement that uses a FOR READ ONLY clause

Using the FOR READ ONLY Clause in Read-Only Mode

Normally, you do not need to include the FOR READ ONLY clause in a SELECT statement. SELECT is a read-only operation by definition, so the FOR READ ONLY clause is usually unnecessary. In certain circumstances, however, you must include the FOR READ ONLY keywords in a SELECT statement.

If you have used the High-Performance Loader (HPL) in express mode to load data into the tables of an ANSI-compliant database, and you have not yet performed a level-0 backup of this data, the database is in read-only mode. When the database is in read-only mode, the database server rejects any attempts by a Select cursor to access the data unless the SELECT or the DECLARE includes a FOR READ ONLY clause. This restriction remains in effect until the user has performed a level-0 backup of the data.

In an ANSI-compliant database, Select cursors are update cursors by default. An update cursor is a cursor that can be used to modify data. These update cursors

are incompatible with the read-only mode of the database. For example, this SELECT statement against the **customer_ansi** table fails:

```
EXEC SQL declare ansi_curs cursor for
select * from customer_ansi;
```

The solution is to include the FOR READ ONLY clause in your Select cursors. The read-only cursor that this clause specifies is compatible with the read-only mode of the database. For example, the following SELECT FOR READ ONLY statement against the **customer_ansi** table succeeds:

```
EXEC SQL declare ansi_read cursor for
select * from customer_ansi for read only;
```

DB–Access executes all SELECT statements with Select cursors, so you must specify FOR READ ONLY in all queries that access data in a read-only ANSI-compliant database. The FOR READ ONLY clause causes DB–Access to declare the cursor for the SELECT statement as a read-only cursor.

For more information on level-0 backups, see your *IBM Informix Backup and Restore Guide*. For more information on Select cursors, read-only cursors, and update cursors, see “DECLARE ” on page 2-286.

For more information on the express mode of the HPL of Dynamic Server, see the *IBM Informix High-Performance Loader User’s Guide*.

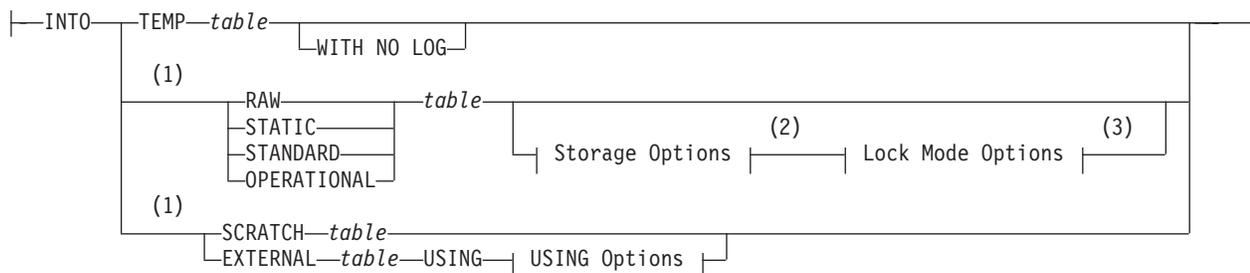
Syntax That Is Incompatible with the FOR READ ONLY Clause

If you attempt to include both the FOR READ ONLY clause and the FOR UPDATE clause in the same SELECT statement, the SELECT statement fails. For information on declaring a read-only cursor for a SELECT statement that does not include a FOR READ ONLY clause, see “DECLARE ” on page 2-286.

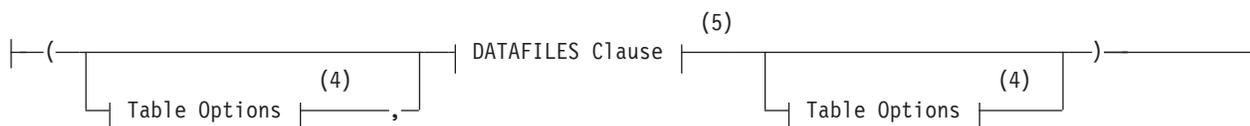
INTO Table Clauses

Use the INTO Table clauses to specify a temporary or external table to receive the data that the SELECT statement retrieves.

INTO Table Clauses:



USING Options:



Notes:

- 1 Extended Parallel Server only
- 2 See "Storage Options" on page 2-210
- 3 See "LOCK MODE Options" on page 2-225
- 4 See "Table Options" on page 2-577
- 5 See "DATAFILES Clause" on page 2-108

Element	Description	Restrictions	Syntax
<i>table</i>	Table to receive the results of the query	Must be unique among names of tables, views, and synonyms that you own in the current database	"Database Object Name " on page 5-16

You must have the Connect privilege on a database to create a temporary or external table in that database. The name of a temporary table need not be unique among the names of temporary tables of other users.

Column names in the temporary or external table must be specified in the Projection clause, where you must supply a display label for all expressions that are not simple column expressions. The display label becomes the column name in the temporary or external table. If you do not declare a display label for a column expression, the table uses the column name from the select list of the Projection clause.

The following INTO TEMP example creates the **pushdate** table with two columns, **customer_num** and **slowdate**:

```
SELECT customer_num, call_dtime + 5 UNITS DAY slowdate
FROM cust_calls INTO TEMP pushdate;
```

Results When No Rows are Returned

When you use an INTO Table clause combined with the WHERE clause, and no rows are returned, the SQLNOTFOUND value is 100 in ANSI-compliant databases and 0 in databases that are not ANSI compliant. If the SELECT INTO TEMP...WHERE... statement is a part of a multistatement PREPARE and no rows are returned, the SQLNOTFOUND value is 100 for both ANSI-compliant databases and databases that are not ANSI-compliant.

This release of Dynamic Server continues to process the remaining statements of a multistatement prepared object after encountering the SQLNOTFOUND value of 100. You can maintain the legacy behavior, however, of not executing the remaining prepared statements by setting the **IFX_MULTIPREPSTMT** environment variable to 1.

Restrictions with INTO Table Clauses in ESQL/C

In Informix ESQL/C, do not use the INTO clause with an INTO *table* clause. If you do, no results are returned to the program variables and the **SQLCODE** variable is set to a negative value.

INTO TEMP Clause

The INTO TEMP clause creates a temporary table to hold the query results. The default initial and next extents for a temporary table are four pages. The temporary table must be accessible by the built-in RSAM access method of the database server; you cannot specify another access method.

If you use the same query results more than once, using a temporary table saves time. In addition, using an INTO TEMP clause often gives you clearer and more understandable SELECT statements.

Data values in a temporary table are static; they are not updated as changes are made to the tables that were used to build the temporary table. You can use the CREATE INDEX statement to create indexes on a temporary table.

A logged, temporary table exists until one of the following situations occurs:

- The application disconnects from the database.
- A DROP TABLE statement is issued on the temporary table.
- The database is closed.

For Dynamic Server, if your database does not have transaction logging, the temporary table behaves in the same way as a table created with the WITH NO LOG option.

If you specify more than one temporary dbpace in the **DBSPACETEMP** environment variable (or if this is not set, in the DBSPACETEMP configuration parameter), the INTO TEMP clause loads the rows of the results set of the query into each of these dbspaces in round-robin fashion. For more information about the storage location of temporary tables that queries with the INTO TEMP clause create, see “Where Temporary Tables are Stored” on page 2-237.

Using SELECT INTO to Create a New Permanent Table

Using Extended Parallel Server, you can use the SELECT INTO statement to create a new permanent table based on the result set of a SELECT statement.

When using SELECT INTO to create a new table, you must specify its type. You can optionally specify storage and lock mode options for the new table.

The column names of the new permanent table are the names specified in the select list. If an asterisk (*) appears as the select list, it is expanded to all the columns of the corresponding tables or views in the SELECT statement.

All expressions other than simple column expressions must have a display label. This is used as the name of the column in the new table. If a column expression has no display label, the table uses the column name. If there are duplicate labels or column names in the select list, an error is returned.

Using the WITH NO LOG Option

Use the WITH NO LOG option to reduce the overhead of transaction logging because operations on nonlogging temporary tables are not logged.

A nonlogging temporary table exists until one of the following events occurs:

- The application disconnects.
- A DROP TABLE statement is issued on the temporary table.

Because nonlogging temporary tables do not disappear when the database is closed, you can use a nonlogging temporary table to transfer data from one database to another while the application remains connected. The behavior of a temporary table that you create with the WITH NO LOG option of the INTO TEMP clause is the same as that of a scratch table.

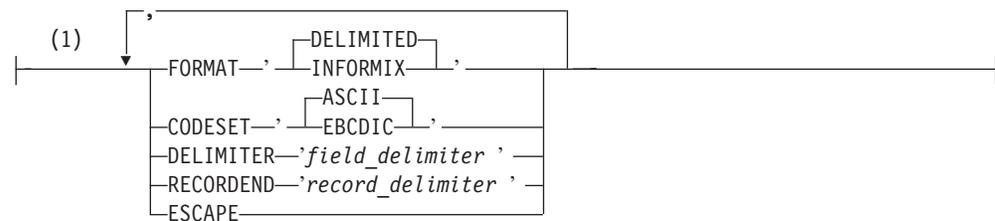
For more information about temporary tables, see “CREATE Temporary TABLE ” on page 2-231.

INTO EXTERNAL Clause (XPS)

The INTO EXTERNAL clause unloads query results into an external table, creating a default external table description that you can use when you subsequently reload the files. This combines the functionality of the CREATE EXTERNAL TABLE . . . SAMEAS and INSERT INTO . . . SELECT statements.

Table Options: The SELECT statement supports a subset of the CREATE EXTERNAL TABLE syntax for “Table Options” on page 2-109. Use the Table Options clause of the SELECT INTO EXTERNAL statement to specify the format of the unloaded data in the external table.

Table Options:



Notes:

- 1 Extended Parallel Server only

Element	Description	Restrictions	Syntax
<i>field_delimiter</i>	Character to separate fields. Default is pipe () character	See “Specifying Delimiters” on page 2-578.	“Quoted String” on page 4-189
<i>record_delimiter</i>	Character to separate records	See “Specifying Delimiters” on page 2-578.	“Quoted String” on page 4-189

The following table describes the keywords that apply to unloading data. If you want to specify additional table options in the external-table description for the purpose of reloading the table later, see “Table Options” on page 2-109.

In the SELECT ... INTO EXTERNAL statement, you can specify all table options that are discussed in the CREATE EXTERNAL TABLE statement except the fixed-format option.

You can use the INTO EXTERNAL clause when the format type of the created data file is either delimited text (if you use the DELIMITED keyword) or text in Informix internal data format (if you use the INFORMIX keyword). You cannot use it for a fixed-format unload.

Keyword

Effect

CODESET

Specifies the type of code set. Options are EBCDIC or ASCII.

DELIMITER

Specifies the character that separates fields in a delimited text file

ESCAPE

Directs the database server to recognize ASCII special characters embedded in ASCII-text-based data files

If you do not specify ESCAPE when you load data, the database server does not check the character fields in text data files for embedded special characters.

If you do not specify ESCAPE when you unload data, the database server does not create embedded hexadecimal characters in text fields.

FORMAT

Specifies the format of the data in the data files

RECORDEND

Specifies the character that separates records in a delimited text file

Specifying Delimiters: If you do not set the **RECORDEND** environment variable, the default value for *record_delimiter* is the newline character (CTRL-J).

If you use a non-printing character as a delimiter, encode it as the octal representation of the ASCII character. For example, '\006' can represent CTRL-F.

For more information on external tables, see "CREATE EXTERNAL TABLE (XPS) " on page 2-104.

INTO SCRATCH Clause (XPS)

Extended Parallel Server supports the INTO Scratch clause. This can improve performance, because scratch tables are not logged. A scratch table does not support indexes or constraints. It persists until the application disconnects, or a DROP TABLE statement is issued on the temporary table.

Because scratch tables do not disappear when the database is closed, you can use a scratch table to transfer data from one database to another while the application remains connected. A scratch table is identical to a temporary table that is created with the WITH NO LOG option. For more information about scratch tables, see "CREATE Temporary TABLE " on page 2-231.

UNION Operator

Place the UNION operator between two SELECT statements to combine the queries into a single query. You can string several SELECT statements together using the UNION operator. Corresponding items do not need to have the same name. Omitting the ALL keyword excludes duplicate rows.

Restrictions on a Combined SELECT

Several restrictions apply on the queries that you can connect with a UNION operator, as the following list describes:

- In Informix ESQL/C, you cannot use an INTO clause in a compound query unless exactly one row is returned, and you are not using a cursor. In this case, the INTO clause must be in the first SELECT statement.
- The number of items in the Projection clause of each query must be the same, and the corresponding items in each Projection clause must have compatible data types.
- The Projection clause of each query cannot specify BYTE or TEXT columns. (This restriction does not apply to UNION ALL operations.)

- If you use an ORDER BY clause, it must follow the last Projection clause, and you must refer to the item ordered by integer, not by identifier. Sorting takes place after the set operation is complete.
- You can store the combined results of a UNION operator in a temporary table, but the INTO TEMP clause can appear only in the final SELECT statement.

A *UNION subquery* is a query that includes the UNION operator within a subquery. The following additional restrictions affect UNION subqueries:

- The CREATE VIEW statement cannot specify a UNION subquery to define the view.
- UNION subqueries cannot be triggering events. If a valid UNION subquery specifies a column on which a Select trigger has been defined, the query succeeds, but the trigger (or the INSTEAD OF trigger on a view) is ignored.
- Only columns in the local database are valid in a UNION subquery. You cannot reference a remote table or view in a UNION subquery.
- General expressions that include host variables are not valid on the left of the ALL, ANY, IN, NOT IN and SOME operators in a query that includes a UNION subquery. An expression that consists solely of a single host variable. however, is valid in this context.

For example, the following query is valid under the above restriction:

```
SELECT col1 FROM tab1 WHERE ? <= ALL
    (SELECT col2 FROM tab2 UNION SELECT col3 FROM tab3);
```

In this example, the expression to the left of ALL is a single host variable (?), which is the only expression involving host variables that is supported before the ALL, ANY, IN, NOT IN, or SOME operators in a query that includes a UNION subquery.

In contrast, the following example shows an invalid query:

```
SELECT col1 FROM tab1 WHERE (? + 8) <= ALL
    (SELECT col2 FROM tab2 UNION SELECT col3 FROM tab3);
```

This query fails because an operand of the relational <= operator to the left of the ALL operator is (? + 8), an arithmetic expression that includes a host variable, which is not valid syntax in a UNION subquery.

Expressions that do not contain host variables are not subject to this restriction. Thus, the following query (which includes the same UNION subquery) is valid:

```
SELECT col1 FROM tab1 WHERE (col1 + 8) <= ALL
    (SELECT col2 FROM tab2 UNION SELECT col3 FROM tab3);
```

Duplicate Rows in a Combined SELECT

If you use the UNION operator alone, the duplicate rows are removed from the complete set of rows. That is, if multiple rows contain identical values in each column, only one row is retained. If you use the UNION ALL operator, all the selected rows are returned (the duplicates are not removed).

The next example uses UNION ALL to join two SELECT statements without removing duplicates. The query returns a list of all the calls that were received during the first quarter of 2007 and the first quarter of 2008.

```
SELECT customer_num, call_code FROM cust_calls
    WHERE call_dtime BETWEEN
        DATETIME (2007-1-1) YEAR TO DAY
        AND DATETIME (2007-3-31) YEAR TO DAY
```

```

UNION ALL
SELECT customer_num, call_code FROM cust_calls
  WHERE call_dtime BETWEEN
    DATETIME (2008-1-1)YEAR TO DAY
    AND DATETIME (2008-3-31) YEAR TO DAY;

```

If you want to remove duplicates, use the UNION operator without the keyword ALL in the query. In the preceding example, if the combination 101 B were returned in both SELECT statements, a UNION operator would cause the combination to be listed once. (If you want to remove duplicates within each SELECT statement, use the DISTINCT keyword in the Projection clause, as described in “Projection Clause” on page 2-531.)

UNION in Subqueries

You can use the UNION and UNION ALL operators in subqueries of SELECT statements within the WHERE clause, the FROM clause, and in collection subqueries. In this release of Dynamic Server, however, subqueries that include UNION or UNION ALL are not supported in the following contexts:

- In the definition of a view
- In the event or in the Action clause of a trigger
- With the FOR UPDATE clause or with an Update cursor
- In a distributed query (accessing tables outside the local database)

For more information about collection subqueries, see “Collection Subquery ” on page 4-3. For more information about the FOR UPDATE clause, see “FOR UPDATE Clause” on page 2-572.

In a combined subquery, the database server can resolve a column name only within the scope of its qualifying table reference. The following query, for example, returns an error:

```

SELECT * FROM t1 WHERE EXISTS
  (SELECT a FROM t2
   UNION
   SELECT b FROM t3 WHERE t3.c IN
     (SELECT t4.x FROM t4 WHERE t4.4 = t2.z));

```

Here **t2.z** in the innermost subquery cannot be resolved, because **z** occurs outside the scope of reference of the table reference **t2**. Only column references that belong to **t4**, **t3**, or **t1** can be resolved in the innermost subquery. The scope of a table reference extends downwards through subqueries, but not across the UNION operator to sibling SELECT statements.

Related Information

Because the SELECT statement is “the Q in SQL,” most features of the database server directly or indirectly support SELECT operations, which are central to relational and object-relational databases. (So this section is not comprehensive.)

For task-oriented discussions of the SELECT statement, see the *IBM Informix Guide to SQL: Tutorial*.

For a discussion of the GLS aspects of the SELECT statement, see the *IBM Informix GLS User's Guide*.

For information on how to access row and collection values with Informix ESQL/C host variables, see the discussion of complex data types in the *IBM Informix ESQL/C Programmer's Manual*.

SET ALL_MUTABLES

Use the SET ALL_MUTABLES statement to specify the modifiability of all mutable environment variables in a user session.

Only Extended Parallel Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```
▶▶ SET ALL_MUTABLES TO MUTABLE  
IMMUTABLE ▶▶
```

Usage

By default, users can reset the values of some environment variables during a session to support their usage requirements. This capability can sometimes conflict, however, with the resource usage policy set by the DBA. By changing the mutability property of environment variables, a DBA can have more accurate control over resource allocation during a user session.

- SET ALL_MUTABLES TO MUTABLE causes all environment settings that have the mutability property to be modifiable during a user session.
- SET ALL_MUTABLES TO IMMUTABLE prevents modification of any environment settings that have the mutability property during a user session.

Settings of the following features frequently have the most impact on resource utilization within a user session, and have the mutability property defined:

- BOUND_IMPL_PDQ
- CLIENT_TZ
- COMPUTE_QUOTA
- DEFAULT TABLE_SPACE
- DEFAULT TABLE_TYPE
- IMPLICIT_PDQ
- MAXSCAN
- PDQPRIORITY
- TEMP_TAB_EXT_SIZE
- TEMP_TAB_NEXT_SIZE
- TMPSPACE_LIMIT

Users can reset the mutability of specific features by using the corresponding SET statement. For example, PDQPRIORITY can be made immutable by the SET PDQPRIORITY IMMUTABLE statement. SET DEFAULT TABLE_TYPE MUTABLE can change the default table type. SET ALL_MUTABLES can reset the mutability all of the above features in a single statement.

Typically, the mutability property is set by the DBA in a **sysdbopen** procedure to control the usage of resources. For a discussion of setting the mutability property in **sysdbopen** procedure, see *IBM Informix Performance Guide*.

The DBA can modify all variables, irrespective of their mutability setting. The command

```
onstat -g mut session_id
```

displays the current mutability settings for the specified session.

The command

```
onmode -q
```

can change the mutability of a specified variable associated with an existing session. See the *IBM Informix Administrator's Reference* for details of how to use the **onstat** and **onmode** utilities.

By default, all environment variables are MUTABLE.

Related Information

Related statements: "SET Default Table Space " on page 2-601, "SET Default Table Type " on page 2-602, "SET ENVIRONMENT " on page 2-615, and "SET PDQPRIORITY " on page 2-643

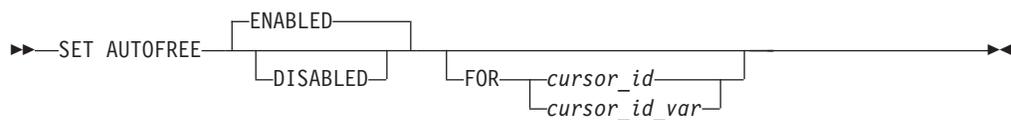
See also the section " Setting Environment Variables in SYSTEM Commands" on page 3-46, which describes how to use the "SYSTEM " on page 3-44 statement of SPL to change the settings of environment variables for the current user session.

SET AUTOFREE

Use the SET AUTOFREE statement to instruct the database server to enable or disable a memory-management feature that can free the memory allocated for a cursor automatically, as soon as the cursor is closed.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL. You can use this statement only with Informix ESQL/C.

Syntax



Element	Description	Restrictions	Syntax
<i>cursor_id</i>	Name of a cursor for which Autofree is to be reset	Must already be declared within the program	"Identifier " on page 5-21
<i>cursor_id_var</i>	Host variable that holds the value of <i>cursor_id</i>	Must store a <i>cursor_id</i> already declared in the program	Must conform to language-specific rules for names.

Usage

When the Autofree feature is enabled for a cursor, and the cursor is subsequently closed, you do not need to explicitly use the FREE statement to release the memory that the database server allocated for the cursor. If you issue SET AUTOFREE but specify no option, the default is ENABLED.

The SET AUTOFREE statement that enables the Autofree feature must appear before the OPEN statement that opens a cursor. The SET AUTOFREE statement does not affect the memory allocated to a cursor that is already open. After a cursor is Autofree enabled, you cannot open that cursor a second time.

Globally Affecting Cursors with SET AUTOFREE

If you include no FOR *cursor_id* or FOR *cursor_id_var* clause, then the scope of SET AUTOFREE is all subsequently-declared cursors in the program (or more precisely, all cursors declared before a subsequent SET AUTOFREE statement with no FOR clause globally resets the Autofree feature). This example enables the Autofree feature for all subsequent cursors in the program:

```
EXEC SQL set autofree;
```

The next example disables the Autofree feature for all subsequent cursors:

```
EXEC SQL set autofree disabled;
```

Using the FOR Clause to Specify a Specific Cursor

If you specify FOR *cursor_id* or FOR *cursor_id_var*, then SET AUTOFREE affects only the cursor that you specify after the FOR keyword.

This option allows you to override a global setting for all cursors. For example, if you issue a SET AUTOFREE ENABLED statement for all cursors in a program, you can issue a subsequent SET AUTOFREE DISABLED FOR statement to disable the Autofree feature for a specific cursor.

In the following example, the first statement enables the Autofree feature for all cursors, while the second statement disables the Autofree feature for the cursor named **x1**:

```
EXEC SQL set autofree enabled;  
EXEC SQL set autofree disabled for x1;
```

Here the **x1** cursor must have been declared but not yet opened.

Associated and Detached Statements

When a cursor is automatically freed, its associated prepared statement (or associated statement) is also freed.

The term *associated statement* has a special meaning in the context of the Autofree feature. A cursor is associated with a prepared statement if it is the first cursor that you declare with the prepared statement, or if it is the first cursor that you declare with the statement after the statement is detached.

The term *detached statement* has a special meaning in the context of the Autofree feature. A prepared statement is detached if you do not declare a cursor with the statement, or if the cursor with which the statement is associated was freed.

If the Autofree feature is enabled for a cursor that has an associated prepared statement, and that cursor is closed, the database server frees the memory allocated to the prepared statement as well as the memory allocated for the cursor. Suppose that you enable the Autofree feature for the following cursor:

```
/*Cursor associated with a prepared statement */
EXEC SQL prepare sel_stmt 'select * from customer';
EXEC SQL declare sel_curs2 cursor for sel_stmt;
```

When the database server closes the **sel_curs2** cursor, it automatically performs the equivalent of the following FREE statements:

```
FREE sel_curs2;
FREE sel_stmt;
```

Because memory for the **sel_stmt** statement is freed automatically, you cannot declare a new cursor on it unless you prepare the statement again.

Closing Cursors Implicitly

A potential problem exists with cursors that have the Autofree feature enabled. In a database that is not ANSI-compliant, if you do not close a cursor explicitly and then open it again, the cursor is closed implicitly. This implicit closing of the cursor triggers the Autofree feature. The second time the cursor is opened, the database server generates an error message (cursor not found) because the cursor is already freed.

Related Information

Related statements: "CLOSE " on page 2-77, "DECLARE " on page 2-286, "FETCH " on page 2-367, "FREE " on page 2-379, "OPEN " on page 2-459, and "PREPARE " on page 2-468

For more information on the Autofree feature, see the *IBM Informix ESQL/C Programmer's Manual*.

SET COLLATION

Use the SET COLLATION statement to specify a new collating order for the session, superseding the collation implied by the **DB_LOCALE** environment variable setting. SET NO COLLATION restores the default collation.

This statement is an extension to the ANSI/ISO standard for SQL. You can use this statement with Informix ESQL/C.

Syntax

```
►► SET { COLLATION locale | NO COLLATION } ►►
```

Element	Description	Restrictions	Syntax
<i>locale</i>	Name of a locale whose collating order is to be used in this session	Must be the name of a locale that the database server can access	"Quoted String" on page 4-189

Usage

As the *IBM Informix GLS User's Guide* explains, the database server uses locale files to specify the character set, the collating order, and other conventions of some natural language to display and manipulate character strings and other data values. The *collating order* of the database locale is the sequential order in which the database server sorts character strings.

If you set no value for **DB_LOCALE**, the default locale, based on United States English, is **en_us.8859-1** for UNIX, or Code Page 1252 for Windows systems. Otherwise, the database server uses the **DB_LOCALE** setting as its locale. **SET COLLATION** overrides the collating order of **DB_LOCALE** at runtime for all database servers previously accessed in the same session.

The new collating order remains in effect for the rest of the session, or until you issue another **SET COLLATION** statement. Other sessions are not affected, but database objects that you created with a non-default collation use whatever collating order was in effect at their time of their creation.

By default, the collating order is the code-set order, but some locales also support a *locale-specific* order. In most contexts, only NCHAR and NVARCHAR data values can be sorted according to a locale-specific collating order.

Specifying a Collating Order with SET COLLATION

SET COLLATION replaces the current collating order with that of the specified *locale* for all database servers previously accessed in the current session. For example, this specifies the collating order of the German language:

```
EXEC SQL set collation "de_de.8859-1";
```

If the next action of a database server in this session sorted NCHAR or NVARCHAR values, this would follow the German collating order.

Suppose that, in the same session, the following **SET NO COLLATION** statement restores the **DB_LOCALE** setting for the collating order:

```
EXEC SQL set no collation;
```

After **SET NO COLLATION** executes, subsequent collation in the same session is based on the **DB_LOCALE** setting. Any database objects that you created using the German collating order, however, such as check constraints, indexes, prepared objects, triggers, or UDRs, will continue to apply German collation to NCHAR and NVARCHAR data types.

Restrictions on SET COLLATION

Although **SET COLLATION** enables you to change the collating order of the database server dynamically within a session, you should be aware of these limitations on the effects of the **SET COLLATION** statement.

- Only collation performed by the database server is affected. Client processes that sort data are not affected by **SET COLLATION**.
- Only the current session is affected. Other sessions are not affected directly by your **SET COLLATION** statements (but any database objects that you create will sort in their creation-time collating order).
- Changing the collating order does not change the code set. The database server always uses the code set specified by **DB_LOCALE**.
- Only NCHAR and NVARCHAR values sort in locale-specific order.

Because SET COLLATION changes only the collating order, rather than the current locale or code set, you cannot use this statement to insert character data from different locales, such as French and Japanese, into the same database. You must use Unicode if the database needs to store characters from two or more languages that require inherently different code sets. The database can store characters from the dissimilar character sets of more than one natural language only if you set **DB_LOCALE** to a Unicode locale when the database is created.

Collation Performed by Database Objects

Although the database reverts to the **DB_LOCALE** collating order after the session ends (or after you execute SET NO COLLATION), objects that you create using a non-default collation persist in the database. You can create, for example, multiple indexes on the same set of columns, called *multilingual indexes*, using different collating orders that SET COLLATION specifies.

Only one clustered index, however, can exist on a given set of columns.

Only one unique constraint or primary key can exist on a given set of columns, but you can create multiple unique indexes on the same set of columns, if each index has a different collation order.

The query optimizer ignores indexes that apply any collation other than the current session collation to NCHAR or NVARCHAR columns when calculating the cost of a query.

The collating order of an attached index must be the same as that of its table, and this must be the default collating order specified by **DB_LOCALE**.

The ALTER INDEX statement cannot change the collation of an index. Any previous SET COLLATION statement is ignored when ALTER INDEX executes.

When you compare values from CHAR columns with NCHAR columns, Dynamic Server casts the CHAR value to NCHAR, and then applies the current collation. Similarly, before comparing VARCHAR and NVARCHAR values, Dynamic Server first casts the VARCHAR values to NVARCHAR.

When synonyms are created for remote tables or views, the participating databases must have the same collating order. Existing synonyms, however, can be used in other databases that support SET COLLATION and the collating order of the synonym, regardless of the **DB_LOCALE** setting.

Check constraints, cursors, prepared objects, triggers, and SPL routines that sort NCHAR or NVARCHAR values use the collation that was in effect at the time of their creation, if this is different from the **DB_LOCALE** setting.

The effect on performance is sensitive to how many different collations are used when creating database objects that sort in a localized order.

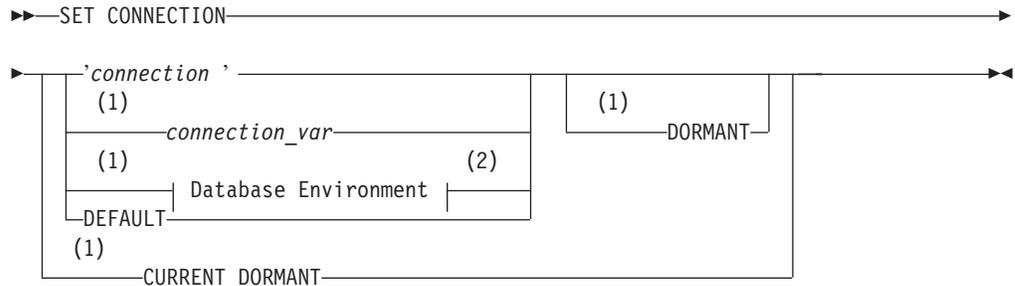
Related Information

For information on locales, see the *IBM Informix GLS User's Guide*.

SET CONNECTION

Use the SET CONNECTION statement to reestablish a connection between an application and a database environment and to make the connection current. You can also use this statement with the DORMANT option to put the current connection in a dormant state. Use this statement with Informix ESQL/C.

Syntax



Notes:

- 1 Informix extension
- 2 See "Database Environment" on page 2-84

Element	Description	Restrictions	Syntax
<i>connection</i>	Name of the initial connection that the CONNECT statement made	The database must already exist	"Quoted String" on page 4-189
<i>connection_var</i>	Host variable that contains the value of <i>connection</i>	Must be a character data type	Language specific

Usage

You can use SET CONNECTION to make a dormant connection the current connection or to make the current connection dormant.

SET CONNECTION is not valid as a prepared statement.

Making a Dormant Connection the Current Connection

If you use the SET CONNECTION statement without the DORMANT option, *connection* must represent a dormant connection. A *dormant connection* is a connection that is established but is not current.

The SET CONNECTION statement, with no DORMANT option, makes the specified dormant connection the current one. The connection that the application specifies must be dormant. The connection that is current when the statement executes becomes dormant.

The SET CONNECTION statement in the following example makes connection con1 the current connection and makes con2 a dormant connection:

```
CONNECT TO 'stores_demo' AS 'con1';  
...  
CONNECT TO 'demo' AS 'con2';  
...  
SET CONNECTION 'con1';
```

A dormant connection has a *connection context* associated with it. When an application makes a dormant connection current, it reestablishes that connection to a database environment and restores its connection context. (For more information on connection context, see the “CONNECT ” on page 2-83 statement on page “CONNECT ” on page 2-83.) Reestablishing a connection is comparable to establishing the initial connection, except that it typically avoids authenticating the permissions for the user again, and it avoids reallocating resources associated with the initial connection. For example, the application does not need to reprepare any statements that have previously been prepared in the connection, nor does it need to redeclare any cursors.

Making a Current Connection the Dormant Connection

In the SET CONNECTION *connection* DORMANT statement, *connection* must represent the current connection. The SET CONNECTION statement with the DORMANT option makes the specified current connection a dormant connection. For example, the following SET CONNECTION statement makes connection con1 dormant:

```
SET CONNECTION 'con1' DORMANT;
```

The SET CONNECTION statement with the DORMANT option generates an error if you specify a connection that is already dormant. For example, if connection con1 is current and connection con2 is dormant, the following SET CONNECTION statement returns an error message:

```
SET CONNECTION 'con2' DORMANT;
```

The following SET CONNECTION statement executes successfully:

```
SET CONNECTION 'con1' DORMANT;
```

Dormant Connections in a Single-Threaded Environment

In a single-threaded Informix ESQL/C application (one that does not use threads), the DORMANT option makes the current connection dormant. Using this option makes single-threaded Informix ESQL/C applications upwardly compatible with thread-safe Informix ESQL/C applications. A single-threaded environment, however, can have only one active connection while the program executes.

Dormant Connections in a Thread-Safe Environment

In a thread-safe Informix ESQL/C application, the DORMANT option makes an active connection dormant. Another thread can now use the connection by issuing the SET CONNECTION statement without the DORMANT option. A thread-safe environment can have many threads (concurrent pieces of work performing particular tasks) in one Informix ESQL/C application, and each thread can have one active connection.

An active connection is associated with a particular thread. Two threads cannot share the same active connection. Once a thread makes an active connection dormant, that connection is available to other threads. A dormant connection is still established but is not currently associated with any thread. For example, if the

connection named con1 is active in the thread named thread_1, the thread named thread_2 cannot make connection con1 its active connection until thread_1 has made connection con1 dormant.

The following code fragment from a thread-safe Informix ESQL/C program shows how a particular thread within a thread-safe application makes a connection active, performs work on a table through this connection, and then makes the connection dormant so that other threads can use the connection:

```
thread_2()
{ /* Make con2 an active connection */
  EXEC SQL connect to 'db2' as 'con2';
  /*Do insert on table t2 in db2*/
  EXEC SQL insert into table t2 values(10);
  /* make con2 available to other threads */
  EXEC SQL set connection 'con2' dormant;
}
```

If a connection to a database environment was initiated using the CONNECT . . . WITH CONCURRENT TRANSACTION statement, any thread that subsequently connects to that database environment can use an ongoing transaction. In addition, if an open cursor is associated with such a connection, the cursor remains open when the connection is made dormant.

Threads within a thread-safe Informix ESQL/C application can use the same cursor by making the associated connection current, even though only one thread can use the connection at any given time.

Identifying the Connection

If the application did not specify a connection name in the initial CONNECT statement, you must use a database environment (such as a database name or a database pathname) as the connection name. For example, the following SET CONNECTION statement uses a database environment for the connection name because the CONNECT statement does not use a connection name. For information about quoted strings that specify a database environment, see “Database Environment” on page 2-84.

```
CONNECT TO 'stores_demo';
...
CONNECT TO 'demo';
...
SET CONNECTION 'stores_demo';
```

If a connection to a database server was assigned a connection name, however, you must use the connection name to reconnect to the database server. An error is returned if you use a database environment rather than the connection name when a connection name exists.

DEFAULT Option

The DEFAULT option specifies the default connection for a SET CONNECTION statement. The default connection is one of the following connections:

- An explicit default connection (a connection established with the CONNECT TO DEFAULT statement)
- An implicit default connection (any connection established with the DATABASE or CREATE DATABASE statements)

Use SET CONNECTION without a DORMANT option to reestablish the default connection, or with that option to make the default connection dormant.

For more information, see “The DEFAULT Connection Specification” on page 2-88 and “The Implicit Connection with DATABASE Statements” on page 2-88.

CURRENT Keyword

Use the CURRENT keyword with the DORMANT option of the SET CONNECTION statement as a shorthand form of identifying the current connection. The CURRENT keyword replaces the current connection name. If the current connection is con1, the following two statements are equivalent:

```
SET CONNECTION 'con1' DORMANT;
```

```
SET CONNECTION CURRENT DORMANT;
```

When a Transaction is Active

Without the DORMANT keyword, the SET CONNECTION statement implicitly puts the current connection in the dormant state.

When you issue a SET CONNECTION statement with the DORMANT keyword, the SET CONNECTION statement explicitly puts the current connection in the dormant state. In both cases, the statement can fail if a connection that becomes dormant has an uncommitted transaction. If the connection that becomes dormant has an uncommitted transaction, the following conditions apply:

- If the connection was established using the WITH CONCURRENT TRANSACTION clause of the CONNECT statement, SET CONNECTION succeeds and puts the connection in a dormant state.
- If the connection was not established by the WITH CONCURRENT TRANSACTION clause of the CONNECT statement, SET CONNECTION fails and cannot set the connection to a dormant state, and the transaction in the current connection continues to be active. The statement generates an error and the application must decide whether to commit or roll back the active transaction.

Related Information

Related statements: “CONNECT ” on page 2-83, “DISCONNECT ” on page 2-320, and “DATABASE ” on page 2-281

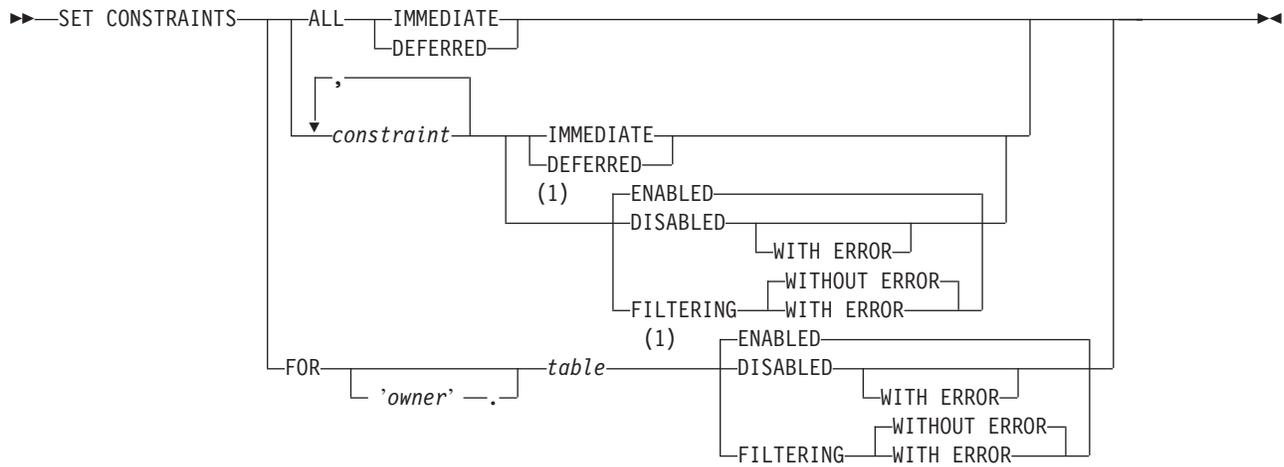
For a discussion of the SET CONNECTION statement and thread-safe applications, see the *IBM Informix ESQL/C Programmer's Manual*.

SET CONSTRAINTS

Use the SET CONSTRAINTS statements to specify how some or all of the constraints on a table are processed. Constraint-mode options include these:

- Whether constraints are checked at the statement level (IMMEDIATE) or at the transaction level (DEFERRED)
- Whether to enable or disable constraints
- Whether to change the filtering mode of constraints.

Syntax



Notes:

1 Informix extension and Dynamic Server only

Element	Description	Restrictions	Syntax
<i>constraint</i>	Constraint whose mode is to be reset	All constraints must exist and must all be defined on the same table	"Identifier " on page 5-21
<i>owner</i>	Owner of <i>table</i>	Must own <i>table</i>	" Owner Name " on page 5-45
<i>table</i>	Table whose constraint mode is to be reset for all constraints	Table must exist in the database	"Identifier " on page 5-21

Usage

The SET CONSTRAINTS keywords begin the SET Transaction Mode statement, which is described in "SET Transaction Mode " on page 2-661. Only Dynamic Server supports the SET Transaction Mode statement, which is an extension to the ANSI/ISO standard for SQL.

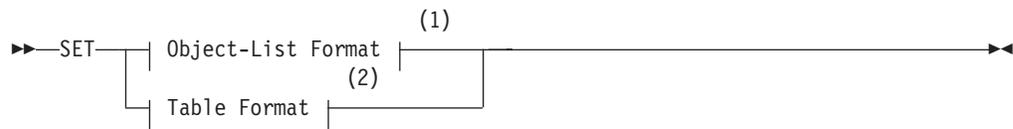
The SET CONSTRAINTS keywords can also begin a special case of the SET Database Object Mode statement, which is an extension to the ANSI/ISO standard for SQL. The SET Database Object Mode statement can also enable or disable a trigger or index, or change the filtering mode of a unique index. For the complete syntax and semantics of the SET INDEX statement, see "SET Database Object Mode."

SET Database Object Mode

Use the SET Database Object Mode statement to change the filtering mode of constraints of unique indexes, or to enable or disable constraints, indexes, and triggers.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL. To specify whether constraints are checked at the statement level or at the transaction level, see "SET Transaction Mode " on page 2-661.

Syntax



Notes:

- 1 See “ Object-List Format” on page 2-593
- 2 See “ Table Format” on page 2-593

Usage

In the context of this statement, *database object* has the restricted meaning of an *index*, a *trigger*, or a *constraint*, rather than the more general meaning of this term that the description of the “Database Object Name ” on page 5-16 segment defines in Chapter 5, “Other Syntax Segments,” on page 5-1.

The scope of the SET Database Object Mode statement is restricted to constraints, indexes, or triggers in the local database to which the session is currently connected. After you change the mode of an object, the new mode is in effect for all sessions of that database, and persists until another SET Database Object Mode statement changes it again, or until the object is dropped from the database.

Only two object modes are available for triggers and for indexes that allow duplicate values:

- enabled
- disabled

For constraints and for unique indexes, you can also specify two additional modes:

- filtering without integrity-violation errors
- filtering with integrity-violation errors

At any given time, an object must be in exactly one of these modes. These modes, which are sometimes called *object states*, are described in the section “Definitions of Database Object Modes” on page 2-595.

The **sysobjstate** system catalog table lists all of the constraint, index, and trigger objects in the database, and the current mode of each object. For information on the **sysobjstate** table, see the *IBM Informix Guide to SQL: Reference*.

Privileges Required for Changing Database Object Modes

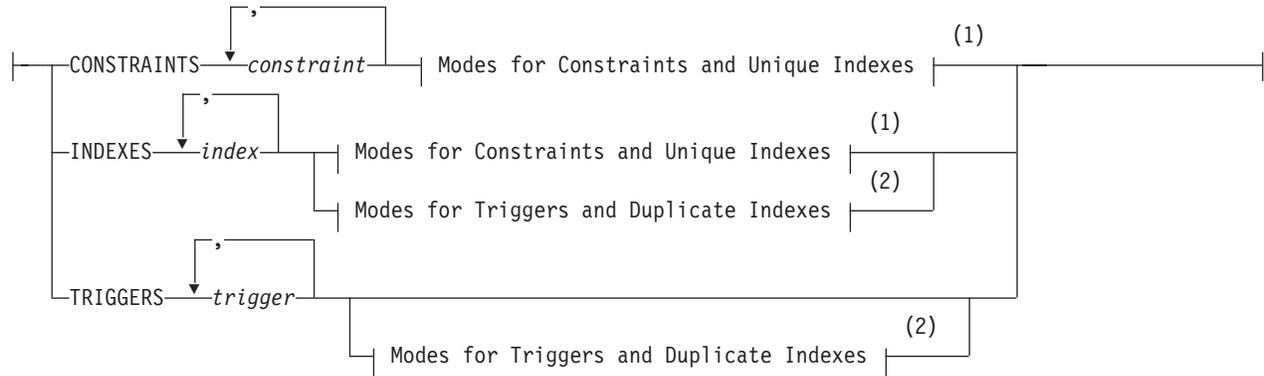
To change the mode of a constraint, index, or trigger, you must have the necessary access privileges. You must meet at least one of these requirements:

- You must have the DBA privilege on the database.
- You must be the owner of the table on which the database object is defined and you must also have the Resource privilege on the database.
- You must have the Alter privilege on the table on which the database object is defined and you must also have the Resource privilege on the database.

Object-List Format

Use the object-list format to change the mode for one or more constraint, index, or trigger.

Object-List Format:



Notes:

- 1 See "Modes for Constraints and Unique Indexes" on page 2-594
- 2 See "Modes for Triggers and Duplicate Indexes" on page 2-597

Element	Description	Restrictions	Syntax
<i>constraint</i>	Name of a constraint whose mode is to be set	Must be a local constraint, and all constraints in the list must be defined on the same table	"Identifier " on page 5-21
<i>index</i>	Name of an index whose mode is to be set	Must be a local index, and all indexes in the list must be defined on the same table	"Identifier " on page 5-21
<i>trigger</i>	Name of a trigger whose mode is to be set	Must be a local trigger, and all triggers in the list must be defined on the same table or view	"Identifier " on page 5-21

For example, to change the mode of the unique index **unq_ssn** on the **cust_subset** table to filtering, enter the following statement:

```
SET INDEXES unq_ssn FILTERING;
```

You can also use the object-list format to change the mode for a list of constraints, indexes, or triggers that are defined on the same table. Assume that four triggers are defined on the **cust_subset** table: **insert_trig**, **update_trig**, **delete_trig**, and **execute_trig**. Also assume that all four triggers are enabled. To disable all triggers except **execute_trig**, enter this statement:

```
SET TRIGGERS insert_trig, update_trig, delete_trig DISABLED;
```

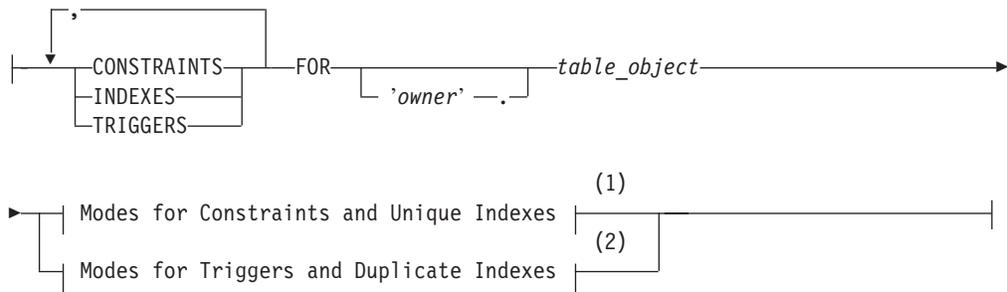
If **my_trig** is a disabled INSTEAD OF trigger on a view, the following statement enables that trigger:

```
SET TRIGGERS my_trig ENABLED;
```

Table Format

Use the table format to change the mode of all database objects of a specified type that have been defined on the same table or view.

Table Format:



Notes:

- 1 See “ Modes for Constraints and Unique Indexes”
- 2 See “ Modes for Triggers and Duplicate Indexes” on page 2-597

Element	Description	Restrictions	Syntax
<i>owner</i>	Owner of <i>table</i>	Must own <i>table</i>	“ Owner Name ” on page 5-45
<i>table_object</i>	Table or view on which objects are defined	Must be a local table or view. Objects defined on a temporary table cannot be set to disabled or filtering modes.	“Identifier ” on page 5-21

This example disables all constraints defined on the **cust_subset** table:

```
SET CONSTRAINTS FOR cust_subset DISABLED;
```

In table format, you can change the modes of more than one database object type with a single statement. For example, this enables all constraints, indexes, and triggers that are defined on the **cust_subset** table:

```
SET CONSTRAINTS, INDEXES, TRIGGERS FOR cust_subset ENABLED;
```

In Dynamic Server 10.00 and in earlier versions, you cannot use the SET TRIGGERS option of the SET Database Object Mode statement to disable an inherited trigger selectively within a table hierarchy. In this release, however, disabling a trigger on a table within a hierarchy has no effect on inherited triggers. For example, the following statement disables all triggers on the specified *subtable*, but the statement has no effect on triggers on table objects that are above or below *subtable* within a table hierarchy:

```
SET TRIGGERS FOR subtable DISABLED
```

Modes for Constraints and Unique Indexes

You can specify enabled or disabled mode for a constraint or a unique index.

Modes for Constraints and Unique Indexes:



If you specify no mode for a constraint in a CREATE TABLE, ALTER TABLE, or SET Database Object Mode statement, the constraint is enabled by default. If you do not specify the mode for an index in the CREATE INDEX or SET Database Object Mode statement, the index is enabled by default.

Definitions of Database Object Modes

You can use database object modes to control the effects of INSERT, DELETE, and UPDATE statements. Your choice of mode affects the tables whose data you are manipulating, the behavior of the database objects defined on those tables, and the behavior of the data manipulation statements themselves.

Enabled Mode

Constraints, indexes, and triggers are enabled by default. The CREATE TABLE, ALTER TABLE, CREATE INDEX, and CREATE TRIGGER statements create database objects in enabled mode, unless you specify another mode. When a database object is enabled, the database server recognizes the existence of the database object and takes the database object into consideration while it executes an INSERT, DELETE, or UPDATE statement (or for Select triggers, a SELECT statement). Thus, an enabled constraint is enforced, an enabled index updated, and an enabled trigger is executed when the trigger event takes place.

When you enable constraints and unique indexes, if a violating row exists, the data manipulation statement fails (that is, no rows are changed) and the database server returns an error message.

Disabled Mode

When a database object is disabled, the database server ignores it during the execution of an INSERT, DELETE, SELECT, or UPDATE statement. A disabled constraint is not enforced, a disabled index is not updated, and a disabled trigger is not executed when the trigger event takes place. When you disable constraints and unique indexes, any data manipulation statement that violates the restriction of the constraint or unique index succeeds (that is, the target row is changed), and the database server does not return an error message.

You can use the disabled mode to add a new constraint or new unique index to an existing table, even if some rows in the table do not satisfy the new integrity specification. Disabling can also be efficient in LOAD operations.

For information on adding a constraint, see “Adding a Constraint That Existing Rows Violate (IDS)” on page 2-61 in the ALTER TABLE statement. For information on adding a unique index, see “Adding a Unique Index When Duplicate Values Exist in the Column” on page 2-140 in the CREATE INDEX statement.

Filtering Mode

When a constraint or unique index is in filtering mode, the INSERT, DELETE, or UPDATE statement succeeds, but the database server enforces the constraint or the unique-index requirement by writing any failed rows to the violations table associated with the target table. Diagnostic information about the constraint violation or unique-index violation is written to the diagnostics table associated with the target table.

In data manipulation operations, filtering mode has the following specific effects on INSERT, UPDATE, and DELETE statements:

- A constraint violation during an INSERT statement causes the database server to make a copy of the nonconforming record and write it to the violations table. The database server does not write the nonconforming record to the target table. If the INSERT statement is not a singleton INSERT, the rest of the insert operation proceeds with the next record.
- A constraint violation or unique-index violation during an UPDATE statement causes the database server to make a copy of the existing record that was to be updated and write it to the violations table. The database server also makes a copy of the new record and writes it to the violations table, but the actual record is not updated in the target table. If the UPDATE statement is not a singleton update, the rest of the update operation proceeds with the next record.
- A constraint violation or unique-index violation during a DELETE statement causes the database server to make a copy of the record that was to be deleted and write it to the violations table. The database server does not delete the actual record in the target table. If the DELETE statement is not a singleton delete, the rest of the delete operation proceeds with the next record.

In all of these cases, the database server sends diagnostic information about each constraint violation or unique-index violation to the diagnostics table associated with the target table.

For information on the structure of the records that the database server writes to the violations and diagnostics tables, see “Structure of the Violations Table” on page 2-667 and “Structure of the Diagnostics Table (IDS)” on page 2-672.

Starting and Stopping the Violations and Diagnostics Tables:

You must use the START VIOLATIONS TABLE statement to start the violations and diagnostics tables for the target table on which the database objects are defined, either before you set any database objects that are defined on the table to the filtering mode, or after you set database objects to filtering, but before any users issue INSERT, DELETE, or UPDATE statements.

If you want to stop the database server from filtering bad records to the violations table and sending diagnostic information about each bad record to the diagnostics table, you must issue a STOP VIOLATIONS TABLE statement.

For further information on these statements, see “START VIOLATIONS TABLE” on page 2-663 and “STOP VIOLATIONS TABLE ” on page 2-676.

Error Options for Filtering Mode: When you set the mode of a constraint or unique index to filtering, you can specify one of two error options. These error options control whether the database server displays an integrity-violation error message when it encounters bad records during execution of data manipulation statements:

- The WITH ERROR option instructs the database server to return a referential integrity-violation error message after executing an INSERT, DELETE, or UPDATE statement in which one or more of the target rows causes a constraint violation or a unique-index violation.
- The WITHOUT ERROR option is the default. This option prevents the database server from issuing a referential integrity-violation error message to the user after an INSERT, DELETE, or UPDATE statement causes a constraint violation or a unique-index violation.

Effect of Filtering Mode on the Database: The net effect of the filtering mode is that the contents of the target table always satisfy all constraints on the table and any unique-index requirements on the table.

In addition, the database server does not lose any data values that violate a constraint or unique-index requirement, because non-conforming records are sent to the violations table, and diagnostic information about those records is sent to the diagnostics table.

Furthermore, when filtering mode is in effect, insert, delete, and update operations on the target table do not fail when the database server encounters bad records. These operations succeed in adding all the good records to the target table. Thus, filtering mode is appropriate for large-scale batch updates of tables. The user can fix records that violate constraints and unique-index requirements after the fact. The user does not need to fix the bad records before the batch update to avoid losing the bad records during the batch update.

Modes for Triggers and Duplicate Indexes

You can specify the modes for triggers or duplicate indexes.

Modes for Triggers and Duplicate Indexes:



If you specify no mode for an index or for a trigger when you create it or in a subsequent SET Database Object Mode statement, the object is enabled by default.

Examples

```
SET CONSTRAINTS u100_1 ENABLED;
```

```
SET CONSTRAINTS r104_11 FILTERING WITH ERROR;
```

```
SET CONSTRAINTS FOR orders DISABLED;
```

Related Information

Related statements: “ALTER TABLE” on page 2-44, “CREATE TABLE ” on page 2-191, “CREATE INDEX ” on page 2-124, “CREATE TRIGGER” on page 2-239, “START VIOLATIONS TABLE” on page 2-663, and “STOP VIOLATIONS TABLE ” on page 2-676

For a discussion of object modes and violation detection and examples that show how database object modes work when users execute data manipulation statements on target tables or add new constraints and indexes to target tables, see the *IBM Informix Guide to SQL: Tutorial*.

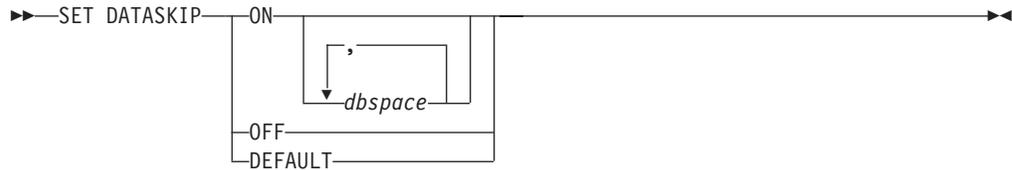
For information on the system catalog tables associated with the SET Database Object Mode statement, see the descriptions of the **sysobjstate** and **sysviolations** tables in the *IBM Informix Guide to SQL: Reference*.

SET DATASKIP

Use the SET DATASKIP statement to control whether the database server skips a dbspace that is unavailable during the processing of a transaction.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax



Element	Description	Restrictions	Syntax
<i>dbspace</i>	Name of the skipped dbspace	Must exist at time of execution	"Identifier " on page 5-21

Usage

SET DATASKIP allows you to reset at runtime the Dataskip feature, which controls whether the database server skips a dbspace that is unavailable (for example, due to a media failure) in the course of processing a transaction.

In Informix ESQL/C, the warning flag `sqlca.sqlwarn.sqlwarn6` is set to W if a dbspace is skipped. See also the *IBM Informix ESQL/C Programmer's Manual*.

In Dynamic Server, this statement applies only to tables that are fragmented across dbspaces or partitions. It does not apply to blobspaces nor to sbspaces.

Specifying SET DATASKIP ON without including a *dbspace* instructs the database server to skip any dbspaces in the fragmentation list that are unavailable. You can use the `onstat -d` or `-D` options to determine whether a dbspace is down.

When you specify SET DATASKIP ON *dbspace*, you are instructing the database server to skip the specified *dbspace* if it is unavailable.

If you specify SET DATASKIP OFF, the Dataskip feature is disabled. If you specify SET DATASKIP DEFAULT, the database server uses the setting for the Dataskip feature from the ONCONFIG file.

Circumstances When a Dbspace Cannot Be Skipped

The database server cannot skip a dbspace under certain conditions. The following list outlines those conditions:

- Referential constraint checking

When you want to delete a parent row, the child rows must also be available for deletion, and must exist in an available fragment.

When you want to insert a new child row, the parent row must be found in the available fragments.

- Updates

When you perform an update that moves a record from one fragment to another, both fragments must be available.

- Inserts

When you try to insert records in a expression-based fragmentation strategy and the dbspace is unavailable, an error is returned.

When you try to insert records in a round-robin fragment-based strategy, and a dbspace is down, the database server inserts the rows into any available dbspace.

When no dbspace is available, an error is returned.

- Indexing

When you perform updates that affect the index, such as when you insert or delete rows, or update an indexed column, the index must be available.

When you try to create an index, the dbspace you want to use must be available.

- Serial keys

The first fragment is used to store the current serial-key value internally. This is not visible to you except when the first fragment becomes unavailable and a new serial key value is required, which can happen during INSERT statements.

Examples

The following skips **dbsp1** for the current session:

```
SET DATASKIP ON dbsp1;
```

The following sets the value of **DATASKIP** to the value specified in **onconfig**:

```
SET DATASKIP DEFAULT;
```

The following switches **DATASKIP** off so that all dbspaces are used.

```
SET DATASKIP OFF;
```

Related Information

Related statement: “ SET ALL_MUTABLES” on page 2-581

For additional information about the Dataskip feature, see your *IBM Informix Administrator's Guide*.

SET DEBUG FILE

Use the SET DEBUG FILE statement to identify the file that is to receive the runtime trace output of an SPL routine.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax

```
▶▶—SET DEBUG FILE TO—'filename'—  
                        └──filename_var──┘ └──WITH APPEND──┘  
                        └──expression──┘
```

Element	Description	Restrictions	Syntax
<i>expression</i>	Expression that returns a filename	Must be a valid filename	"Expression" on page 4-40
<i>filename</i>	Pathname of the file that contains the output of the TRACE statement	See "Using the WITH APPEND Option"	"Quoted String" on page 4-189.
<i>filename_var</i>	Host variable storing <i>filename</i> string	Must be a character type	Language specific

Usage

This statement specifies the file to which the database server writes the output from the TRACE statement in the SPL routine. Each time the TRACE statement is executed, the trace information is added to this output file.

Using the WITH APPEND Option

The output file that you specify in the SET DEBUG FILE statement can be a new file or existing file. If you specify an existing file, its current contents are deleted when you issue the SET DEBUG FILE TO statement. The first execution of a TRACE statement sends trace output to the beginning of the file.

If you include the WITH APPEND option, the current contents of the file are preserved when you issue the SET DEBUG FILE statement. The first execution of a TRACE statement adds the new trace output to the end of the file.

If you specify a new file in the SET DEBUG FILE TO statement, it makes no difference whether you include the WITH APPEND option. The first execution of a TRACE statement sends trace output to the beginning of the new file whether you include or omit the WITH APPEND option.

Closing the Output File

To close the file that the SET DEBUG FILE TO statement opened, issue another SET DEBUG FILE TO statement with another filename. You can then read or edit the contents of the first file.

Redirecting Trace Output

You can use the SET DEBUG FILE TO statement outside an SPL routine to direct the trace output of the SPL routine to a file. You can also use this statement within an SPL routine to redirect its own output.

Location of the Output File

If you execute the SET DEBUG FILE statement with a simple filename on a local database, the output file is located in your current directory. If your current database is on a remote database server, the output file is located in your home directory on the remote database server. If you provide a full pathname for the debug file, the file is placed in the directory that you specify on the remote database server. If you do not have write permissions in the directory, you receive an error.

The following example sends the output of the SET DEBUG FILE TO statement to a file called **debug.out**:

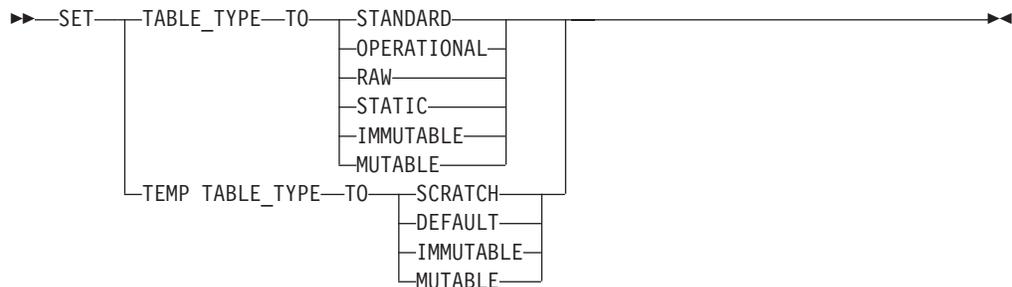
```
SET DEBUG FILE TO 'debug' || '.out';
```

SET Default Table Type

Use the SET Default Table Type statement to specify the default table type for tables (or temporary tables) that you subsequently create in the same session.

Only Extended Parallel Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Usage

If CREATE TABLE specifies no table type, the default table type is STANDARD. The SET TABLE_TYPE statement can change this default for subsequent CREATE TABLE statements in the current session. Similarly, you can use the SET TEMP TABLE_TYPE to change the default temporary table type.

These statements have no effect on tables for which you explicitly specify a table type in the statement that creates the new table or temporary table.

Because the CREATE Temporary TABLE statement requires an explicit table type, the SET TEMP TABLE_TYPE statement only affects SQL operations that create a temporary implicitly, such as in executing join operations, SELECT statements with the GROUP BY or ORDER BY clause, and index builds.

The effect of SET Default Table Type persists until the end of the session, or until you issue this statement again to specify a new default table type.

The SET TABLE_TYPE TO STANDARD statement and the SET TEMP TABLE_TYPE TO DEFAULT statements restore the default behavior.

The DBA can use the keyword MUTABLE or IMMUTABLE to enable (with MUTABLE) or to prevent (with IMMUTABLE) changes by users to the default table type during a user session. The DBA typically sets the default table type in a **sysdbopen** procedure and specifies IMMUTABLE if it should not be changed during a user session.

Although the scope of these statements is the current session, they can be used to have a database-wide effect. The next example shows how to do this by using SPL routines to establish a default table type at connect time:

```
CREATE PROCEDURE public.sysdbopen()  
  SET TABLE_TYPE TO RAW;  
  SET TEMP TABLE_TYPE TO SCRATCH;  
  SET TABLE_SPACE TO other_tables;
```

```

...
END PROCEDURE;

CREATE PROCEDURE helene.sysdbopen()
EXECUTE PROCEDURE public.sysdbopen();
SET ROLE marketing;
SET TABLE_SPACE TO marketing_dbslice;
END PROCEDURE;

```

For information about the `sysdbopen()` procedure, see “IFX_REPLACE_MODULE Function” on page 6-54.

Related Information

Related statements: “CREATE TABLE ” on page 2-191, “CREATE TEMP TABLE ” on page 2-230, “ SET ALL_MUTABLES” on page 2-581, “SET Default Table Space ” on page 2-601.

For more information on table types that can be specified in the CREATE TABLE statement, see “CREATE TABLE ” on page 2-191. For more information about temporary tables, see “CREATE Temporary TABLE ” on page 2-231.

SET DEFERRED_PREPARE

Use the SET DEFERRED_PREPARE statement to control whether a client process postpones sending a PREPARE statement to the database server until the OPEN or EXECUTE statement is sent. Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL. You can use this statement only with Informix ESQL/C.

Syntax

```

>> SET DEFERRED_PREPARE { ENABLED | DISABLED } <<

```

Usage

By default, the SET DEFERRED_PREPARE statement causes the application program to delay sending the PREPARE statement to the database server until the OPEN or EXECUTE statement is executed. In effect, the PREPARE statement is bundled with the other statement so that one round-trip of messages, instead of two, is sent between the client and the server. This Deferred-Prepare feature can affect the following series of Dynamic SQL statement:

- PREPARE, DECLARE, OPEN statement blocks that operate with the FETCH or PUT statements
- PREPARE followed by the EXECUTE or EXECUTE IMMEDIATE statement

You can specify ENABLED or DISABLED options for SET DEFERRED_PREPARE.

If you specify no option, the default is ENABLED. The following example enables the Deferred-Prepare feature by default:

```
EXEC SQL set deferred_prepare;
```

The ENABLED option enables the Deferred-Prepare feature within the application. The following example explicitly specifies the ENABLED option:

```
EXEC SQL set deferred_prepare enabled;
```

After an application issues SET DEFERRED_PREPARE ENABLED, the Deferred-Prepare feature is enabled for subsequent PREPARE statements in the application. The application then exhibits the following behavior:

- The sequence PREPARE, DECLARE, OPEN sends the PREPARE statement to the database server with the OPEN statement. If the prepared statement has syntax errors, the database server does not return error messages to the application until the application declares a cursor for the prepared statement and opens the cursor.
- The sequence PREPARE, EXECUTE sends the PREPARE statement to the database server with the EXECUTE statement. If a prepared statement contains syntax errors, the database server does not return error messages to the application until the application attempts to execute the prepared statement.

If Deferred-Prepare is enabled in a PREPARE, DECLARE, OPEN statement block that contains a DESCRIBE statement, the DESCRIBE statement must follow the OPEN statement rather than the PREPARE statement. If the DESCRIBE follows PREPARE, the DESCRIBE statement results in an error.

Use the DISABLED option to disable the Deferred-Prepare feature within the application. The following example specifies the DISABLED option:

```
EXEC SQL set deferred_prepare disabled;
```

If you specify the DISABLED option, the application sends each PREPARE statement to the database server when the PREPARE statement is executed.

Example of SET DEFERRED_PREPARE

The following code fragment shows a SET DEFERRED_PREPARE statement with a PREPARE, EXECUTE statement block. In this case, the database server executes the PREPARE and EXECUTE statements all at once:

```
EXEC SQL BEGIN DECLARE SECTION;
    int a;
EXEC SQL END DECLARE SECTION;
EXEC SQL allocate descriptor 'desc';
EXEC SQL create database test;
EXEC SQL create table x (a int);

/* Enable Deferred-Prepare feature */
EXEC SQL set deferred_prepare enabled;
/* Prepare an INSERT statement */
EXEC SQL prepare ins_stmt from 'insert into x values(?)';
a = 2;
EXEC SQL EXECUTE ins_stmt using :a;
if (SQLCODE)
    printf("EXECUTE : SQLCODE is %d\n", SQLCODE);
```

Using Deferred-Prepare with OPTOFC

You can use the Deferred-Prepare and Open-Fetch-Close Optimization (OPTOFC) features in combination. The OPTOFC feature delays sending the OPEN message to the database server until the FETCH message is sent. The following situations occur if you enable the Deferred-Prepare and OPTOFC features at the same time:

- If the text of a prepared statement contains syntax errors, the error messages are not returned to the application until the first FETCH statement is executed.

- A DESCRIBE statement cannot be executed until after the FETCH statement.
- You must issue an ALLOCATE DESCRIPTOR statement before a DESCRIBE or GET DESCRIPTOR statement can be executed.

The database server performs an internal execution of a SET DESCRIPTOR statement which sets the TYPE, LENGTH, DATA, and other fields in the system descriptor area. You can specify a GET DESCRIPTOR statement after the FETCH statement to see the data that is returned.

Related Information

Related statements: “DECLARE ” on page 2-286, “DESCRIBE” on page 2-311, “EXECUTE ” on page 2-348, “OPEN ” on page 2-459, and “PREPARE ” on page 2-468

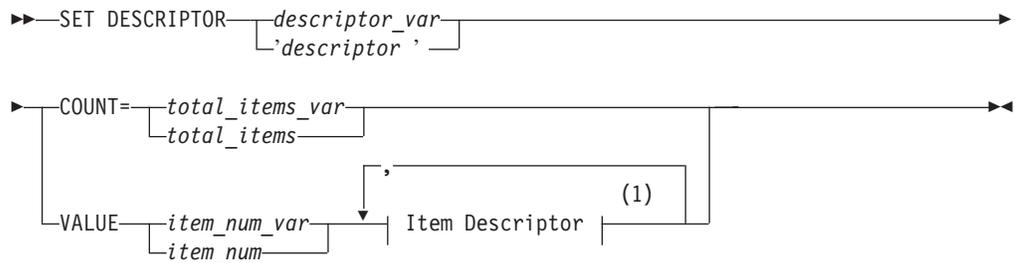
For a task-oriented discussion of the PREPARE statement and dynamic SQL, see the *IBM Informix Guide to SQL: Tutorial*.

For more information about concepts that relate to the SET DEFERRED_PREPARE statement, see the *IBM Informix ESQL/C Programmer’s Manual*.

SET DESCRIPTOR

The SET DESCRIPTOR statement sets values in a system-descriptor area (SDA). Use this statement with Informix ESQL/C.

Syntax



Notes:

- 1 See “Item Descriptor” on page 2-606

Element	Description	Restrictions	Syntax
<i>descriptor</i>	String that identifies the SDA to which values are assigned	System-descriptor area (SDA) must be previously allocated	“Quoted String” on page 4-189
<i>descriptor_var</i>	Host variable that stores <i>descriptor</i>	Same restrictions as <i>descriptor</i>	Language specific
<i>item_num</i>	Unsigned integer that specifies ordinal position of an item descriptor in the SDA	$0 < item_num \leq (\text{number of item descriptors specified when SDA was allocated})$	“Literal Number ” on page 4-185
<i>item_num_var</i>	Host variable that stores <i>item_num</i>	Same restrictions as <i>item_num</i>	Language specific
<i>total_items</i>	Unsigned integer that specifies how many items the SDA describes	Same restrictions as <i>item_num</i>	“Literal Number ” on page 4-185
<i>total_items_var</i>	Host variable that stores <i>total_items</i>	Same restrictions as <i>total_items</i>	Language specific

Usage

The SET DESCRIPTOR statement can be used after you have described SELECT, EXECUTE FUNCTION, EXECUTE PROCEDURE, ALLOCATE DESCRIPTOR, or INSERT statements with the DESCRIBE ... USING SQL DESCRIPTOR statement.

SET DESCRIPTOR can assign values to a system-descriptor area in these cases:

- To set the **COUNT** field of a system-descriptor area to match the number of items for which you are providing descriptions in the system-descriptor area
- To set the item descriptor for each value for which you are providing descriptions in the system-descriptor area
- To modify the contents of an item-descriptor field

If an error occurs during the assignment to any identified system-descriptor fields, the contents of all identified fields are set to 0 or NULL, depending on the data type of the variable.

Using the COUNT Clause

Use the COUNT clause to set the number of items that are to be used in the system-descriptor area. If you allocate a system-descriptor area with more items than you are using, you need to set the **COUNT** field to the number of items that you are actually using. The following example shows a fragment of an Informix ESQL/C program:

```
EXEC SQL BEGIN DECLARE SECTION;
    int count;

EXEC SQL END DECLARE SECTION;

EXEC SQL allocate descriptor 'desc_100'; /*allocates for 100 items*/
    count = 2;
EXEC SQL set descriptor 'desc_100' count = :count;
```

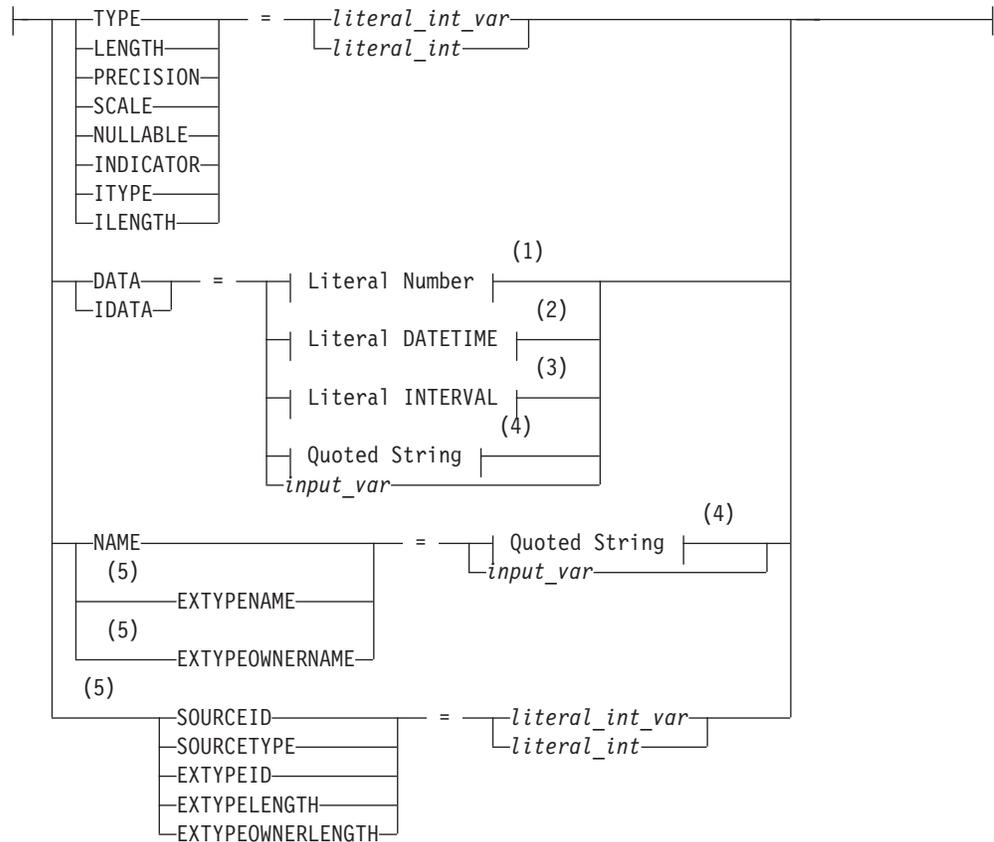
Using the VALUE Clause

Use the VALUE clause to assign values from host variables into fields of a system-descriptor area. You can assign values for items for which you are providing a description (such as parameters in a WHERE clause), or you can modify values for items after you use a DESCRIBE statement to fill the fields for an UPDATE or INSERT statement.

Item Descriptor

Use the Item Descriptor portion of the SET DESCRIPTOR statement to set value for an individual field in a system-descriptor area.

Item Descriptor:



Notes:

- 1 See "Literal Number " on page 4-185
- 2 See "Literal DATETIME " on page 4-181
- 3 See "Literal INTERVAL" on page 4-183
- 4 See "Quoted String" on page 4-189
- 5 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>input_var</i>	Host variable storing data for the specified item descriptor field	Must be appropriate for the specified field	Language-specific
<i>literal_int</i>	Integer value (> 0) assigned to the specified item descriptor field	Restrictions depend on the keyword to the left of = symbol	"Literal Number " on page 4-185
<i>literal_int_var</i>	Variable having <i>literal_int</i> value	Same as for <i>literal_int</i>	Language-specific

For information on codes that are valid for the TYPE or ITYPE fields and their meanings, see "Setting the TYPE or ITYPE Field" on page 2-608.

For the restrictions that apply to other field types, see the individual headings for field types under "Using the VALUE Clause" on page 2-606.

Setting the TYPE or ITYPE Field

Use these integer values to set the value of **TYPE** or **ITYPE** for each item.

SQL Data Type	Integer Value	X-Open Integer Value	SQL Data Type	Integer Value	X-Open Integer Value
CHAR	0	1	MONEY	8	–
SMALLINT	1	4	DATETIME	10	–
INTEGER	2	5	BYTE	11	–
FLOAT	3	6	TEXT	12	–
SMALLFLOAT	4	–	VARCHAR	13	–
DECIMAL	5	3	INTERVAL	14	–
SERIAL	6	–	NCHAR	15	–
DATE	7	–	NVARCHAR	16	–

The following table lists integer values that represent additional data types available with Dynamic Server.

SQL Data Type	Integer Value	SQL Data Type	Integer Value
INT8	17	COLLECTION	23
SERIAL8	18	Varying-length OPAQUE type	40
SET	19	Fixed-length OPAQUE type	41
MULTISET	20	LVARCHAR (client-side only)	43
LIST	21	BOOLEAN	45
ROW	22		

The same **TYPE** constants can also appear in the **syscolumns.coltype** column in the system catalog; see *IBM Informix Guide to SQL: Reference*.

For code that is easier to maintain, use the predefined constants for these SQL data types instead of their actual integer values. These constants are defined in the **\$INFORMIX/incl/public/sqltypes.h** header file. You cannot, however, use the actual constant name in the SET DESCRIPTOR statement. Instead, assign the constant to an integer host variable and specify the host variable in the SET DESCRIPTOR statement file.

The following example shows how you can set the **TYPE** field in Informix ESQL/C:

```
main()
{
EXEC SQL BEGIN DECLARE SECTION;
    int itemno, type;
EXEC SQL END DECLARE SECTION;
...
EXEC SQL allocate descriptor 'desc1' with max 5;
...
type = SQLINT; itemno = 3;
EXEC SQL set descriptor 'desc1' value :itemno type = :type;
}
```

This information is identical for **ITYPE**. Use **ITYPE** when you create a dynamic program that does not comply with the X/Open standard.

Compiling Without the -xopen Option: If you compile without the **-xopen** option, the normal Informix SQL code is assigned for **TYPE**. You must be careful not to mix normal and X/Open modes, because errors can result. For example, if a data type is not defined under X/Open mode, but is defined under normal mode, executing a **SET DESCRIPTOR** statement can result in an error.

Setting the TYPE Field in X/Open Programs: In X/Open mode, you must use the X/Open set of integer codes for the data type in the **TYPE** field.

If you use the **ILENGTH**, **IDATA**, or **ITYPE** fields in a **SET DESCRIPTOR** statement, a warning message appears. The warning indicates that these fields are not standard X/Open fields for a system-descriptor area.

For code that is easier to maintain, use the predefined constants for these X/Open SQL data types instead of their actual integer value. These constants are defined in the **\$INFORMIX/incl/public/sqlxtype.h** header file.

Using DECIMAL or MONEY Data Types: If you set the **TYPE** field for a **DECIMAL** or **MONEY** data type, and you want to use a scale or precision other than the default values, set the **SCALE** and **PRECISION** fields. You do not need to set the **LENGTH** field for a **DECIMAL** or **MONEY** item; the **LENGTH** field is set accordingly from the **SCALE** and **PRECISION** fields.

Using DATETIME or INTERVAL Data Types: If you set the **TYPE** field for a **DATETIME** or **INTERVAL** value, the **DATA** field can be a **DATETIME** or **INTERVAL** literal or a character string. If you use a character string, the **LENGTH** field must be the encoded qualifier value.

To determine the encoded qualifiers for a **DATETIME** or **INTERVAL** character string, use the **datetime** and **interval** macros in the **datetime.h** header file.

If you set **DATA** to a host variable of **DATETIME** or **INTERVAL**, you do not need to set **LENGTH** explicitly to the encoded qualifier integer.

Setting the DATA or IDATA Field

When you set the **DATA** or **IDATA** field, use the appropriate type of data (character string for **CHAR** or **VARCHAR**, integer for **INTEGER**, and so on).

If any value other than **DATA** is set, the value of **DATA** is undefined. You cannot set the **DATA** or **IDATA** field for an item without setting **TYPE** for that item. If you set the **TYPE** field for an item to a character type, you must also set the **LENGTH** field. If you do not set the **LENGTH** field for a character item, you receive an error.

Setting the LENGTH or ILENGTH Field

If your **DATA** or **IDATA** field contains a character string, you must specify a value for **LENGTH**. If you specify **LENGTH=0**, **LENGTH** is automatically set to the maximum length of the string. The **DATA** or **IDATA** field can contain a literal character string of up to 368-bytes, or a character string derived from a character variable of a **CHAR** or **VARCHAR** data type. This provides a method to determine dynamically the length of a string in the **DATA** or **IDATA** field.

If a DESCRIBE statement precedes a SET DESCRIPTOR statement, **LENGTH** is automatically set to the maximum length of the character field that is specified in your table.

This information is identical for **ILENGTH**. Use **ILENGTH** when you create a dynamic program that does not comply with the X/Open standard.

Setting the INDICATOR Field

If you want to put a NULL value into the system-descriptor area, set the **INDICATOR** field to -1 and do not set the **DATA** field.

If you set the **INDICATOR** field to 0 to indicate that the data is not NULL, you must set the **DATA** field.

Setting Opaque-Type Fields (IDS)

The following item-descriptor fields provide information about a column that has an opaque type as its data type:

- The **EXTYPEID** field stores the extended identifier for the opaque type. This integer value must correspond to a value in the **extended_id** column of the **sysxdtypes** system catalog table.
- The **EXTYPENAME** field stores the name of the opaque type. This character value must correspond to a value in the **name** column of the row with the matching **extended_id** value in the **sysxdtypes** system catalog table.
- The **EXTYPELENGTH** field stores the length of the opaque-type name. This integer value is the length, in bytes, of the string in the **EXTYPENAME** field.
- The **EXTYPEOWNERNAME** field stores the name of the opaque-type owner. This character value must correspond to a value in the **owner** column of the row with the matching **extended_id** value in the **sysxdtypes** system catalog table.
- The **EXTYPEOWNERLENGTH** field stores the length of the value in the **EXTYPEOWNERNAME** field. This integer value is the length, in bytes, of the string in the **EXTYPEOWNERNAME** field.

For more information on the **sysxdtypes** system catalog table, see the *IBM Informix Guide to SQL: Reference*.

Setting Distinct-Type Fields

The following item-descriptor fields provide information about a column that has a distinct type as its data type:

- The **SOURCEID** field stores the extended identifier for the source data type. Set this field if the source type of the distinct type is an opaque data type. This integer value must correspond to a value in the **source** column for the row of the **sysxdtypes** system catalog table whose **extended_id** value matches that of the distinct type you are setting.
- The **SOURCETYPE** field stores the data type constant for the source data type. This value is the data type constant for the built-in data type that is the source type for the distinct type. The codes for the **SOURCETYPE** field are the same as those for the **TYPE** field (page “Setting the TYPE or ITYPE Field” on page 2-608). This integer value must correspond to the value in the **type** column for the row of the **sysxdtypes** system catalog table whose **extended_id** value matches that of the distinct type you are setting.

For more information on the **sysxdtypes** system catalog table, see the *IBM Informix Guide to SQL: Reference*.

Modifying Values Set by the DESCRIBE Statement

You can use a DESCRIBE statement to modify the contents of a system-descriptor area after it is set.

After you use DESCRIBE on a SELECT or an INSERT statement, you must check to determine whether the **TYPE** field is set to either 11 or 12 to indicate a TEXT or BYTE data type. If **TYPE** contains an 11 or a 12, you must use the SET DESCRIPTOR statement to reset **TYPE** to 116, which indicates FILE type.

Related Information

Related statements: "ALLOCATE DESCRIPTOR " on page 2-2, "DEALLOCATE DESCRIPTOR " on page 2-284, "DECLARE " on page 2-286, "DESCRIBE" on page 2-311, "EXECUTE " on page 2-348, "FETCH " on page 2-367, "GET DESCRIPTOR " on page 2-381, "OPEN " on page 2-459, "PREPARE " on page 2-468, and "PUT " on page 2-479

For more information on system-descriptor areas, refer to the *IBM Informix ESQL/C Programmer's Manual*.

SET ENCRYPTION PASSWORD

Use the SET ENCRYPTION PASSWORD statement to define or reset a *session password* for encryption and decryption of character, BLOB, or CLOB values. Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL. You can use this statement with ESQL/C.

Syntax

```
►►—SET ENCRYPTION PASSWORD—'password' —————►►  
└──────────WITH HINT—'hint' ─────────┘
```

Element	Description	Restrictions	Syntax
<i>hint</i>	String that GETHINT returns from an encrypted argument	$(0 \text{ byte}) \leq \textit{hint} \leq (32 \text{ bytes})$. Do not include the <i>password</i> in the <i>hint</i> .	"Expression" on page 4-40
<i>password</i>	Password (or a multi-word phrase) for data encryption	$(6 \text{ bytes}) \leq \textit{password} \leq (120 \text{ bytes})$. Do not specify your login password.	"Expression" on page 4-40

Usage

The SET ENCRYPTION PASSWORD statement declares a password to support data confidentiality through built-in functions that use the Triple-DES or AES algorithms for encryption and decryption. These functions enable the database to store sensitive data in an encrypted format that prevents anyone who cannot provide the secret password from viewing, copying, or modifying encrypted data.

The password is not stored as plain text in the database and is not accessible to the DBA. This security feature is independent of the Trusted Facility feature.

Important: By default, communication between client systems and Dynamic Server is in plain text. Unless the database is accessible only by a secure network, the DBA must enable the encryption communication support module (ENCCSM) to

provide data encryption between the database server and any client system. Otherwise, an attacker might read the password and use it to access encrypted data.

If the network is not secure, all of the database servers in a distributed query need ENCCSM enabled, so that the password is not transmitted as plain text. For information about how to enable a communication support module (CSM), see your *IBM Informix Administrator's Guide*.

Operations on encrypted data tend to be slower than corresponding operations on plain text data, but use of this feature has no effect on unencrypted data.

The SET ENCRYPTION PASSWORD statements can be prepared, and EXECUTE IMMEDIATE can process a prepared SET ENCRYPTION PASSWORD statement.

Storage Requirements for Encryption

Use the ENCRYPT_AES or ENCRYPT_TDES built-in functions to encrypt data. Encrypted values of character data types are stored in BASE64 format (also called Radix-64). For character data, this requires significantly more storage than the corresponding unencrypted data. Omitting the *hint* can reduce encryption overhead by more than 50 bytes for each encrypted value. It is the responsibility of the user to make sufficient storage space available for encrypted values.

The following table lists the data types that can be encrypted, and built-in functions that you can use to encrypt and decrypt values of those data types:

Original Data Type	Encrypted Data Type	BASE64 Format	Decryption Function
CHAR	CHAR	Yes	DECRYPT_CHAR
NCHAR	NCHAR	Yes	DECRYPT_CHAR
VARCHAR	VARCHAR	Yes	DECRYPT_CHAR
NVARCHAR	NVARCHAR	Yes	DECRYPT_CHAR
LVARCHAR	LVARCHAR	Yes	DECRYPT_CHAR
BLOB	BLOB	No	DECRYPT_BINARY
CLOB	BLOB	No	DECRYPT_CHAR

You cannot encrypt a column of the IDSSECURITYLABEL data type.

If the encrypted VARCHAR (or NVARCHAR) value is longer than the 255 byte maximum size for those data types, the encryption function returns a CHAR (or NCHAR) value of sufficient size to store the encrypted value.

DECRYPT_BINARY and DECRYPT_CHAR both return the same value from encrypted CHAR, NCHAR, VARCHAR, NVARCHAR, or LVARCHAR values. No built-in encryption or decryption functions support BYTE or TEXT data types, but you can use BLOB data types to encrypt very large strings.

Warning: If the declared size of a database column in which you intend to store encrypted data is smaller than the encrypted data length, truncation occurs when you insert the encrypted data into the column. The truncated data cannot subsequently be decrypted, because the data length indicated in the header of the encrypted string does not match what the column stores. To avoid truncation,

make sure that any column storing encrypted strings has sufficient length. (See the cross-reference in the next paragraph for details of how to calculate encrypted string lengths.)

Besides the unencrypted data length, the storage required for encrypted data depends on the encoding format, on whether you specify a *hint*, and on the block size of the encryption function. For a formula to estimate the encrypted size, see "Calculating storage requirements for encrypted data" on page "Calculating storage requirements for encrypted data" on page 4-111.

Specifying a Session Password and Hint

The required *password* specification can be quoted strings or other character expression that evaluates to a string whose length is at least 6 bytes but no more than 128 bytes. The optional *hint* can specify a string no longer than 32 bytes.

The password or *hint* can be a single word or several words. The *hint* should be a word or phrase that helps you to remember the *password*, but does not include the *password*. You can subsequently execute the built-in **GETHINT** function (with an encrypted value as its argument) to return the plain text of *hint*.

The following ESQL/C program fragment defines a routine that includes the SET ENCRYPTION PASSWORD statement and executes DML statements:

```
process_ssn( )
{
EXEC SQL BEGIN DECLARE SECTION;
char password[128];
char myhint[33];
char myid[16], myssn[16];
EXEC SQL END DECLARE SECTION;
. . .
EXEC SQL SET ENCRYPTION PASSWORD :password WITH HINT :myhint;
. . .
EXEC SQL INSERT INTO tab1 VALUES (':abcd', ENCRYPT_AES("111-22-3333")) ;
EXEC SQL SELECT Pid, DECRYPT(ssn, :password) INTO :myid, :myssn;
. . .
EXEC SQL SELECT GETHINT(ssn) INTO :myhint, WHERE id = :myid;
}
```

Levels of Encryption

You can use SET ENCRYPTION PASSWORD with encryption and decryption functions to support these granularities of encryption in the database.

- **Column-Level Encryption:** All values in a given column of a database table are encrypted using the same password, the same encryption algorithm, and the same encryption mode. (In this case, you can save disk space by storing the *hint* outside the encrypted column, rather than repeating it in every row.)
- **Cell-Level Encryption:** Values of a given column in different rows of the same database table are encrypted using different passwords, or different encryption algorithms, or different encryption modes. This technique is sometimes necessary to protect personal data. (*Row-column level* encryption and *set-column level* encryption are both synonyms for cell-level encryption.)

Cell-level encryption can cause substantial maintenance costs. If you implement this level of encryption, your application is responsible for determining which rows contain encrypted data and for using the correct code to handle the data. The built-in decryption functions of Dynamic Server fail with error -26005 if they are applied to unencrypted data. The simplest way to avoid this error is to use column-level encryption rather than cell-level encryption.

If you do not use encryption functions, people might enter unencrypted data into columns that are meant to contain encrypted data. To ensure that data entered into a field is always encrypted, use views and INSTEAD OF triggers.

Protecting Passwords

Passwords and hints that you declare with SET ENCRYPTION PASSWORD are not stored as plain text in any table of the system catalog, which also maintains no record of which columns or tables contain encrypted data. To prevent other users from accessing the plain text of encrypted data or of a password, however, you must avoid actions that might compromise the secrecy of a password:

- Do not create a functional index using a decryption function. (This would store plain-text data in the database, defeating the purpose of encryption.)
- On a network that is not secure, always work with encrypted data, or use session encryption, because the SQL communication between client and server sends passwords, hints, and the data to be encrypted as plain text.
- Do not store passwords in a trigger or in a UDR that exposes the password to the public.
- Do not set the session password prior to creating any view, trigger, procedure, or UDR. Set the session password only when you use the object. Otherwise, the password might be visible in the schema to other users, and queries executed by other users might return unencrypted data. The following example shows a procedure that includes an encrypted password:

```
-- reset session encryption password
set encryption password null;

-- create procedure without password
create procedure p1 ();
  insert into tab2 select (decrypt_char (col1))
  from tab1;
end procedure;

-- set session encryption password
set encryption password ("PASSWD2");

-- insert data
insert into tab1 values (encrypt_aes ('WXY'));

-- call procedure
```

Output from the SET EXPLAIN statement always displays the *password* and *hint* parameters as XXXXX, rather than displaying actual *password* or *hint* values.

Related Information

For more information about built-in functions for encrypting and decrypting data, see “Encryption and Decryption Functions” on page 4-107.

For information on setting the environment variable **INFORMIXCONCSMCFG**, refer to the *IBM Informix Guide to SQL: Reference*.

Usage

SET ENVIRONMENT specifies environment options that can affect queries or manage resource use by the routine in which the statement is executed. Some options can override the setting of an environment variable or of a configuration parameter during the session in which SET ENVIRONMENT is issued. For example, on Dynamic Server, the SET ENVIRONMENT OPTCOMPIND '2' statement instructs the query optimizer to use cost as the basis for subsequent join plans during the session, even if this behavior conflicts with the current setting of the OPTCOMPIND environment variable.

For several environment options that Extended Parallel Server supports, the following keywords have similar effects

- OFF disables the specified option
- ON enables the option
- DEFAULT sets the option to its default value
- MUTABLE makes the option modifiable during a user session
- IMMUTABLE prevents modification of the option during a user session.

The arguments that follow the option name depend on the syntax of the option. The option name and its ON, OFF, and DEFAULT keywords are not quoted and are not case sensitive. All other arguments must be enclosed between single (') or double (") quotation marks. If a quoted string is a valid argument for an environment option, the argument is case sensitive.

For Dynamic Server, if you specify an undefined environment option name, error -19840 is returned. If you specify an invalid *value* for a valid option, an option-specific error is returned. (for example, error -19843, Invalid IFX_AUTO_REPREPARE value specified). The SET ENVIRONMENT statement can enable only the environment options that are described in sections that follow.

For information about the performance implications of the SET ENVIRONMENT options, refer to the *IBM Informix Performance Guide*.

BOUND_IMPL_PDQ Environment Option (XPS)

If IMPLICIT_PDQ is set to ON or to a value, use the BOUND_IMPL_PDQ environment option to specify that the allocated memory should be bounded by the current explicit PDQPRIORITY value or range. If IMPLICIT_PDQ is OFF, then BOUND_IMPL_PDQ is ignored. For example, you might execute the following statement to force the database server to use explicit PDQPRIORITY values as guidelines in allocating memory if the IMPLICIT_PDQ environment option has already been set:

```
SET ENVIRONMENT BOUND_IMPL_PDQ ON;
```

If you set both IMPLICIT_PDQ and BOUND_IMPL_PDQ, then the explicit PDQPRIORITY value determines the upper limit of memory that can be allocated to a query. If PDQPRIORITY is specified as a range, the database server grants memory within the range specified.

For detailed information, see the *IBM Informix Performance Guide*.

CLIENT_TZ Environment Option (XPS)

Use the CLIENT_TZ environment option to set the timezone offset from GMT to use in the current session for evaluating the built-in CURRENT and TODAY functions. If you specify OFF or DEFAULT, values returned from CURRENT and TODAY use the timezone of the database server.

To specify a timezone offset, use the following format for the CLIENT_TZ *value*, and enclose the *value* between a pair of single (') or double (") quotation marks:

► [+] hour : minute ◄

Element	Description	Restrictions	Syntax
<i>hour</i>	Number of hours in timezone offset from Greenwich Mean Time (GMT)	Must be an integer in the range $-14 < hour < 14$	"Literal Number " on page 4-185
<i>minute</i>	Additional minutes in timezone offset from Greenwich Mean Time (GMT)	Must be a 2-digit integer in the range 00 to 59 inclusive	"Literal Number " on page 4-185

A positive (+) offset implies a time zone whose location is west of zero longitude (= GMT). A negative offset implies a time zone east of zero longitude. The following statement specifies an offset of five hours and fifteen minutes:

```
SET ENVIRONMENT CLIENT_TZ '+5:15';
```

If the statement in this statement executed successfully, then subsequent calls to the CURRENT function in the same session return a time-of-day value that is 5 hours and 15 minutes greater than the default value with no offset.

The following example illustrates the effect on CURRENT output after a UDR is invoked that uses SET ENVIRONMENT to reset CLIENT_TZ from an initial offset from GMT of -8 hours to a -5 hour offset:

```
1> create procedure p1()
1> set environment CLIENT_TZ '-5:00';
1> end procedure;
```

Procedure created.

```
2> select current from systables where tabid = 1;
```

```
(expression)
2007-01-17 09:21:44.000
```

```
3> execute procedure p1();
Procedure executed.
```

```
4> select current from systables where tabid = 1;
```

```
(expression)
2007-01-17 12:21:53.000
```

This example also illustrates how the scope of the setting for CLIENT_TZ, like all the other SET ENVIRONMENT options of Extended Parallel Server, is the entire session, rather than a scope local to the routine that issues the SET ENVIRONMENT statement. The new setting persists until the session ends, or until a subsequent SET ENVIRONMENT statement resets the CLIENT_TZ value.

If a `CLIENT_TZ` value has been previously specified, the keyword `ON` tells the database server to use that value.

The `CLIENT_TZ` value that you specify in the `SET ENVIRONMENT` statement temporarily overrides any `CLIENT_TZ` setting of the `IBM_XPS_PARAMS` environment variable for the duration of the current session.

The `onstat -g ses` command can display the current `CLIENT_TZ` setting.

In a distributed transaction that you initiate from Extended Parallel Server, but in which remote Dynamic Server instances evaluate `CURRENT` or `TODAY` values, the `CLIENT_TZ` setting is ignored by Dynamic Server.

COMPUTE_QUOTA Environment Option (XPS)

Use the `COMPUTE_QUOTA` environment option to allow the optimizer to use only one CPU VP (virtual processor) on each coserver for each query operator instead of using all CPU VPs of each coserver.

To turn on this environment option, execute the following statement:

```
SET ENVIRONMENT COMPUTE_QUOTA ON;
```

IMPLICIT_PDQ Environment Option (XPS)

Use the `IMPLICIT_PDQ` environment option to allow the database server to determine the amount of memory allocated to a query. Unless `BOUND_IMPL_PDQ` is also set, the database server ignores the current explicit setting of `PDQPRIORITY`. It does not allocate more memory, however, than is available when `PDQPRIORITY` is set to 100, as determined by $\text{MAX_PDQPRIORITY} / 100 * \text{DS_TOTAL_MEMORY}$.

This environment option is `OFF` by default.

If you set *value* between 1 and 100, the database server scales its estimate by the specified value. If you set a low value, the amount of memory assigned to the query is reduced, which might increase the amount of query-operator overflow.

For example, to request the database server to determine memory allocations for queries and distribute memory among query operators according to their needs, enter the following statement:

```
SET ENVIRONMENT IMPLICIT_PDQ ON;
```

To require the database server to use explicit `PDQPRIORITY` settings as the upper bound and optional lower bound of memory that it grants to a query, set the `BOUND_IMPL_PDQ` environment option.

MAXSCAN Environment Option (XPS)

Use the `MAXSCAN` environment option to change the default number of scan threads on each coserver. The default is three scan threads for each CPU VP. You might want to reduce the number of scan threads on a coserver if the `GROUP`, `JOIN`, and other operators above the scan are not producing rows quickly enough to keep the default number of scan threads busy.

For some queries, you might want to increase the number of scan threads on each coserver. For example, if each coserver has three CPU VPs, the database server can

create nine scan threads on each coserver. To increase the number to four threads for each CPU VP, execute the following statement:

```
SET ENVIRONMENT MAXSCAN '12';
```

MAXSCAN is automatically set to 1 if the COMPUTE_QUOTA environment option is enabled or if the isolation level is set to Cursor Stability and the database server can use a pipe operator.

TMPSPACE_LIMIT Environment Option (XPS)

Use the TMPSPACE_LIMIT environment option to specify the amount of temporary space on each coserver that a query can use for query operator overflow. Temporary space limits do not affect the creation of temporary tables. The limits apply only to the query-operator overflow that occurs when a query cannot get enough memory to complete execution.

By default, TMPSPACE_LIMIT is OFF, and a query can use all available temporary space for operator overflow. If the DS_TOTAL_TMPSPACE configuration parameter is not set, then setting TMPSPACE_LIMIT has no effect.

If you enter a value between '0' and '100' as an argument to the TMPSPACE_LIMIT option, the database server sets the temporary space quota to the percent of available temporary space that the ONCONFIG parameter DS_TOTAL_TMPSPACE, specifies.

- If the value is 100, queries can use only the amount of temporary space on each coserver that DS_TOTAL_TMPSPACE specifies.
- If the value is 0, query operators cannot overflow to temporary space.
- If you do not specify a value, a query can use all available temporary space on each coserver, as limited by DS_TOTAL_TMPSPACE, for query operator overflow.

To require queries to use only the amount of temporary space specified by the setting of DS_TOTAL_TMPSPACE on each coserver, execute this statement:

```
SET ENVIRONMENT TMPSPACE_LIMIT ON;
```

To limit queries to 50 percent of DS_TOTAL_TMPSPACE on each coserver, execute the following statement:

```
SET ENVIRONMENT TMPSPACE_LIMIT "50";
```

TEMP_TAB_EXT_SIZE and TEMP_TAB_NEXT_SIZE Options (XPS)

Use the TEMP_TAB_EXT_SIZE and TEMP_TAB_NEXT_SIZE environment options to specify the number of kilobytes to allocate as the first and next extent sizes, respectively, for system-generated temporary tables.

The following example shows the syntax for setting the first and next extent sizes using these environment options:

```
SET ENVIRONMENT TEMP_TAB_EXT_SIZE '64';  
SET ENVIRONMENT TEMP_TAB_NEXT_SIZE '128';
```

These statements set the first extent size of a generated temporary table to 64 kilobytes and the next extent size to 128 kilobytes.

The minimum value of these options is four times the page size on your system. If you specify a size below this minimum, the database server instead uses a default size of four pages. For flex inserts, the default size is 32 pages or 128 kilobytes.

The maximum value for TEMP_TAB_EXT_SIZE and TEMP_TAB_NEXT_SIZE is the maximum value of a chunk size.

Use the DEFAULT keyword to reset the values of these environments options back to the system defaults.

Important: The database server calculates the extent sizes of temporary tables that it creates for hash tables and sorts. The SET ENVIRONMENT statement has no effect on extent sizes for these tables.

EXTDIRECTIVES Environment Option (IDS)

You can use the EXTDIRECTIVES environment option of the SET ENVIRONMENT statement to enable or disable external directives during a session. The EXTDIRECTIVES option overwrites the value for enabling or disabling the external directive that is set in the ONCONFIG file.

To disable the external directives during a session, specify 0, off, or OFF as the value for SET ENVIRONMENT EXTDIRECTIVES.

To enable the external directives during a session, specify 1, on, or ON as the value for SET ENVIRONMENT EXTDIRECTIVES.

To enable the default values specified in the EXT_DIRECTIVES configuration parameter and the client-side IFX_EXTDIRECTIVES environment variable, specify DEFAULT as the value for SET ENVIRONMENT EXTDIRECTIVES.

IFX_AUTO_REPREPARE Environment Option (IDS)

The IFX_AUTO_REPREPARE environment option can reduce the incidence of SQL error -710 in databases that dynamic SQL applications access. This error might be issued after a cursor attempts to execute a prepared object or an SPL routine after DDL operations have changed the schema of a table that the prepared object or the stored procedure references.

After some changes to a database table, such as adding an enabled index, enabling the IFX_AUTO_REPREPARE feature can avoid -710 errors, and can reduce the need for users to issue the PREPARE statement explicitly to recompile prepared objects or to reoptimize stored procedures explicitly with the UPDATE STATISTICS statement. When this environment option is enabled, and the changes to table schema do not require reissuing the DESCRIBE statement, the database server automatically identifies and recompiles prepared statements and SPL routines that reference the modified table.

The value that the SET ENVIRONMENT IFX_AUTO_REPREPARE statement specifies can enable or disable this automatic recompilation feature:

- If the specified IFX_AUTO_REPREPARE value is '1' or 'ON' or 'on', then automatic recompilation is enabled.
- If the specified IFX_AUTO_REPREPARE value is '0' or 'OFF' or 'off', then automatic recompilation is disabled.

The following statement enables automatic recompilation after DDL operations on tables that prepared objects or SPL routines reference:

```
SET ENVIRONMENT IFX_AUTO_REPREPARE '1';
```

This overrides the setting of the AUTO_REPREPARE configuration parameter, if it is zero or 'None', for the remainder of the current session, or until you reset IFX_AUTO_REPREPARE.

The database server might not detect some changes to a table schema that invalidate prepared objects or SPL routines, even when IFX_AUTO_REPREPARE is enabled. For example, changes to a table schema by one session causes concurrent sessions to receive error -710 when they attempt to read the same table after obtaining a shared lock.

Enabling IFX_AUTO_REPREPARE has no effect on prepared statements and SPL routines that reference tables in which DDL operations change the number of columns in the table, or change the data type of a column. After these schema changes, you typically must reissue the DESCRIBE statement, the PREPARE statement, and (for cursors associated with routines) the UPDATE STATISTICS statement for any routines that reference the table whose schema has been modified.

If you are satisfied with how your applications currently handle errors from schema changes, you can disable automatic recompilation, as in this example:

```
SET ENVIRONMENT IFX_AUTO_REPREPARE 'OFF';
```

If enabling the IFX_AUTO_REPREPARE session environment variable results in an error, that error is passed back to the application.

For more information about the AUTO_REPREPARE configuration parameter, see your *IBM Informix Administrator's Reference*. For a discussion of the effects of the IFX_AUTO_REPREPARE and AUTO_REPREPARE settings on cursors and on queries, see your *IBM Informix Performance Guide*.

OPTCOMPIND Environment Option (IDS)

The OPTCOMPIND environment option can improve the performance of databases that are used for both decision support and online transaction processing. Use this option to specify join methods for the query optimizer to use in subsequent queries.

- If the value is '0' then the query optimizer uses a nested-loop join where possible, rather than a sort-merge join or a hash join.
- If the value is '1' and the transaction isolation level is Repeatable Read, the optimizer behaves as in setting '0', as described above; for any other isolation level, it behaves like setting '2', as described next.
- If the value is '2' then the query optimizer does not necessarily prefer nested-loop joins, but bases its decision entirely on the estimated cost, regardless of the transaction isolation mode.

For example, the following statement replaces whatever OPTCOMPIND setting was previously in effect with a purely cost-based optimizer strategy:

```
SET ENVIRONMENT OPTCOMPIND '2';
```

Use the DEFAULT keyword to restore the system default value, as described in the OPTCOMPIND section of the *IBM Informix Guide to SQL: Reference*.

For performance implications of the OPTCOMPIND option, see your *IBM Informix Performance Guide*.

The scope of the OPTCOMPIND setting that the SET ENVIRONMENT 'OPTCOMPIND' statement specifies is local to the routine that issues the statement, rather than persisting for the entire session. After that routine terminates, the setting reverts to the system default value that the OPTCOMPIND environment variable specifies.

No other option to SET ENVIRONMENT of Dynamic Server or of Extended Parallel Server has a scope that is local to the routine. The settings of all the other SET ENVIRONMENT options persist until the session ends, or until another SQL statement resets their value.

USELASTCOMMITTED Environment Option (IDS)

The USELASTCOMMITTED environment option can improve concurrency in sessions that use the Committed Read, Dirty Read, Read Committed, or Read Uncommitted isolation levels by reducing the risk of locking conflicts when two or more sessions attempt to access the same row in a table whose locking granularity is row-level locking.

The SET ENVIRONMENT USELASTCOMMITTED statement can specify whether queries and other operations that encounter exclusive locks that other sessions hold while changing data values should use the most recently committed version of the data, rather than wait for the lock to be released.

This statement can override the USELASTCOMMITTED configuration parameter setting for the duration of the current session. You can use the SET ISOLATION statement to override the USELASTCOMMITTED session environment setting.

The USELASTCOMMITTED option can have any one of four values:

- If the value is 'COMMITTED READ' then the database server reads the most recently committed version of the data when it encounters an exclusive lock while attempting to read a row in the Committed Read or Read Committed isolation level.
- If the value is 'DIRTY READ' then the database server reads the most recently committed version of the data if it encounters an exclusive lock while attempting to read a row in the Dirty Read or Read Committed isolation level.
- If the value is 'ALL' then the database server reads the most recently committed version of the data if it encounters an exclusive lock while attempting to read a row in the Committed Read, Dirty Read, Read Committed, or Read Uncommitted isolation level.
- If the value is 'NONE', this value disables the USELASTCOMMITTED feature that can access the last committed version of data in a locked row. Under this setting, if your session encounters an exclusive lock when attempting to read a row in the Committed Read, Dirty Read, Read Committed, or Read Uncommitted isolation level, your transaction cannot read that row until the concurrent transaction that holds the an exclusive lock is committed or rolled back.

For example, the following statements specify the Committed Read isolation mode and replace the explicit or default USELASTCOMMITTED configuration parameter setting with a setting that reads the most recently committed version of the data in rows on which concurrent readers hold an exclusive lock:

```
SET ISOLATION COMMITTED READ;  
SET ENVIRONMENT USELASTCOMMITTED 'ALL';
```

Any SPL routine can use these statements to specify the Committed Read Last Committed transaction isolation level during a session. These statements enable SQL operations that read data to use the last committed version when an exclusive lock is encountered during an operation that reads a row. This can avoid deadlock situations or other locking errors when another session is attempting to modify the same row. It does not reduce the risk of locking conflicts with other sessions that are writing to tables, or with concurrent DDL transactions that hold implicit or explicit locks on a user table or on a system catalog table.

For example, the following statements within a `PUBLIC.sysdbopen` or `user.sysdbopen` procedure specify at connection time the Committed Read isolation mode and replace the explicit or default `USELASTCOMMITTED` configuration parameter setting with a setting that reads the most recently committed version of the data in tables on which concurrent readers hold an exclusive lock:

```
SET ISOLATION COMMITTED READ;  
SET ENVIRONMENT USELASTCOMMITTED 'ALL';
```

Besides `sysdbopen()`, any SPL routine can use these statements to specify the Committed Read Last Committed transaction isolation level during a session. These statements enable SQL operations that read data to use the last committed version when an exclusive lock is encountered during an operation that reads a table. This can avoid deadlock situations or other locking errors when another session is attempting to modify the same row or table. It does not reduce the risk of locking conflicts with other sessions that are writing to tables, or with concurrent DDL transactions that hold implicit or explicit locks on a user table or on a system catalog table.

In cross-server distributed queries, if the isolation level of the session that issued the query has the `LAST COMMITTED` isolation level option in effect, but one or more of the participating databases does not support this `LAST COMMITTED` feature, then the entire transaction conforms to the Committed Read or Dirty Read isolation level of the session that issued the transaction, without the `LAST COMMITTED` option enabled.

For information about additional restrictions that can prevent a transaction from reading the most recently committed data from a table locked by another transaction while `USELASTCOMMITTED` is enabled, see “The `LAST COMMITTED` Option to Committed Read” on page 2-634.

For more information about the `USELASTCOMMITTED` configuration parameter, see your *IBM Informix Administrator's Reference*.

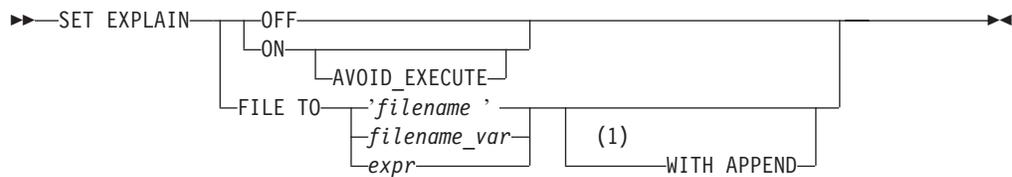
Related Information

Related statements: “`SET ALL_MUTABLES`” on page 2-581, “`SET ISOLATION`” on page 2-631, “`SET PDQPRIORITY`” on page 2-643

SET EXPLAIN

Use the `SET EXPLAIN` statement to enable (or disable) recording measurements of queries in the current session, including the plan of the query optimizer, an estimate of the number of rows returned, the relative cost of the query, and additional query measurement information.

Syntax



Notes:

- 1 Extended Parallel Server only

Element	Description	Restrictions	Syntax
<i>expr</i>	Expression that returns a <i>filename</i> specification	Must return a string satisfying the restrictions on <i>filename</i>	"Expression" on page 4-40
<i>filename</i>	Path and filename of the file to receive the output. For the default, see "Location of the Output File" on page 2-600.	Must conform to operating-system rules. If an existing file, see the WITH APPEND Option" on page 2-600.	"Quoted String" on page 4-189
<i>filename_var</i>	Host variable that stores <i>filename</i>	Must be a character data type	Language specific

Usage

Output from a SET EXPLAIN ON statement is directed to the appropriate file until you issue a SET EXPLAIN OFF statement or until the program ends. If you do not enter a SET EXPLAIN statement, then the default behavior is OFF, and the database server does not generate measurements for queries.

The SET EXPLAIN statement executes during the database server optimization phase, which occurs when you initiate a query. For queries that are associated with a cursor, if the query is prepared and does not have host variables, optimization occurs when you prepare it. Otherwise, optimization occurs when you open the cursor.

The SET EXPLAIN statement provides various measurements of the work involved in performing a query.

Option Effect

ON Generates measurements for each subsequent query and writes the results to an output file in the current directory. If the file already exists, new output is appended to the existing file.

AVOID_EXECUTE

Prevents a SELECT, INSERT, UPDATE, or DELETE statement from executing while the database server prints the query plan to an output file

OFF Terminates activity of the SET EXPLAIN statement, so that measurements for subsequent queries are no longer generated or written to the output file

FILE TO

Generates measurements for each subsequent query and allows you to specify the location for the explain output file. In Extended Parallel Server, if the specified file already exists, new output overwrites the contents of the file unless you use the WITH APPEND option.

The following example writes a query plan in the default **sqexplain.out** file for subsequent queries in the current session:

```
SET EXPLAIN ON;
```

The following example suspends writing query plan to a file in the current session:

```
SET EXPLAIN OFF;
```

Using the AVOID_EXECUTE Option

The SET EXPLAIN ON AVOID_EXECUTE statement activates the Avoid Execute option for a session, or until the next SET EXPLAIN OFF (or ON) without AVOID_EXECUTE. The AVOID_EXECUTE keyword prevents DML statements from executing; instead, the database server prints the query plan to an output file. If you activate AVOID_EXECUTE for a query that contains a remote table, the query does not execute at either the local or remote site.

The following example stores the output in the specified file.

```
SET EXPLAIN ON AVOID_EXECUTE;  
SET EXPLAIN FILE TO 'tmp/explain.out' WITH APPEND;
```

This form of the SET EXPLAIN statement is valid for XPS only.

When AVOID_EXECUTE is set, the database server sends a warning message. If you are using DB-Access, it displays a text message

```
Warning! avoid_execute has been set
```

for any select, delete, update or insert query operations. From ESQL, the **sqlwarn.sqlwarn7** character is set to 'W'.

Use the SET EXPLAIN ON or the SET EXPLAIN OFF statement to turn off the AVOID_EXECUTE option. The SET EXPLAIN ON statement turns off the AVOID_EXECUTE option but continues to generate a query plan and writes the results to an output file.

If you issue the SET EXPLAIN ON AVOID_EXECUTE statement in an SPL routine, the SPL routine and any DDL statements still execute, but the DML statements inside the SPL routine do not execute. The database server prints the query plan of the SPL routine to an output file. To turn off this option, you must execute the SET EXPLAIN ON or the SET EXPLAIN OFF statement outside the SPL routine. If you execute the SET EXPLAIN ON AVOID_EXECUTE statement before you execute an SPL routine, the DML statements inside the SPL routine do not execute, and the database server does not print a query plan of the SPL routine to an output file.

Nonvariant functions in a query are still evaluated when AVOID_EXECUTE is in effect, because the database server calculates these functions before optimization.

For example, the **func()** function is evaluated, even though the following SELECT statement is not executed:

```
SELECT * FROM orders WHERE func(10) > 5;
```

For other performance implications of the AVOID_EXECUTE option, see your *IBM Informix Performance Guide*.

If you execute the SET EXPLAIN ON AVOID_EXECUTE statement before you open a cursor in an Informix ESQL/C program, each FETCH operation returns the message that the row was not found. If you execute SET EXPLAIN ON

AVOID_EXECUTE after an Informix ESQL/C program opens a cursor, however, this statement has no effect on the cursor, which continues to return rows.

Using the FILE TO Option

When you execute a SET EXPLAIN FILE TO statement, explain output is implicitly turned on. The default filename for the output is **sqexplain.out** until changed by a SET EXPLAIN FILE TO statement. Once changed, the filename remains set until the end of the session or until it is changed by another SET EXPLAIN FILE TO statement.

The filename can be any valid combination of optional path and filename. If no path component is specified, the file is placed in your current directory. The permissions for the file are owned by the current user.

The output file that you specify in the SET EXPLAIN statement can be a new file or an existing file. In Dynamic Server, if the FILE TO clause specifies an existing file, the new output is appended to that file.

Using the WITH APPEND Option (XPS)

In Extended Parallel Server, if you specify an existing file, the current contents of that file are purged when you issue the SET EXPLAIN FILE TO statement, which sends its output to the beginning of the specified file.

If you include the WITH APPEND option, the current contents of the file are preserved when you issue the SET EXPLAIN FILE TO statement. The execution of a SET EXPLAIN . . . WITH APPEND statement appends output to the end of the file.

If you specify a new file in the SET EXPLAIN FILE TO statement, it makes no difference whether you include the WITH APPEND option. The first execution of the statement sends output to the beginning of the new file.

Dynamic Server does not support the WITH APPEND keywords, but if an output file already exists, any subsequent query plan is appended to that file, so that the functionality of the WITH APPEND feature is provided by default.

Default Name and Location of the Output File on UNIX

When you issue the SET EXPLAIN ON statement, the plan that the optimizer chooses for each subsequent query is written to the **sqexplain.out** file by default.

If the output file does not exist when you issue SET EXPLAIN ON, the database server creates the output file. If the output file already exists when you issue the SET EXPLAIN ON statement, subsequent output is appended to the file.

If the client application and the database server are on the same computer, the **sqexplain.out** file is stored in your current directory. If you are using a Version 5.x or earlier client application and the **sqexplain.out** file does not appear in the current directory, check your home directory for the file. When the current database is on another computer, the **sqexplain.out** file is stored in your home directory on the remote host.

Default Name and Location of the Output File on Windows

On Windows, SET EXPLAIN ON writes the plan that the optimizer chooses for each subsequent query to file %INFORMIXDIR%\sqexpln\username.out where *username* is the user login.

SET EXPLAIN Output

By examining the SET EXPLAIN output file, you can determine if steps can be taken to improve the performance of the query. The following table lists terms that can appear in the output file and their significance.

Table 2-3. Output file terms

Term	Significance
Query	<p>Displays the executed query and indicates whether SET OPTIMIZATION was set to HIGH or LOW. If you SET OPTIMIZATION to LOW, the output displays the following uppercase string as the first line: QUERY: {LOW}</p> <p>If you SET OPTIMIZATION to HIGH, the output of SET EXPLAIN displays the following uppercase string as the first line: QUERY:</p>
Directives followed	<p>Lists the directives set for the query</p> <p>If the syntax for a directive is incorrect, the query is processed without the directive. In that case, the output shows DIRECTIVES NOT FOLLOWED in addition to DIRECTIVES FOLLOWED.</p> <p>For more information on the directives specified after this term, see the “Optimizer Directives ” on page 5-35 or “SET OPTIMIZATION ” on page 2-641.</p> <p>If a DELETE or UPDATE statement specifies an uncorrelated subquery in the WHERE clause, the set of qualifying rows returned by the subquery is materialized as a temporary table, and the output of SET EXPLAIN displays within parentheses the following message: (Temp Table For Subquery)</p>
Estimated cost	<p>An estimate of the amount of work for the query</p> <p>The optimizer uses an estimate to compare the cost of one path with another. The estimate is a number the optimizer assigns to the selected access method. This number does not translate directly into time and cannot be used to compare different queries. It can be used, however, to compare changes made for the same query. When data distributions are used, a query with a higher estimate generally takes longer to run than one with a smaller estimate.</p> <p>In the case of a query and a subquery, two estimated cost figures are returned; the query figure also includes the subquery cost. The subquery cost is shown so that you can see the cost that is associated with only the subquery.</p>
Estimated number of rows returned	<p>An estimate of the number of rows to be returned</p> <p>This number is based on information in the system catalog tables.</p>
Numbered list	<p>The order in which tables are accessed, followed by the access method used (index path or sequential scan)</p> <p>When a query involves table inheritance, all the tables are listed under the supertable in the order in which they were accessed.</p>

Table 2-3. Output file terms (continued)

Term	Significance
Index keys	<p>The columns used as filters or indexes; the column name used for the index path or filter is indicated. The notation (<i>Key Only</i>) indicates that all the desired columns are part of the index key, so a key-only read of the index could be substituted for a read of the actual table.</p> <p>The <i>Lower Index Filter</i> shows the key value where the index read begins; and the <i>Upper Index Filter</i> is shown for the key value where the index read stops. The <i>Index Key Filters</i> show filters that will be applied on retrieved index key values. If the query uses an index self-join path, the <i>Index Self Join Keys</i> shows the leading index key columns used as self-join keys, and the <i>Lower bound</i> and <i>Upper bound</i> show the boundaries of the leading index key columns.</p>
Join method	<p>When the query involves a join between two tables, the join method that the optimizer used (Nested Loop or Dynamic Hash) is shown at the bottom of the output for that query.</p> <p>When the query involves a dynamic join of two tables, if the output contains the words <i>Build Outer</i>, the hash table is built on the first table listed (called the build table). If the words <i>Build Outer</i> do not appear, the hash table is built on the second table listed.</p>
Query statistics	<p>When the EXPLAIN_STAT configuration parameter is set to 1, this section shows the number of rows returned, the number of rows estimated in the query plan, the time required, calls to iterator functions, and the estimated cost of scan and join operations on table objects.</p>

If the query uses a collating order other than the default for the **DB_LOCALE** setting, then the **DB_LOCALE** setting and the name of the other locale that is the basis for the collation in the query (as specified by the SET COLLATION statement) are both included in the output file. Similarly, if an index is not used because of its collation, the output file indicates this.

Complete-Connection Level Settings

The SET EXPLAIN statement supports *complete-connection level* settings. This means that values in the local session environment at the time of connection are propagated to all new or resumed transactions of the following types:

- transactions within the local database
- distributed transactions across databases of the same server instance
- distributed transactions across databases of two or more database server instances
- global transactions with XA-compliant data sources that are registered in the local database

If you change the SET EXPLAIN setting within a transaction, the new value is propagated back to the local environment and also to all subsequent new or resumed transactions.

Examples of SET EXPLAIN Output

The following SQL statements cause Dynamic Server to write the query plans of the UPDATE statement (and of its subquery) to the default **sqexplain.out** output file:

```

DATABASE stores_demo;
SET EXPLAIN ON;
UPDATE orders SET ship_charge = ship_charge + 2.00
  WHERE customer_num IN
    (SELECT orders.customer_num FROM orders
     WHERE orders.ship_weight < 50);
CLOSE DATABASE;

```

The following information is displayed in the resulting output:

QUERY:

```

-----
update orders set ship_charge = ship_charge + 2.00
where customer_num in
(select orders.customer_num from orders where
  orders.ship_weight < 50)

```

Estimated Cost: 4
Estimated # of Rows Returned: 8

1) informix.orders: INDEX PATH

(1) Index Keys: customer_num (Serial, fragments: ALL)
Lower Index Filter: informix.orders.customer_num = ANY

Subquery:

```

-----
Estimated Cost: 2
Estimated # of Rows Returned: 8
(Temp Table For Subquery)

```

1) informix.orders: SEQUENTIAL SCAN

Filters: informix.orders.ship_weight < 50.00

The next example is based on the following SQL statements, which include a DELETE operation:

```

DATABASE stores_demo;
SET EXPLAIN ON;
DELETE FROM catalog WHERE stock_num IN
  (SELECT stock.stock_num FROM stock, catalog WHERE
   stock.stock_num = catalog.stock_num
   AND stock.unit_price < 50);
CLOSE DATABASE;

```

Below is the resulting output:

QUERY:

```

-----
DELETE FROM catalog WHERE stock_num IN
  (SELECT stock.stock_num from stock, catalog
   WHERE stock.stock_num = catalog.stock_num
   AND stock.unit_price < 50);

```

Estimated Cost: 19
Estimated # of Rows Returned: 37

1) ajay.catalog: INDEX PATH

(1) Index Keys: stock_num manu_code (Serial, fragments: ALL)
Lower Index Filter: ajay.catalog.stock_num = ANY

Subquery:

```

-----
Estimated Cost: 12
Estimated # of Rows Returned: 44
(Temp Table For Subquery)

```

- 1) ajay.stock: SEQUENTIAL SCAN
 Filters: ajay.stock.unit_price < \$50.00
- 2) ajay.catalog: INDEX PATH
 (1) Index Keys: stock_num manu_code
 (Key-Only) (Serial, fragments: ALL)
 Lower Index Filter:
 ajay.stock.stock_num = ajay.catalog.stock_num

NESTED LOOP JOIN

Related Information

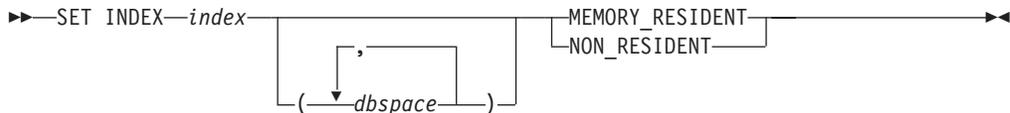
“SET OPTIMIZATION ” on page 2-641, “UPDATE STATISTICS ” on page 2-705

For a description of the EXPLAIN and AVOID_EXECUTE optimizer directives, see “Explain-Mode Directives” on page 5-43. For discussions of SET EXPLAIN and of analyzing the output of the optimizer, see your *IBM Informix Performance Guide*.

SET INDEX

Use the SET INDEX statement to specify that one or more fragments of an index be resident in shared memory as long as possible. Only Extended Parallel Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Element	Description	Restrictions	Syntax
<i>dbspace</i>	Dbpace in which to store the fragment	Must exist	“Identifier ” on page 5-21
<i>index</i>	Index for which to change residency state	Must exist	“Identifier ” on page 5-21

Usage

This statement was formerly supported by Dynamic Server, but it is ignored in current releases. Beginning with Version 9.40, Dynamic Server determines the residency status of indexes and tables automatically.

The SET INDEX statement is a special case of the SET Residency statement. The SET Residency statement can also specify how long a table fragment remains resident in shared memory.

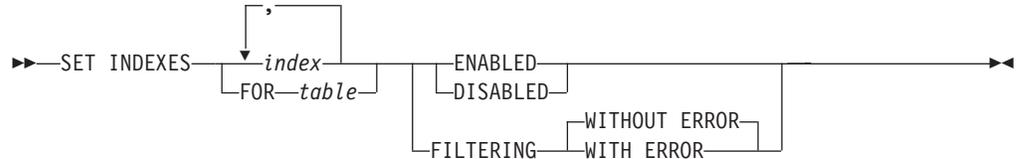
For the complete syntax and semantics of the SET INDEX statement, see “SET Residency ” on page 2-646.

SET INDEXES

Use the SET INDEXES statement to enable or disable an index, or to change the filtering mode of a unique index.

Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL. Do not confuse the SET INDEXES statement with the SET INDEX statement.

Syntax



Element	Description	Restrictions	Syntax
<i>table</i>	Table whose indexes are all to be enabled, disabled, or changed in their filtering mode	Must exist	"Identifier " on page 5-21
<i>index</i>	Index to be enabled, disabled, or changed in its filtering mode	Must exist	"Identifier " on page 5-21

Usage

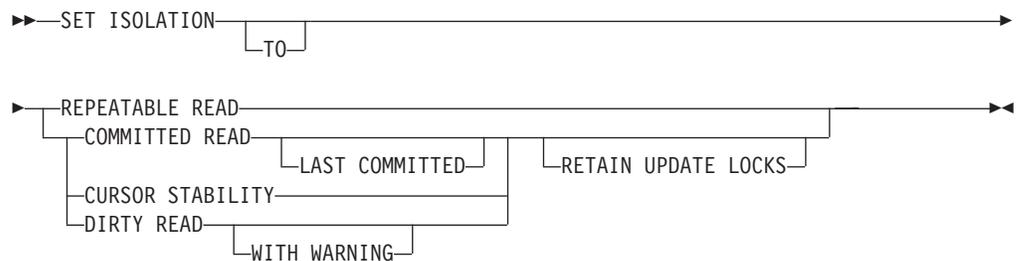
The SET INDEXES statement is a special case of the SET Database Object Mode statement. The SET Database Object Mode statement can also enable or disable a trigger or constraint or change the filtering mode of a constraint.

For the complete syntax and semantics of the SET INDEXES statement, see "SET Database Object Mode" on page 2-591.

SET ISOLATION

Use the SET ISOLATION statement to define the degree of concurrency among processes that attempt to access the same rows simultaneously. This statement is an extension to the ANSI/ISO standard for SQL.

Syntax



Usage

The SET ISOLATION statement is an Informix extension to the ANSI SQL-92 standard. The SET ISOLATION statement can change the enduring isolation level for the session. If you want to set isolation levels through an ANSI-compliant statement, use the SET TRANSACTION statement instead. For a comparison of these two statements, see “SET TRANSACTION ” on page 2-656.

The TO keyword is optional, and has no effect.

SET ISOLATION provides the same functionality as the ISO/ANSI-compliant SET TRANSACTION statement for isolation levels of DIRTY READ (called UNCOMMITTED in SET TRANSACTION), COMMITTED READ, and REPEATABLE READ (called SERIALIZABLE in SET TRANSACTION).

The database *isolation_level* affects read concurrency when rows are retrieved from the database. The isolation level specifies the phenomena that can occur during execution of concurrent SQL transactions. The following phenomena are possible:

- **Dirty Read.** SQL transaction T1 modifies a row. SQL transaction T2 then reads that row before T1 performs a COMMIT. If T1 then performs a ROLLBACK, T2 will have read a row that was never committed, and therefore can be considered never to have existed.
- **Non-Repeatable Read.** SQL transaction T1 reads a row. SQL transaction T2 then modifies or deletes that row and performs a COMMIT. If T1 then attempts to reread that row, T1 might receive the modified value or discover that the row has been deleted.
- **Phantom Row.** SQL transaction T1 reads the set of rows N that satisfy some search condition. SQL transaction T2 then executes SQL statements that generate one or more new rows that satisfy the search condition used by SQL transaction T1. If T1 then repeats the original read with the same search condition, T1 receives a different set of rows.

The database server uses shared locks to support different levels of isolation among processes attempting to access data.

The update or delete process always acquires an exclusive lock on the row that is being modified. The level of isolation does not interfere with rows that you are updating or deleting. If another process attempts to update or delete rows that you are reading with an isolation level of Repeatable Read, that process is denied access to those rows.

In Informix ESQL/C, cursors that are open when SET ISOLATION executes might or might not use the new isolation level when rows are retrieved. Any isolation level that was set from the time the cursor was opened until the application fetches a row might be in effect. The database server might have read rows into internal buffers and internal temporary tables using the isolation level that was in effect at that time. To ensure consistency and reproducible results, close any open cursors before you execute the SET ISOLATION statement.

You can issue the SET ISOLATION statement from a client computer only after a database is opened.

Complete-Connection Level Settings

The SET ISOLATION statement supports *complete-connection level* settings. This means that values in the local session environment at the time of connection are propagated to all new or resumed transactions. These can include the following types of transactions:

- transactions within the local database,
- distributed transactions across databases of the same server instance,
- distributed transactions across databases of two or more database server instances,
- global transactions with XA-compliant data sources that are registered in the local database.

If you change the isolation level within a transaction, the new value is propagated back to the local environment and also to all subsequent new or resumed transactions.

Informix Isolation Levels

The following definitions explain the critical characteristics of each isolation level, from the lowest level of isolation to the highest.

Using the Dirty Read Isolation Level

Use the Dirty Read option to copy rows from the database whether or not there are locks on them. The program that fetches a row places no locks and it respects none. Dirty Read is the only isolation level available to databases that do not implement transaction logging.

This isolation level is most appropriate for static tables that are used for queries of tables where data is not being modified, because it provides no isolation. With Dirty Read, the program might return an uncommitted row that was inserted or modified within a transaction that has subsequently rolled back, or a *phantom row* that was not visible when you first read the query set, but that materializes in the query set before a subsequent read within the same transaction. (Only the Repeatable Read isolation level prevents access to phantom rows. Only Dirty Read provides access to uncommitted rows from concurrent transactions that might subsequently be rolled back.)

The optional WITH WARNING keywords instruct the database server to issue a warning when DML operations that use the Dirty Read isolation level might return an uncommitted row or a phantom row. The transaction in the following example uses this isolation level:

```
BEGIN WORK;  
SET ISOLATION TO DIRTY READ WITH WARNING;  
...  
COMMIT WORK;
```

The Dirty Read isolation level is sensitive to the current setting of the USELASTCOMMITTED configuration parameter and of the USELASTCOMMITTED session environment variable. For information about the behavior of the Dirty Read isolation level when either of these are set to DIRTY READ or to ALL, see “The LAST COMMITTED Option to Committed Read” on page 2-634.

When you use High Availability Data Replication, the database server effectively uses Dirty Read isolation on the HDR Secondary Server, regardless of the specified SET ISOLATION or SET TRANSACTION isolation level.

Using the Committed Read Isolation Level

Use the Committed Read option to guarantee that every retrieved row is committed in the table at the time that the row is retrieved. This option does not place a lock on the fetched row. Committed Read is the default level of isolation in a database with logging that is not ANSI compliant.

Committed Read is appropriate when each row is processed as an independent unit, without reference to other rows in the same table or in other tables.

The LAST COMMITTED Option to Committed Read: Use the LAST COMMITTED keyword option of the Committed Read isolation level to reduce the risk of exclusive row-level locks held by other sessions either causing applications to fail with locking errors, or preventing applications from reading a locked row until after a concurrent transaction is committed or rolled back. In contexts where an application attempts to read a row on which another session holds an exclusive lock, these keywords instruct the database server to return the most recently committed version of the row, rather than wait for the lock to be released.

This feature takes effect implicitly in all user sessions that use the Committed Read isolation level of the SET ISOLATION statement, or that use the Read Committed isolation level of the ANSI/ISO-compliant SET TRANSACTION statement, under any of the following circumstances:

- if the USELASTCOMMITTED configuration parameter is set to 'COMMITTED READ' or to 'ALL'
- if the SET ENVIRONMENT statement sets the USELASTCOMMITTED session environment variable to 'COMMITTED READ' or to 'ALL'.

This feature also takes effect implicitly in all user sessions that use the Dirty Read isolation level of the SET ISOLATION statement, or that use the Read Uncommitted isolation level of the ANSI/ISO-compliant SET TRANSACTION statement, under any of the following circumstances:

- if the USELASTCOMMITTED configuration parameter is set to 'DIRTY READ' or to 'ALL'
- if the SET ENVIRONMENT statement sets the USELASTCOMMITTED session environment variable to 'DIRTY READ' or to 'ALL'.

Enabling this feature cannot eliminate the possibility of locking conflicts, but they reduce the number of scenarios in which other sessions reading the same row can cause an error. The LAST COMMITTED keywords are only effective with concurrent read operations. They cannot prevent locking conflicts or errors that can occur when concurrent sessions attempt to write to the same row.

This feature has no effect on Committed Read or Dirty Read behavior in contexts where no "last committed" version of the table is available, including these:

- The database does not support transaction logging
- The table was created with the LOCK MODE PAGE keywords, or has been altered to have a locking mode of PAGE
- The IFX_DEF_TABLE_LOCKMODE environment variable is set to 'PAGE'
- The DEF_TABLE_LOCKMODE configuration parameter is set to 'PAGE'
- The LOCK TABLE statement has explicitly set an exclusive lock on the table
- An uncommitted DDL statement has implicitly set an exclusive lock on the table
- The table is a system catalog table on which an uncommitted DDL statement has implicitly set an exclusive lock

- The table has columns of complex data types or of user-defined data types
- The table is a RAW table
- A DataBlade module is accessing the table
- The table was created using the Virtual Table Interface.

User-defined access methods are not required to support the LAST COMMITTED feature.

The scope of LAST COMMITTED semantics is neither statement-based nor transaction-based. This isolation level has the same instant-in-time scope that the Committed Read isolation level has without the LAST COMMITTED option. For example, when a query is executed twice within a single transaction with LAST COMMITTED in effect, different results might be returned by the same query, if other DML transactions that were operating on the same data are committed in the interval between the two submissions of the query. This instantaneous nature of the semantics of Committed Read and of Committed Read Last Committed exactly implements the ANSI/ISO Read Committed isolation level.

Using the Cursor Stability Isolation Level

Use the Cursor Stability option to place a shared lock on the fetched row, which is released when you fetch another row or close the cursor. Another process can also place a shared lock on the same row, but no process can acquire an exclusive lock to modify data in the row. Such row stability is important when the program updates another table based on the data it reads from the row.

If you set the isolation level to Cursor Stability, but you are not using a transaction, the Cursor Stability acts like the Committed Read isolation level.

Using the Repeatable Read Isolation Level

Use the Repeatable Read option to place a shared lock on every row that is selected during the transaction. Another process can also place a shared lock on a selected row, but no other process can modify any selected row during your transaction, nor insert a row that meets the search criteria of your query during your transaction. If you repeat the query during the transaction, you reread the same information. The shared locks are released only when the transaction commits or rolls back. Repeatable Read is the default isolation level in an ANSI-compliant database.

Repeatable Read isolation places the largest number of locks and holds them the longest. Therefore, it is the level that reduces concurrency the most.

Default Isolation Levels

The default isolation level for a particular database is established when you create the database according to database type. The following list describes the default isolation level for each database type.

Isolation Level

Database Type

Dirty Read

Default level in a database without logging

Committed Read

Default level in a logged database that is not ANSI compliant

Repeatable Read

Default level in an ANSI-compliant database

The default level remains in effect until you issue a SET ISOLATION statement. After a SET ISOLATION statement executes, the new isolation level remains in effect until one of the following events occurs:

- You enter another SET ISOLATION statement.
- You open another database that has a default isolation level different from the level that your last SET ISOLATION statement specified.
- The program ends.

For a Dynamic Server database that is not ANSI-compliant, unless you explicitly set the USELASTCOMMITTED configuration parameter, the LAST COMMITTED feature is not in effect for the default isolation levels. The SET ENVIRONMENT statement or the SET ISOLATION statement can override this default and enable LAST COMMITTED for the current session.

Using the RETAIN UPDATE LOCKS Option

Use the RETAIN UPDATE LOCKS option to affect the behavior of the database server when it handles a SELECT ... FOR UPDATE statement.

In a database with the isolation level set to Dirty Read, Committed Read, or Cursor Stability, the database server places an update lock on a fetched row of a SELECT ... FOR UPDATE statement. When you turn on the RETAIN UPDATE LOCKS option, the database server retains the update lock until the end of the transaction rather than releasing it at the next subsequent FETCH or when the cursor is closed. This option prevents other users from placing an exclusive lock on the updated row before the current user reaches the end of the transaction.

You can use this option to achieve the same locking effects but avoid the overhead of dummy updates or the repeatable read isolation level.

You can turn this option on or off at any time during the current session.

You can turn the option off by resetting the isolation level without using the RETAIN UPDATE LOCKS keywords, as in the following example.

```
BEGIN WORK;  
SET ISOLATION TO  
    COMMITTED READ LAST COMMITTED RETAIN UPDATE LOCKS;  
...  
COMMIT WORK;  
BEGIN WORK;  
SET ISOLATION TO COMMITTED READ LAST COMMITTED ;  
...  
COMMIT WORK;
```

For more information on update locks, see “Locking Considerations” on page 2-692.

Turning the Option OFF During a Transaction: If you set the RETAIN UPDATE LOCKS option to OFF after a transaction has begun, but before the transaction has been committed or rolled back, several update locks might still exist.

Switching OFF the feature does not directly release any update lock. When you turn this option off, the database server reverts to normal behavior for the three isolation levels. That is, a FETCH statement releases the update lock placed on a row by the immediately preceding FETCH statement, and a closed cursor releases the update lock on the current row.

Update locks placed by earlier FETCH statements are not released unless multiple update cursors are present within the same transaction. In this case, a subsequent FETCH could also release older update locks of other cursors.

Effects of Isolation Levels

You cannot set the transaction isolation level in a database that does not have logging. Every retrieval in such a database occurs as a Dirty Read.

The data retrieved from a BYTE or TEXT column can vary, depending on the transaction isolation level. Under Dirty Read or Committed Read levels of isolation, a process can read a BYTE or TEXT column that is either deleted (if the delete is not yet committed) or in the process of being deleted. Under these isolation levels, deleted data is readable under certain conditions. For information about these conditions, see the *IBM Informix Administrator's Guide*.

When you use DB–Access, as you use higher levels of isolation, lock conflicts occur more frequently. For example, if you use Cursor Stability, more lock conflicts occur than if you use Committed Read.

Using a scroll cursor in an Informix ESQL/C transaction, you can force consistency between your temporary table and the database table either by setting the level to Repeatable Read or by locking the entire table during the transaction.

If you use a scroll cursor WITH HOLD in a transaction, you cannot force consistency between your temporary table and the database table. A table-level lock or locks that are set by Repeatable Read are released when the transaction is completed, but the scroll cursor with hold remains open beyond the end of the transaction. You can modify released rows as soon as the transaction ends, but retrieved data in the temporary table might be inconsistent with the actual data.

Attention: Do not use nonlogging tables within a transaction. If you need to use a nonlogging table within a transaction, either set the isolation level to Repeatable Read or else lock the table in Exclusive mode to prevent concurrency problems.

Related Information

Related statements: “CREATE DATABASE ” on page 2-97, “SET LOCK MODE,” and “SET TRANSACTION ” on page 2-656

For a discussion of how to set the transaction isolation level, see the *IBM Informix Guide to SQL: Tutorial*.

SET LOCK MODE

Use the SET LOCK MODE statement to define how the database server handles a process that tries to access a locked row or table.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax

```
►► SET LOCK MODE TO { NOT WAIT | WAIT } [seconds] ►►
```

Element	Description	Restrictions	Syntax
<i>seconds</i>	Maximum number of seconds that a process waits for a lock to be released before issuing an error	Valid only if shorter than system default	"Literal Number " on page 4-185

Usage

This statement can direct the response of the database server in the following ways when a process tries to access a locked row or table.

Lock Mode

Effect

NOT WAIT

Database server ends the operation immediately and returns an error code. This condition is the default.

WAIT Database server suspends the process until the lock releases.

WAIT *seconds*

Database server suspends the process until the lock releases or until the waiting period ends. If the lock remains after the waiting period, the operation ends and an error code is returned.

In the following example, the user specifies that if the process requests a locked row, the operation should end immediately and an error code should be returned:

```
SET LOCK MODE TO NOT WAIT;
```

In the following example, the user specifies that the process should be suspended until the lock is released:

```
SET LOCK MODE TO WAIT;
```

The next example sets an upper limit of 17 seconds on the length of any wait:

```
SET LOCK MODE TO WAIT 17;
```

WAIT Clause

The WAIT clause causes the database server to suspend the process until the lock is released or until a specified number of seconds have passed without the lock being released.

The database server protects against the possibility of a deadlock when you request the WAIT option. Before the database server suspends a process, it checks whether suspending the process could create a deadlock. If the database server discovers that a deadlock could occur, it ends the operation (overruling your instruction to wait) and returns an error code. In the case of either a suspected or actual deadlock, the database server returns an error.

Cautiously use the unlimited waiting period that was created when you specify the WAIT option without *seconds*. If you do not specify an upper limit, and the process that placed the lock somehow fails to release it, suspended processes could wait indefinitely. Because a true deadlock situation does not exist, the database server does not take corrective action.

In a network environment, the DBA uses the ONCONFIG parameter DEADLOCK_TIMEOUT to establish a default value for *seconds*. If you use a SET

LOCK MODE statement to set an upper limit, your value applies only when your waiting period is shorter than the system default.

Complete-Connection Level Settings

The SET LOCK MODE statement supports *complete-connection level* settings. This means that values in the local session environment at the time of connection are propagated to all new or resumed transactions. These can include the following types of transactions:

- transactions within the local database,
- distributed transactions across databases of the same server instance,
- distributed transactions across databases of two or more database server instances,
- global transactions with XA-compliant data sources that are registered in the local database.

If you change the lock mode setting within a transaction, the new value is propagated back to the local environment and also to all subsequent new or resumed transactions.

In releases of Dynamic Server earlier than 9.40.UC8, the SET LOCK MODE statement did not support complete-connection level settings. The process waited for the specified number of seconds only if you acquired locks within the current database server and a remote database server within the same transaction.

Related Information

Related statements: “LOCK TABLE ” on page 2-449, “SET ISOLATION” on page 2-631, “SET TRANSACTION ” on page 2-656, and “UNLOCK TABLE ” on page 2-688

For a description of the two distinct meanings of the term *lock mode* in this document, see the section “Locking Granularity” on page 2-452.

For a discussion of how to set the lock mode behavior of the database server, see the *IBM Informix Guide to SQL: Tutorial*.

To avoid waiting in operations that attempt to read rows on which concurrent sessions hold exclusive row-level locks, you can also use the LAST COMMITTED feature, either by setting it explicitly in the SET ISOLATION COMMITTED READ statement, or by setting the USELASTCOMMITTED configuration parameter or the USELASTCOMMITTED session environment option. For more information about the LAST COMMITTED feature, see the section “Using the Committed Read Isolation Level” on page 2-634.

SET LOG

Use the SET LOG statement to change your database logging mode from buffered transaction logging to unbuffered transaction logging or vice versa.

This statement is an extension to the ANSI/ISO standard for SQL. Unlike most extensions, the SET LOG statement is not valid in an ANSI-compliant database.

Syntax



Usage

You activate transaction logging when you create a database or add logging to an existing database. These transaction logs can be buffered or unbuffered.

Buffered logging is a type of logging that holds transactions in a memory buffer until the buffer is full, regardless of when the transaction is committed or rolled back. The database server provides this option to speed up operations by reducing the number of disk writes.

Attention: You gain a marginal increase in efficiency with buffered logging, but you incur some risk. In the event of a system failure, the database server cannot recover any completed transactions in the memory buffer that were not written to disk.

The SET LOG statement in the following example changes the transaction logging mode to buffered logging:

```
SET BUFFERED LOG;
```

Unbuffered logging is a type of logging that does not hold transactions in a memory buffer. As soon as a transaction ends, the database server writes the transaction to disk. If a system failure occurs when you are using unbuffered logging, you recover all completed transactions, but not those still in the buffer. The default condition for transaction logs is unbuffered logging.

The SET LOG statement in the following example changes the transaction logging mode to unbuffered logging:

```
SET LOG;
```

The SET LOG statement redefines the mode for the current session only. The default mode, which the database administrator sets with the **ondblog** utility, remains unchanged.

The buffering option does not affect retrievals from external tables. For distributed queries, a database with logging can retrieve only from databases with logging, but it makes no difference whether the databases use buffered or unbuffered logging.

An ANSI-compliant database cannot use buffered logging.

You cannot change the logging mode of ANSI-compliant databases. If you created a database with the WITH LOG MODE ANSI keywords, you cannot later use the SET LOG statement to change the logging mode to buffered or unbuffered transaction logging.

Related Information

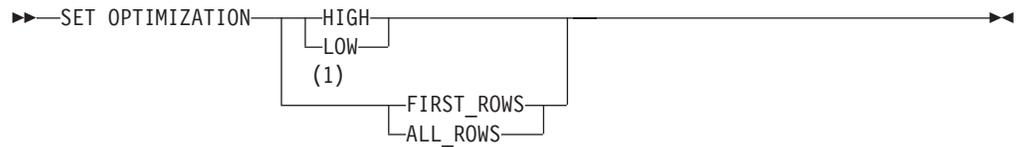
Related statement: "CREATE DATABASE " on page 2-97

SET OPTIMIZATION

Use the SET OPTIMIZATION statement to specify how much time the optimizer spends developing a query plan or specifying optimization goals.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 Dynamic Server only

Usage

You can execute a SET OPTIMIZATION statement at any time. The specified optimization level carries across databases on the current database server. The option that you specify remains in effect until you issue another SET OPTIMIZATION statement or until the program ends. The default database server optimization level for the amount of time that the query optimizer spends determining the query plan is HIGH.

On Dynamic Server, the default optimization goal is ALL_ROWS. Although you can set only one option at a time, you can issue two SET OPTIMIZATION statements: one that specifies the time the optimizer spends to determine the query plan and one that specifies the optimization goal of the query.

HIGH and LOW Options

The HIGH and LOW options determine how much time the query optimizer spends to determine the query plan:

- HIGH

This option directs the optimizer to use a sophisticated cost-based algorithm that examines all reasonable query-plan choices and selects the best overall alternative.

For large joins, this algorithm can incur more overhead than you desire. In extreme cases, you can run out of memory.

- LOW

This option directs the optimizer to use a less sophisticated but faster to design optimization algorithm, based on the lowest-cost path at each stage. This algorithm eliminates unlikely join strategies during the early stages of optimization and reduces the time and resources spent during optimization.

When you specify the LOW level of optimization, the database server might not select the optimal strategy because that strategy was eliminated from consideration during the early stages of the algorithm.

FIRST_ROWS and ALL_ROWS Options (IDS)

The FIRST_ROWS and ALL_ROWS options relate to the optimization goal of the query:

- FIRST_ROWS

This option directs the optimizer to choose the query plan that returns the first result record as soon as possible, ignoring plans that would sort records or create a hash table.

- ALL_ROWS

This option directs the optimizer to choose the query plan that returns all the records as quickly as possible.

You can also specify the optimization goal of a specific query with the optimization-goal directive. For more information, see “Optimizer Directives” on page 5-35.

Optimizing SPL Routines

For SPL routines that remain unchanged or change only slightly, you might want to set the SET OPTIMIZATION statement to HIGH when you create the SPL routine. This step stores the best query plans for the SPL routine. Then execute a SET OPTIMIZATION LOW statement before you execute the SPL routine. The SPL routine then uses the optimal query plans and runs at the more cost-effective rate.

Examples

The following example shows optimization across a network. The **central** database (on the **midstate** database server) is to have LOW optimization; the **western** database (on the **rockies** database server) is to have HIGH optimization.

```
CONNECT TO 'central@midstate';
SET OPTIMIZATION LOW;
SELECT * FROM customer;
CLOSE DATABASE;
CONNECT TO 'western@rockies';
SET OPTIMIZATION HIGH;
SELECT * FROM customer;
CLOSE DATABASE;
CONNECT TO 'wyoming@rockies';
SELECT * FROM customer;
```

Here the **wyoming** database is to have HIGH optimization because it resides on the same database server as the **western** database. The code does not need to re-specify the optimization level for the **wyoming** database because the **wyoming** database resides on the **rockies** database server like the **western** database.

The following example directs the Dynamic Server optimizer to use the most time to determine a query plan, and to then return the first rows of the result as soon as possible:

```
SET OPTIMIZATION LOW;
SET OPTIMIZATION FIRST_ROWS;
SELECT lname, fname, bonus
   FROM sales_emp, sales
  WHERE sales.empid = sales_emp.empid AND bonus > 5,000
     ORDER BY bonus DESC;
```

Related Information

Related statements: “SET EXPLAIN ” on page 2-623, “SET ENVIRONMENT ” on page 2-615, and “UPDATE STATISTICS ” on page 2-705

For information on other methods by which you can alter the query plan of the Dynamic Server optimizer, see “ Optimizer Directives ” on page 5-35.

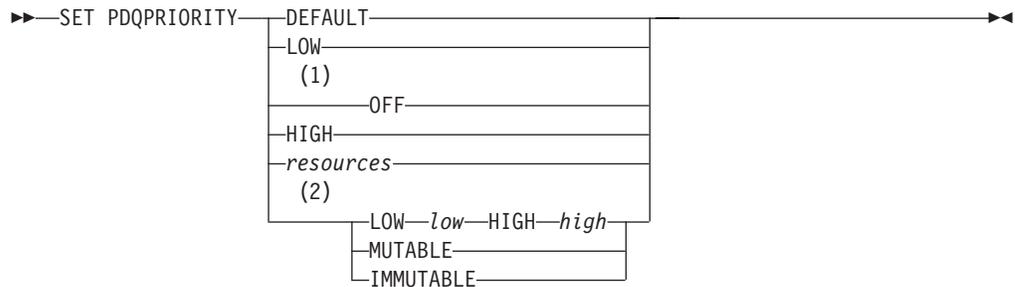
For more information on how to optimize queries, see your *IBM Informix Performance Guide*.

SET PDQPRIORITY

The SET PDQPRIORITY statement enables an application to set the query priority level dynamically within a routine.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 Dynamic Server only
- 2 Extended Parallel Server only

Element	Description	Restrictions	Syntax
<i>high</i>	Integer value that specifies the desired resource allocation	Must be in the range 1 to 100, and not less than the <i>low</i> value	“Literal Number ” on page 4-185
<i>low</i>	Integer that specifies the minimum resource allocation	Must be in the range 1 to 100, and not greater than the <i>high</i> value	“Literal Number ” on page 4-185
<i>resources</i>	Integer that specifies the query priority level and the percent of resources to process the query	Can range from -1 to 100. See also “Allocating Database Server Resources” on page 2-644.	“Literal Number ” on page 4-185

Usage

The SET PDQPRIORITY statement overrides the **PDQPRIORITY** environment variable (but has lower precedence than the MAX_PDQPRIORITY configuration parameter). The scope of SET PDQPRIORITY is local to the routine, and does not affect other routines within the same session. When a routine that issues this statement terminates, the setting reverts to the system default value.

The SET PDQPRIORITY statement is not supported in SPL routines.

In Dynamic Server, set PDQ priority to a value less than the quotient of 100 divided by the maximum number of prepared statements. For example, if two prepared statements are active, you should set the PDQ priority to less than 50.

In Extended Parallel Server, you can use SET PDQPRIORITY to set PDQ priority at runtime to a value greater than 0 when you need more memory for operations such as sorts, forming groups, and index builds. For guidelines on which values to use, see your *IBM Informix Performance Guide*.

For example, assume that the DBA sets the MAX_PDQPRIORITY parameter to 50. Then a user enters the following SET PDQPRIORITY statement to set the query priority level to 80 percent of resources:

```
SET PDQPRIORITY 80;
```

When it processes the query, the database server uses the MAX_PDQPRIORITY value to factor the query priority level set by the user. The database server silently processes the query with a priority level of 40. This priority level represents 50 percent of the 80 percent of resources that the user specifies.

The following keywords are supported by the SET PDQPRIORITY statement.

Keyword

Effect

DEFAULT

Uses the setting of the **PDQPRIORITY** environment variable

LOW Data values are fetched from fragmented tables in parallel. (In Dynamic Server, when you specify **LOW**, the database server uses no other forms of parallelism.)

OFF PDQ is turned off (Dynamic Server only). The database server uses no parallelism. **OFF** is the default if you use neither the **PDQPRIORITY** environment variable nor the SET PDQPRIORITY statement.

HIGH The database server determines an appropriate **PDQPRIORITY** value, based on factors that include the number of available processors, the fragmentation of the tables being queried, the complexity of the query, and others. IBM reserves the right to change the performance behavior of queries when **HIGH** is specified in future releases.

MUTABLE

PDQPRIORITY can be changed in a user session by the user.

IMMUTABLE

PDQPRIORITY cannot be changed in a user session.

The DBA typically specifies the **MUTABLE** or **IMMUTABLE** setting, which only Extended Parallel Server supports, in a **sysdbopen()** procedure.

Only Extended Parallel Server supports the **HIGH** and **LOW** keywords in the same statement, each followed by a value, to specify a range of resource levels. For more information, see "Using a Range of Values" on page 2-645.

Allocating Database Server Resources

You can specify an integer in the range from -1 to 100 to indicate a query priority level as the percent of database server resources to process the query. Resources include the amount of memory and the number of processors. The higher the number you specify, the more resources the database server uses.

Use of more resources usually indicates better performance for a given query. Using excessive resources, however, can cause contention for resources and remove resources from other queries, so that degraded performance results. With the *resources* option, the following values are numeric equivalents of the keywords that indicate query priority level.

Value Equivalent Keyword-Priority Level

- 1 DEFAULT
- 0 OFF (Dynamic Server only)
- 1 LOW

For Dynamic Server, the following statements are equivalent. The first statement uses the keyword LOW to establish a low query-priority level. The second uses a value of 1 in the *resources* parameter to establish a low query-priority level.

```
SET PDQPRIORITY LOW;
```

```
SET PDQPRIORITY 1;
```

Using a Range of Values

In Extended Parallel Server, when you specify a range of values in SET PDQPRIORITY, you allow the Resource Grant Manager (RGM) some discretion in allocating resources.

The *high* value in the range is the desired resource allocation, while the *low* value is the minimum acceptable resource allocation for the query. If the *low* value exceeds the available system resources, the RGM blocks the query. Otherwise, the RGM chooses the largest PDQ priority in the specified range that does not exceed available resources.

Related Information

For information about configuration parameters and about the Resource Grant Manager, see your *IBM Informix Administrator's Guide* and your *IBM Informix Performance Guide*.

For information about the PDQPRIORITY environment variable, see the *IBM Informix Guide to SQL: Reference*.

SET PLOAD FILE

Use the SET PLOAD FILE statement to prepare a log file for a session of loading or unloading data from or to an external table. The log file records summary statistics about each load or unload job. The log file also lists any reject files created during a load job. Only Extended Parallel Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```
➤ SET PLOAD FILE TO filename [WITH APPEND] ➤
```

Element	Description	Restrictions	Syntax
<i>filename</i>	Name of the log file. If you specify no <i>filename</i> , then log information is written to /dev/null .	If the file cannot be opened for writing, an error results	Platform-dependent

Usage

Each time a session closes, the log file for that session also closes. If you issue more than one SET PLOAD FILE statement within a session, each new statement closes a previously opened log file and opens a new log file.

The WITH APPEND option allows you to append new log information to the existing log file.

The following example specifies a file and includes the WITH APPEND option:

```
SET PLOAD FILE TO '/tmp/pload.out' WITH APPEND;
```

If you invoke a SET PLOAD FILE statement with a simple filename on a local database, the output file is located in your current directory. If your current database is on a remote database server, then the output file is located in your home directory on the remote database server, on the coserver where the initial connection was made. If you provide a full pathname for the file, it is placed in the directory and file specified on the remote server.

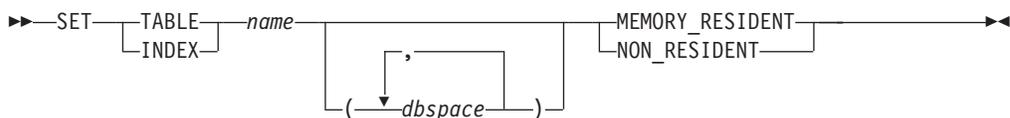
Related Information

Related statement: "CREATE EXTERNAL TABLE (XPS) " on page 2-104

SET Residency

Use the SET Residency statement to specify that one or more fragments of a table or index be resident in shared memory as long as possible. Only Extended Parallel Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Element	Description	Restrictions	Syntax
<i>dbspace</i>	The dbspace where fragment resides	The dbspace must exist	"Identifier " on page 5-21
<i>name</i>	Table or index for which the residency state will be changed	Table or index must exist	"Identifier " on page 5-21

Usage

This statement was formerly supported by Dynamic Server, but it is ignored in current releases. Beginning with Version 9.40, Dynamic Server determines the residency status of indexes and tables automatically.

The SET Residency statement allows you to specify the tables, indexes, and data fragments that you want to remain in the buffer as long as possible. When a free buffer is requested, pages that are declared with the MEMORY_RESIDENT keyword are considered last for page replacement.

The default state is nonresident. The residency state is persistent while the database server is up. That is, each time the database server is started, you must specify the database objects that you want to remain in shared memory.

After a table, index, or data fragment is set to MEMORY_RESIDENT, the residency state remains in effect until one of the following events occurs:

- You use SET Residency to set the database object to NON_RESIDENT.
- The database object is dropped.
- The database server is taken offline.

Only user **informix** can set or change the residency state of a database object.

Residency and the Changing Status of Fragments

If new fragments are added to a resident table, the fragments are not marked automatically as resident. You must issue the SET Residency statement for each new fragment or reissue the statement for the entire table.

Similarly, if a resident fragment is detached from a table, the residency state of the fragment remains unchanged. If you want the residency state to change to nonresident, you must issue the SET Residency statement to declare the specific fragment (or the entire table) as nonresident.

Examples

The next example shows how to set the residency status of an entire table:

```
SET TABLE tab1 MEMORY_RESIDENT;
```

For fragmented tables or indexes, you can specify residency for individual fragments as the following example shows:

```
SET INDEX index1 (dbspace1, dbspace2) MEMORY_RESIDENT;  
SET TABLE tab1 (dbspace1) NON_RESIDENT;
```

This example specifies that the **tab1** fragment in **dbspace1** is not to remain in shared memory while the **index1** fragments in **dbspace1** and **dbspace2** are to remain in shared memory as long as possible.

Related Information

Related statement: "ALTER FRAGMENT " on page 2-7

For information on how to monitor the residency status of tables, indexes, and fragments, refer to your *IBM Informix Administrator's Guide*.

SET ROLE

Use the SET ROLE statement to enable the privileges of a user-defined role. This statement is an extension to the ANSI/ISO standard for SQL.

Syntax



Element	Description	Restrictions	Syntax
<i>role</i>	Name of a role to be enabled	Must already exist in the database and must already have been granted to the user, but cannot be a built-in role. If enclosed between quotation marks, <i>role</i> is case sensitive.	" Owner Name "

Usage

Any user who is granted a role can enable the role by using the SET ROLE statement. You can only enable one role at a time. If you execute the SET ROLE statement after a role is already set, the new role replaces the old role as the current role.

The SET ROLE statement returns an error if the user does not currently hold the role, or if the role is a built-in role. (The access privileges held by a built-in role, such as the EXTEND role or the DBSECADM role, are always in effect, and do not require activation by the SET ROLE statement if the user holds that role.)

Users can be granted a default role for the database instance when the DBA issues the GRANT DEFAULT ROLE statement. If no default role exists for the user in the current database, role NULL or NONE is assigned by default. In this context, NULL and NONE are synonyms. Roles NULL and NONE can have no privileges. To set your role to NULL or NONE disables your current role.

When you use SET ROLE to enable a role, you gain the privileges of the role, in addition to the privileges of PUBLIC and your own privileges. If a role is granted to another role that has been assigned to you, you gain the privileges of both roles, in addition to any privileges of PUBLIC and your own privileges.

After SET ROLE executes successfully, the specified role remains effective until the current database is closed or the user executes another SET ROLE statement. Only the user, however, not the role, retains ownership of any database objects, such as tables, that were created during the session.

A role is in scope only within the current database. You cannot use privileges that you acquire from a role to access data in another database. For example, if you have privileges from a role in the database named **acctg**, and you execute a distributed query over the databases named **acctg** and **inventory**, your query cannot access the data in the **inventory** database unless you were also granted appropriate privileges in the **inventory** database.

If your database supports explicit transactions, you must issue the SET ROLE statement outside a transaction. If your database is ANSI-compliant, SET ROLE must be the first statement of a new transaction. If the SET ROLE statement is executed while a transaction is active, an error occurs. For more information about SQL statements that initiate an implicit transaction, see "SET SESSION AUTHORIZATION and Transactions" on page 2-652.

If the SET ROLE statement is executed as a part of a trigger or SPL routine, and the owner of the trigger or SPL routine was granted the role with the WITH GRANT OPTION, the role is enabled even if you are not granted the role. For example, this code fragment sets a role and then relinquishes it after a query:

```
EXEC SQL set role engineer;
EXEC SQL select fname, lname, project
        INTO :efname, :elname, :eproject FROM projects
        WHERE project_num > 100 AND lname = 'Larkin';
printf ("%s is working on %s\n", efname, eproject);
EXEC SQL set role NULL;
```

Setting the Default Role

The DBA or the owner of the database can issue the GRANT DEFAULT ROLE statement to assign an existing role as the *default role* to a specified list of users or to PUBLIC. Unlike a non-default role, the default role takes effect automatically when the user connects to the database, and does not require the SET ROLE statement to enable it. Each of the three statements in next example respectively performs one of the following operations on a role:

- Declares a role called **Engineer**
- Assigns Select privileges on a the **locomotives** table to the **Engineer** role
- Defines **Engineer** as the default role for the user **jgould**.

```
EXEC SQL CREATE ROLE 'Engineer';
EXEC SQL GRANT SELECT ON locomotives TO 'Engineer';
EXEC SQL GRANT DEFAULT ROLE 'Engineer' TO jgould;
```

If **jgould** subsequently uses the SET ROLE statement to enable some other role, then by executing the following statement, **jgould** replaces that role with **Engineer** as the current role:

```
SET ROLE DEFAULT;
```

If you have no default role, SET ROLE DEFAULT makes NONE your current role, leaving only the privileges that have been granted explicitly to your *username* or to PUBLIC. After GRANT DEFAULT ROLE changes your default role to a new default role, executing SET ROLE DEFAULT restores your most recently granted default role, even if this role was not your default role when you connected to the database.

If one default role is granted to PUBLIC, but a different role is granted as the default role to an individual user, the individually-granted default role takes precedence if that user issues SET ROLE DEFAULT or connects to the database.

Related Information

Related statements: “CREATE ROLE ” on page 2-164, “DROP ROLE ” on page 2-334, “GRANT ” on page 2-395, and “REVOKE” on page 2-494

For a discussion of how to use roles, see the *IBM Informix Guide to SQL: Tutorial*.

SET SCHEDULE LEVEL

The SET SCHEDULE LEVEL statement specifies the scheduling level of a query when queries are waiting to be processed. Only Extended Parallel Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

▶▶—SET SCHEDULE LEVEL—*level*—▶▶

Element	Description	Restrictions	Syntax
<i>level</i>	Integer value that specifies the scheduling priority of a query	Must be between 1 and 100. If the value falls outside the range of 1 and 100, the database server uses the default value of 50.	“Literal Number ” on page 4-185

Usage

The highest priority level is 100. That is, a query at level 100 is more important than a query at level 1. In general, the Resource Grant Manager (RGM) processes a query with a higher scheduling level before a query with a lower scheduling level. The exact behavior of the RGM is also influenced by the setting of the DS_ADM_POLICY configuration parameter.

Examples

```
SET SCHEDULE LEVEL 1;
```

This statement only works for XPS.

Related Information

Related statement: “SET PDQPRIORITY ” on page 2-643

For information about the Resource Grant Manager, see your *IBM Informix Administrator's Guide*.

For information about the DS_ADM_POLICY configuration parameter, see your *IBM Informix Administrator's Reference*.

SET SESSION AUTHORIZATION

The SET SESSION AUTHORIZATION statement lets you change the user name under which database operations are performed in the current session. Only Dynamic Server supports this statement.

Syntax

▶▶—SET SESSION AUTHORIZATION TO—'*user*'—▶▶

Element	Description	Restrictions	Syntax
<i>user</i>	User name by which database operations will be performed in the current session	Must be a valid user name. Delimiters (') are optional	“ Owner Name ” on page 5-45

Usage

This statement allows you to assume the identity of another user, including their DAC and LBAC credentials.

Both the DBA and SETSESSIONAUTH access privilege are required to execute this statement. Unless when you start the session you already hold the SETSESSIONAUTH privilege for PUBLIC (or for the user whose name you specify in the SET SESSION AUTHORIZATION statement), and you also hold the DBA privilege, this statement fails with an error.

If the database server has been converted from a legacy version that did not support label-based access control, users who held the DBA privilege are automatically granted the SETSESSIONAUTH access privilege for PUBLIC in the migration process. If the database server has been initialized as a version that supports LBAC security policies, users who hold the DBSECADM role can grant the SETSESSIONAUTH privilege to other users. Because the security credentials of each user determine what data rows can be accessed in protected tables, the DBSECADM should exercise care in granting the SETSESSIONAUTH privilege and in specifying its scope.

The new identity remains in effect in the current database until you execute SET SESSION AUTHORIZATION again, or until you close the current database. When you use this statement, the specified *user* must have the Connect privilege on the current database. In addition the DBA cannot set the new authorization identifier to PUBLIC, nor to any existing role in the current database.

Setting a session to another user causes a change in a user name in the current active database server. The specified *user*, as far as this database server process is concerned, is completely dispossessed of any privileges while accessing the database server through some administrative utility. Additionally, the new session *user* is not able to initiate any administrative operation (execute a utility, for example) by virtue of the acquired identity.

After the SET SESSION AUTHORIZATION statement successfully executes, any role enabled by a previous user is relinquished. You must use the SET ROLE statement if you wish to assume a role that has been granted to the specified *user*. The database server does not enable the default role of *user* automatically.

After SET SESSION AUTHORIZATION successfully executes, the database server puts any owner-privileged UDRs that the DBA created while using the new authorization identifier in RESTRICTED mode, which can affect access privileges during operations of the UDR on objects in remote databases. For more information on RESTRICTED mode, see the **sysprocedures** system catalog table in the *IBM Informix Guide to SQL: Reference*.

When you assume the identity of another user by executing the SET SESSION AUTHORIZATION statement, you can perform operations in the current database only. You cannot perform an operation on a database object outside the current database, such as a remote table. In addition, you cannot execute a DROP DATABASE or RENAME DATABASE statement, even if the database is owned by the real or effective user.

You can use this statement either to obtain access to the data directly or to grant the database-level or table-level privileges needed for the database operation to proceed. The following example shows how to use the SET SESSION AUTHORIZATION statement to obtain table-level privileges:

```
SET SESSION AUTHORIZATION TO 'cath1';
GRANT ALL ON customer TO mary;
SET SESSION AUTHORIZATION TO 'mary';
UPDATE customer SET fname = 'Carl' WHERE lname = 'Pauli';
```

If you enclose *user* in quotation marks, the name is case sensitive and is stored exactly as you typed it. In an ANSI-compliant database, if you do not use quotation marks as delimiters, the name is stored in uppercase letters.

SET SESSION AUTHORIZATION and Transactions

If your database is not ANSI compliant, you must issue the SET SESSION AUTHORIZATION statement outside a transaction. If you issue the statement within a transaction, you receive an error message.

In an ANSI-compliant database, you can execute the SET SESSION AUTHORIZATION statement only if you have not executed a statement that initiates an implicit transaction (for example, CREATE TABLE or SELECT). Statements that do not initiate an implicit transaction are statements that do not acquire locks or log data (for example, SET EXPLAIN and SET ISOLATION). You can execute the SET SESSION AUTHORIZATION statement immediately after a DATABASE statement or a COMMIT WORK statement.

Related Information

Related statements: "CONNECT " on page 2-83, "DATABASE " on page 2-281, "GRANT " on page 2-395, "REVOKE" on page 2-494, and "SET ROLE " on page 2-647

SET STATEMENT CACHE

Use the SET STATEMENT CACHE statement to turn caching on or off for the current session. Only Dynamic Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax

```
▶▶ SET STATEMENT CACHE { ON | OFF } ▶▶
```

Usage

You can use the SET STATEMENT CACHE statement to turn caching in the SQL statement cache ON or OFF for the current session. The statement cache stores in a buffer identical statements that are repeatedly executed in a session. Only data manipulation language (DML) statements (DELETE, INSERT, UPDATE, or SELECT) can be stored in the statement cache.

This mechanism allows qualifying statements to bypass the optimization stage and parsing stage, and avoid recompiling, which can reduce memory consumption and can improve query processing time.

Precedence and Default Behavior

SET STATEMENT CACHE takes precedence over the **STMT_CACHE** environment variable and the STMT_CACHE configuration parameter. You must enable the SQL statement cache, however, either by setting the STMT_CACHE configuration parameter or by using the **onmode** utility, before the SET STATEMENT CACHE statement can execute successfully.

When you issue a `SET STATEMENT CACHE ON` statement, the SQL statement cache remains in effect until you issue a `SET STATEMENT CACHE OFF` statement or until the program ends. If you do not use `SET STATEMENT CACHE`, the default behavior depends on the setting of the `STMT_CACHE` environment variable or the `STMT_CACHE` configuration parameter.

Turning the Cache ON

Use the `ON` option to enable the SQL statement cache. When the SQL statement cache is enabled, each statement that you execute passes through the SQL statement cache to determine if a matching cache entry is present. If so, the database server uses the cached entry to execute the statement.

If the statement has no matching entry, the database server tests to see if it qualifies for entry into the cache. For the conditions a statement must meet to enter into the cache, see “Statement Qualification” on page 2-654.

Restrictions on Matching Entries in the SQL Statement Cache

When the database server considers whether or not a statement is identical to a statement in the SQL statement cache, the following items must match:

- Lettercase
- Comments
- White space
- Optimization settings
 - `SET OPTIMIZATION` statement options
 - Optimizer directives
 - The `SET ENVIRONMENT OPTCOMPIND` statement options or settings of the `OPTCOMPIND` environment variable, or of the `OPTCOMPIND` configuration parameter in the `ONCONFIG` file. (If conflicting settings exist for the same query, this is the descending order of precedence.)
- Parallelism settings
 - `SET PDQPRIORITY` statement options or settings of the `PDQPRIORITY` environment variable
- Query text strings
- Literals

If an SQL statement is semantically equivalent to a statement in the SQL statement cache but has different literals, the statement is not considered identical and qualifies for entry into the cache. For example, the following `SELECT` statements are not identical:

```
SELECT col1, col2 FROM tab1 WHERE col1=3;
SELECT col1, col2 FROM tab1 WHERE col1=5;
```

In this example, both statements are entered into the SQL statement cache.

Host-variable names, however, are insignificant. For example, the following select statements are considered identical:

```
SELECT * FROM tab1 WHERE x = :x AND y = :y;
SELECT * FROM tab1 WHERE x = :p AND y = :q;
```

In the previous example, although the host names are different, the statements qualify, because the case, query text strings, and white space match. Performance does not improve, however, because each statement has already been parsed and optimized by the PREPARE statement.

Turning the Cache OFF

The OFF option disables the SQL statement cache. When you turn caching OFF for your session, no SQL statement cache code is executed for that session.

The SQL statement cache is designed to save memory in environments where identical queries are executed repeatedly and schema changes are infrequent. If this is not the case, you might want to turn the SQL statement cache off to avoid the overhead of caching. For example, if you have little cache cohesion, that is, when relatively few matches but many new entries into the cache exist, the cache management overhead is high. In this case, turn the SQL statement cache off.

If you know that you are executing many statements that do not qualify for the SQL statement cache, you might want to disable it and avoid the overhead of testing to see if each DML statement qualifies for insertion into the cache.

Statement Qualification

A statement that can be cached in the SQL statement cache (and consequently, one that can match a statement that already appears in the SQL statement cache) must meet all of the following conditions:

- It must be a SELECT, INSERT, UPDATE, or DELETE statement.
- It must contain only non-opaque built-in data types (excluding BLOB, BOOLEAN, BYTE, CLOB, LVARCHAR, and TEXT).
- It must contain only built-in operators.
- It cannot contain user-defined routines.
- It cannot contain temporary or remote tables.
- It cannot contain subqueries in the Projection list.
- It cannot be part of a multistatement PREPARE.
- It cannot have user-privilege restrictions on target columns.
- In an ANSI-compliant database, it must contain fully qualified object names.
- It cannot require re-optimization.

Requiring Re-Execution Before Cache Insertion

A qualified SQL statement is fully inserted into the SQL statement cache only after the database server counts a configurable number of references (which are sometimes called "hits") to that statement. For the default value of zero, a qualified DML statement does not need to be re-executed before it is cached.

Using the STMT_CACHE_HITS configuration parameter, however, the database administrator (DBA) can specify that qualified DML statements must be executed a minimum number of times before they are inserted into the statement cache. By setting this to a value of one (or to a larger value), the DBA excludes one-time-only *ad hoc* queries from full insertion into the SQL statement cache, thereby lowering cache-management overhead.

Enabling or Disabling Insertions After Size Exceeds Configured Limit

The DBA can prevent the insertion of additional qualified SQL statements into the statement cache when the cache size reaches its configured size (as specified by the

STMT_CACHE_SIZE configuration parameter) by setting the configuration parameter STMT_CACHE_NOLIMIT to zero.

Prepared Statements and the Statement Cache

Prepared statements are inherently cached for a single session. That is, if a prepared statement is executed many times (or if a single cursor is opened many times), the same prepared query plan is used by that session. If a session prepares a statement and then executes it many times, its performance is essentially unaffected by using the SQL statement cache, because the statement is optimized just once, during the PREPARE statement.

If other sessions also prepare that same statement, however, or if the first session prepares the statement several times, then the statement cache usually provides a direct performance benefit, because the database server only calculates the query plan once. Of course, the original session might gain a (small) benefit from the statement cache, even if it prepares the statement only once, because other sessions use less memory, and the database server does less work for the other sessions.

Examples

The following example turns on statement caching for the current session:

```
SET STATEMENT CACHE ON;
```

The following turns off statement caching for the current session:

```
SET STATEMENT CACHE OFF;
```

Related Information

For information on optimization settings, see “SAVE EXTERNAL DIRECTIVES ” on page 2-522, “SET OPTIMIZATION ” on page 2-641, and “ Optimizer Directives ” on page 5-35.

For information about the **STMT_CACHE** environment variable, see the *IBM Informix Guide to SQL: Reference*.

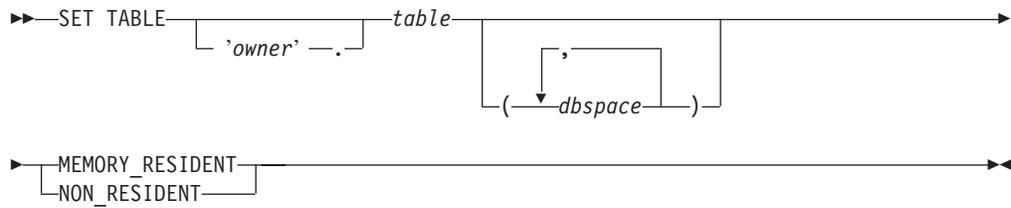
For more information about STMT_CACHE, STMT_CACHE_NUMPOOL, STMT_CACHE_HITS, and other configuration parameters that affect the statement cache, as well as cache-related command-line options of the **onmode** utility, see your *IBM Informix Administrator's Reference*.

For more information on the performance implications of this feature, on when and how to use the SQL statement cache, on how to monitor the cache with the **onstat** options, and on how to tune the configuration parameters, see your *IBM Informix Performance Guide*.

SET TABLE

Use the SET TABLE statement to specify that one or more fragments of a table be resident in shared memory as long as possible. Only Extended Parallel Server supports this statement, which is an extension to the ANSI/ISO standard for SQL.

Syntax



Element	Description	Restrictions	Syntax
<i>dbspace</i>	Name of the dbspace in which the fragment resides	Must exist.	"Identifier " on page 5-21;
<i>owner</i>	The owner of <i>table</i>	Must own the <i>table</i>	" Owner Name " on page 5-45
<i>table</i>	Name of the table for which you want to change the residency state	Must exist	"Identifier " on page 5-21;

Usage

This statement was formerly supported by Dynamic Server, but it is ignored in current releases. Beginning with Version 9.40, Dynamic Server determines the residency status of indexes and tables automatically. In Extended Parallel Server, however, you can use this statement to specify the residency status of table fragments.

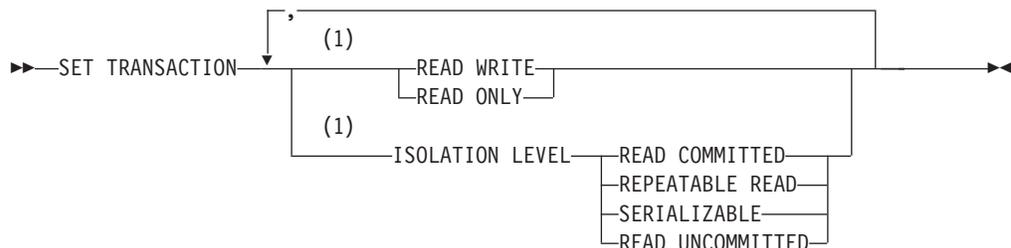
The SET TABLE statement is a special case of the SET Residency statement. The SET Residency statement can also specify how long an index fragment remains resident in shared memory.

For the complete syntax and semantics of the SET TABLE statement, see "SET Residency " on page 2-646.

SET TRANSACTION

Use the SET TRANSACTION statement to define the isolation level and to specify whether the access mode of a transaction is read-only or read-write.

Syntax



Notes:

- 1 Use path no more than once

Usage

SET TRANSACTION is valid only in databases with transaction logging. You can issue a this statement from a client computer only after a database is opened. The transaction isolation level affects concurrency among processes that attempt to access the same rows simultaneously from the database. The database server uses shared locks to support different levels of isolation among processes that are attempting to read data, as the following list shows:

- Read Uncommitted
- Read Committed
- (ANSI) Repeatable Read
- Serializable

The update or delete process always acquires an exclusive lock on the row that is being modified. The level of isolation does not interfere with such rows, but the access mode does affect whether you can update or delete rows.

If another process attempts to update or delete rows that you are reading with an isolation level of Serializable or (ANSI) Repeatable Read, that process will be denied access to those rows.

Comparing SET TRANSACTION with SET ISOLATION

The SET TRANSACTION statement complies with ANSI SQL-92. This statement is similar to the Informix SET ISOLATION statement; however, the SET ISOLATION statement is not ANSI compliant and does not provide access modes. In fact, the isolation levels that you can set with the SET TRANSACTION statement are almost parallel to the isolation levels that you can set with the SET ISOLATION statement, as the following table shows.

SET TRANSACTION Isolation Level	SET ISOLATION Isolation Level
Read Uncommitted	Dirty Read
Read Committed	Committed Read
[<i>Not supported</i>]	Cursor Stability
(ANSI) Repeatable Read	(Informix) Repeatable Read
Serializable	(Informix) Repeatable Read

Another difference between SET TRANSACTION and SET ISOLATION is the behavior of the isolation levels within transactions. You can issue SET TRANSACTION only once for a transaction. Any cursors that are opened during that transaction are guaranteed that isolation level (or access mode, if you are defining an access mode). With SET ISOLATION, after a transaction is started, you can change the isolation level more than once within the transaction.

The following examples illustrate this difference in the behavior of the SET ISOLATION and SET TRANSACTION statements:

```
EXEC SQL BEGIN WORK;  
EXEC SQL SET ISOLATION TO DIRTY READ;  
EXEC SQL SELECT ... ;  
EXEC SQL SET ISOLATION TO REPEATABLE READ;  
EXEC SQL INSERT ... ;  
EXEC SQL COMMIT WORK;    -- Executes without error
```

Compare the previous example to these SET TRANSACTION statements:

```

EXEC SQL BEGIN WORK;
EXEC SQL SET TRANSACTION ISOLATION LEVEL SERIALIZABLE;
EXEC SQL SELECT ... ;
EXEC SQL SET TRANSACTION ISOLATION LEVEL READ COMMITTED;
  -- Produces error 876: Cannot issue SET TRANSACTION
  -- in an active transaction.

```

An additional difference between SET ISOLATION and SET TRANSACTION is the duration of isolation levels. Because SET ISOLATION supports complete-connection level settings, the isolation level specified by SET ISOLATION remains in effect until another SET ISOLATION statement is issued. The isolation level set by SET TRANSACTION only remains in effect until the transaction terminates. Then the isolation level is reset to the default for the database type.

Informix Isolation Levels

The following definitions explain the critical characteristics of each isolation level, from the lowest level of isolation to the highest.

Using the Read Uncommitted Option

Use the Read Uncommitted option to copy rows from the database whether or not locks are present on them. The program that fetches a row places no locks and it respects none. Read Uncommitted is the only isolation level available to databases that do not have transactions.

This isolation level is most appropriate in queries of static tables whose data is not being modified, because it provides no isolation. With Read Uncommitted, the program might return an uncommitted row that was inserted or modified within a transaction that was subsequently rolled back.

The Uncommitted Read isolation level of SET TRANSACTION does not directly support the LAST COMMITTED feature of the Committed Read isolation level of the SET ISOLATION statement. The LAST COMMITTED feature can reduce the risk of locking conflicts when an application attempts to read a row on which another session holds an exclusive lock while modifying data. When this feature is enabled, the database server returns the most recently committed version of the data, rather than wait for the lock to be released.

This feature takes effect implicitly, however, in all user sessions that use the Read Uncommitted isolation level of the SET TRANSACTION statement, under either of the following circumstances:

- if the USELASTCOMMITTED configuration parameter is set to 'DIRTY READ' or to 'ALL'
- if the SET ENVIRONMENT statement set the USELASTCOMMITTED session environment option to 'DIRTY READ' or to 'ALL'.

See the section “The LAST COMMITTED Option to Committed Read” on page 2-634 for more information about the LAST COMMITTED feature and its restrictions.

Using the Read Committed Option

Use the Read Committed option to guarantee that every retrieved row is committed in the table at the time that the row is retrieved. This option does not place a lock on the fetched row. Read Committed is the default level of isolation in a database with logging that is not ANSI compliant.

Read Committed is appropriate when each row of data is processed as an independent unit, without reference to other rows in the same or other tables.

The Read Committed isolation level of SET TRANSACTION does not directly support the LAST COMMITTED feature of the Committed Read isolation level of the SET ISOLATION statement, which can reduce the risk of locking conflicts when an application attempts to read data in a row on which another session holds an exclusive row-level lock. When this feature is enabled, the database server returns the most recently committed version of the data, rather than wait for the lock to be released

This feature takes effect implicitly, however, in all user sessions that use the Read Committed isolation level of the SET TRANSACTION statement, under either of the following circumstances:

- if the USELASTCOMMITTED configuration parameter is set to 'COMMITTED READ' or to 'ALL'
- if the SET ENVIRONMENT statement set the USELASTCOMMITTED session environment variable to 'COMMITTED READ' or to 'ALL'.

See the section “The LAST COMMITTED Option to Committed Read” on page 2-634 for more information about the LAST COMMITTED feature and its restrictions.

Using the Repeatable Read and Serializable Options

The Informix implementation of Repeatable Read and of Serializable are equivalent. The Serializable (or Repeatable Read) option places a shared lock on every row that is selected during the transaction.

Another process can also place a shared lock on a selected row, but no other process can modify any selected row during your transaction or insert a row that meets the search criteria of your query during your transaction.

A *phantom row* is a row that was not visible when you first read the query set, but that materializes in a subsequent read of the query set in the same transaction. Only this isolation level prevents access to a phantom row.

If you repeat the query during the transaction, you reread the same data. The shared locks are released only when the transaction is committed or rolled back. Serializable is the default isolation level in an ANSI-compliant database. Serializable isolation places the largest number of locks and holds them the longest. Therefore, it is the level that reduces concurrency the most.

Default Isolation Levels

The default isolation level is established when you create the database.

Informix Name	ANSI Name	When This Is the Default Level of Isolation
Dirty Read	Read Uncommitted	Database without transaction logging
Committed Read	Read Committed	Databases with logging that are not ANSI-compliant
Repeatable Read	Serializable	ANSI-compliant databases

For a Dynamic Server database that is not ANSI-compliant, unless you explicitly set the USELASTCOMMITTED configuration parameter, the LAST COMMITTED feature is not in effect for the default isolation levels. The SET ENVIRONMENT statement or the SET ISOLATION statement can override this default and enable LAST COMMITTED for the current session.

The default isolation level remains in effect until you issue a `SET TRANSACTION` statement within a transaction. After a `COMMIT WORK` statement completes the transaction or a `ROLLBACK WORK` statement cancels the transaction, the isolation level is reset to the default.

When you use High Availability Data Replication, the database server effectively uses Dirty Read isolation on the HDR Secondary Server, regardless of the specified `SET ISOLATION` or `SET TRANSACTION` isolation level.

Access Modes

Access modes affect read and write concurrency for rows within transactions. Use access modes to control data modification. `SET TRANSACTION` can specify that a transaction is read-only or read-write. By default, transactions are read-write. When you specify a read-only transaction, certain limitations apply. Read-only transactions cannot perform the following actions:

- Insert, delete, or update rows of a table.
- Create, alter, or drop any database object such as schemas, tables, temporary tables, indexes, or SPL routines.
- Grant or revoke access privileges.
- Update statistics.
- Rename columns or tables.

You can execute SPL routines in a read-only transaction as long as the SPL routine does not try to perform any restricted statement.

Effects of Isolation Levels

You cannot set the transaction isolation level in a database that does not have logging. Every retrieval in an unlogged database occurs as a Read Uncommitted.

The data that is obtained during retrieval of `BYTE` or `TEXT` data can vary, depending on the transaction isolation levels. Under Read Uncommitted or Read Committed isolation levels, a process is permitted to read a `BYTE` or `TEXT` column that is either deleted (if the delete is not yet committed) or in the process of being deleted. Under these isolation levels, an application can read a deleted `BYTE` or `TEXT` column when certain conditions exist. For information about these conditions, see the *IBM Informix Administrator's Guide*.

In Informix ESQL/C, if you use a scroll cursor in a transaction, you can force consistency between your temporary table and the database table either by setting the isolation level to Serializable or by locking the entire table. A scroll cursor with hold, however, cannot guarantee the same consistency between the two tables. Table-level locks set by Serializable are released when the transaction is completed, but the scroll cursor with hold remains open beyond the end of the transaction. You can modify released rows as soon as the transaction ends, so the retrieved data in the temporary table might be inconsistent with the actual data.

Warning: Do not use nonlogging tables within a transaction. If you need to use a nonlogging table within a transaction, either set the isolation level to Repeatable Read or lock the table in exclusive mode to prevent concurrency problems.

Related Information

Related statements: “`CREATE DATABASE`” on page 2-97, “`SET ISOLATION`” on page 2-631, and “`SET LOCK MODE`” on page 2-637

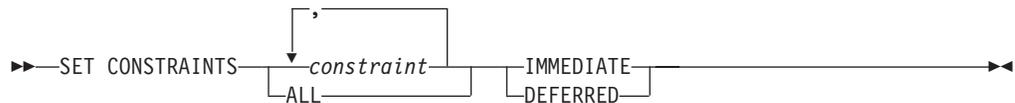
For a discussion of isolation levels and concurrency issues, see the *IBM Informix Guide to SQL: Tutorial*.

SET Transaction Mode

Use the SET Transaction Mode statement to specify whether constraints are checked at the statement level or at the transaction level during the current transaction.

In Dynamic Server, to enable or disable constraints, or to change their filtering mode, see “SET Database Object Mode” on page 2-591.

Syntax



Element	Description	Restrictions	Syntax
<i>constraint</i>	Constraint whose transaction mode is to be changed	All constraints must exist in the same database, which must support logging	“Identifier ” on page 5-21

Usage

This statement is valid only in a database with transaction logging, and its effect is limited to the transaction in which it is executed.

Use the IMMEDIATE keyword to set the transaction mode of constraints to statement-level checking. IMMEDIATE is the default transaction mode of constraints when they are created.

Use the DEFERRED keyword to set the transaction mode to transaction-level checking. You cannot change the transaction mode of a constraint to DEFERRED unless the constraint is currently enabled.

Statement-Level Checking

When you set the transaction mode to IMMEDIATE, statement-level checking is turned on, and all specified constraints are checked at the end of each INSERT, UPDATE, or DELETE statement. If a constraint violation occurs, the statement is not executed.

Transaction-Level Checking

When you set the transaction mode of constraints to DEFERRED, statement-level checking is turned off, and all (or the specified) constraints are not checked until the transaction is committed. If a constraint violation occurs while the transaction is being committed, the transaction is rolled back.

Tip: If you defer checking a primary-key constraint, checking the not-NULL constraint for that column or set of columns is also deferred.

Duration of Transaction Modes

The duration of the transaction mode that the SET Transaction Mode statement specifies is the transaction in which the SET Transaction Mode statement is executed. You cannot execute this statement outside a transaction. Once a COMMIT WORK or ROLLBACK WORK statement is successfully completed, the transaction mode of all constraints reverts to IMMEDIATE.

To switch from transaction-level checking to statement-level checking, you can use the SET Transaction Mode statement to set the transaction mode to IMMEDIATE, or you can use a COMMIT WORK or ROLLBACK WORK statement to terminate your transaction.

Specifying All Constraints or a List of Constraints

You can specify all constraints in the database in the SET Transaction Mode statement, or you can specify a single constraint, or list of constraints.

If you specify the ALL keyword, the SET Transaction Mode statement sets the transaction mode for all constraints in the database. If any statement in the transaction requires that any constraint on any table in the database be checked, the database server performs the checks at the statement level or the transaction level, depending on the setting that you specify in the SET Transaction Mode statement.

If you specify a single constraint name or a list of constraints, the SET Transaction Mode statement sets the transaction mode for the specified constraints only. If any statement in the transaction requires checking of a constraint that you did not specify in the SET Transaction Mode statement, that constraint is checked at the statement level regardless of the setting that you specified in the SET Transaction Mode statement for other constraints.

When you specify a list of constraints, the constraints do not need to be defined on the same table, but they must exist in the same database.

Specifying Remote Constraints

You can set the transaction mode of local constraints or remote constraints. That is, the constraints that are specified in the SET Transaction Mode statement can be constraints that are defined on local tables or constraints that are defined on remote tables.

Examples of Setting the Transaction Mode for Constraints

The following example shows how to defer checking constraints within a transaction until the transaction is complete. The SET Transaction Mode statement in the example specifies that any constraints on any tables in the database are not checked until the COMMIT WORK statement is encountered.

```
BEGIN WORK;  
SET CONSTRAINTS ALL DEFERRED;  
...  
COMMIT WORK;
```

The following example specifies that a list of constraints is not checked until the transaction is complete:

```

BEGIN WORK;
SET CONSTRAINTS update_const, insert_const DEFERRED;
...
COMMIT WORK;

```

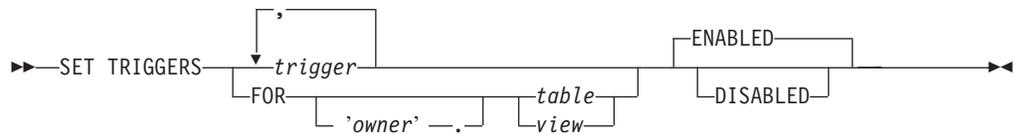
Related Information

Related statements: “ALTER TABLE” on page 2-44 and “CREATE TABLE ” on page 2-191

SET TRIGGERS

Use the SET TRIGGERS statement to enable or disable all or some of the triggers on a table, or all or some of the INSTEAD OF triggers on a view.

Syntax



Element	Description	Restrictions	Syntax
<i>owner</i>	The owner of <i>table</i> or <i>view</i>	Must own the <i>table</i> or <i>view</i>	“Owner Name ” on page 5-45
<i>table</i>	Table whose triggers are all to be enabled or disabled	Must exist	“Identifier ” on page 5-21
<i>trigger</i>	Trigger to be enabled or disabled	Must exist	“Identifier ” on page 5-21
<i>view</i>	View whose INSTEAD OF triggers are all to be enabled or disabled	Must exist	“Identifier ” on page 5-21

Usage

The SET TRIGGERS statement is a special case of the SET Database Object Mode statement. The SET Database Object Mode statement can also enable or disable an index or a constraint, or change the filtering mode of a unique index or of a constraint.

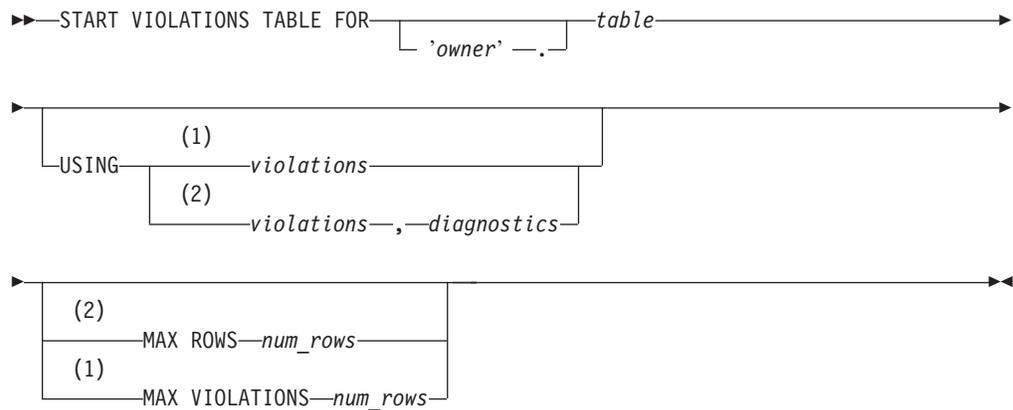
For the complete syntax and semantics of the SET TRIGGERS statement, see “SET Database Object Mode” on page 2-591.

START VIOLATIONS TABLE

Use the START VIOLATIONS TABLE statement to create a violations table and (for Dynamic Server only) a diagnostics table for a specified target table.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

- 1 Extended Parallel Server only
- 2 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>diagnostics</i>	Declares the name of a diagnostics table to be associated with the target <i>table</i> . Default name is <i>table_dia</i> .	Must be unique among names of tables, views, sequences, and synonyms	"Identifier " on page 5-21
<i>num_rows</i>	Maximum number of rows database server (IDS) or any single coserver (XPS) can insert into <i>violations</i> when a single statement is executed on <i>table</i>	Must be an integer in range from 1 to the maximum value of the INTEGER data type	"Literal Number " on page 4-185
<i>owner</i>	The owner of <i>table</i>	Must own <i>table</i>	" Owner Name " on page 5-45
<i>table</i>	Target table for which <i>violations</i> and (for Dynamic Server only) <i>diagnostics</i> are to be created	If USING clause is omitted, no more than 124 bytes	"Identifier " on page 5-21
<i>violations</i>	Violations table to be associated with <i>table</i> . Default name is <i>table_vio</i> .	Same restrictions as <i>diagnostics</i>	"Identifier " on page 5-21

Usage

The database server associates the *violations* table (and for Dynamic Server only the *diagnostics* table) with the *target* table that you specify after the FOR keyword by recording the relationship among the three tables in the **sysviolations** system catalog table. In Extended Parallel Server, the `START VIOLATIONS TABLE` statement creates a violations table, but no diagnostics table is created.

A target table must satisfy these requirements:

- It cannot be external to the database.
- It cannot already be associated with a violations or diagnostics table.
- It cannot be a system catalog table.

The `START VIOLATIONS TABLE` statement creates the special violations table that holds nonconforming rows that fail to satisfy constraints and unique indexes during insert, update, and delete operations on target tables. This statement also

creates the special diagnostics table that contains information about the integrity violations that each row causes in the violations table.

Relationship to SET Database Object Mode Statement

The `START VIOLATIONS TABLE` statement is closely related to the `SET Database Object Mode` statement. If you use `SET Database Object Mode` to set the constraints or unique indexes defined on a table to the `FILTERING` mode, without also using `START VIOLATIONS TABLE`, any rows that violate a constraint or unique-index requirement in data manipulation operations are not filtered out to a violations table. Instead you receive an error message that indicates that you must start a violations table for the target table.

Similarly, if you use the `SET Database Object Mode` statement to set a disabled constraint or disabled unique index to the `ENABLED` or `FILTERING` mode, but you do not use `START VIOLATIONS TABLE` for the table on which the database objects are defined, any rows that do not satisfy the constraint or unique-index requirement are not filtered out to a violations table.

In these cases, to identify the rows that do not satisfy the constraint or unique-index requirement, issue the `START VIOLATIONS TABLE` statement to start the violations and diagnostics tables. Do this before you use the `SET Database Object Mode` statement to set the database objects to the `ENABLED` or `FILTERING` database object mode.

Extended Parallel Server does not support the `SET Database Object Mode` feature, and the concept of database object modes does not exist. Once you use the `START VIOLATIONS TABLE` statement to create a violations table and associate it with a target table, the existence of this violations table causes all violations of constraints and unique-index requirements by insert, delete, and update operations to be recorded in the violations table.

In other words, once you issue a `START VIOLATIONS TABLE` statement, all constraints and unique indexes in a database on Extended Parallel Server behave like filtering-mode constraints and filtering-mode unique indexes in a database on Dynamic Server. For an explanation of filtering-mode constraints and filtering-mode unique indexes, see “Filtering Mode” on page 2-595.

Effect on Concurrent Transactions

If the database has transaction logging, you must issue `START VIOLATIONS TABLE` in isolation. That is, no other transaction can be in progress on a target table when you issue `START VIOLATIONS TABLE` on that table within a transaction. Any transactions that start on the target table after the first transaction has issued the `START VIOLATIONS TABLE` statement will behave the same way as the first transaction with respect to the violations and diagnostics tables. That is, any constraint and unique-index violations by these subsequent transactions will be recorded in the violations and diagnostics tables.

For example, if transaction A operates on table **tab1** and issues a `START VIOLATIONS TABLE` statement on table **tab1**, the database server starts a violations table named **tab1_vio** and filters any constraint or unique-index violations on table **tab1** by transaction A to table **tab1_vio**. If transactions B and C start on table **tab1** after transaction A has issued the `START VIOLATIONS TABLE` statement, the database server also filters any constraint and unique-index violations by transactions B and C to table **tab1_vio**.

The result is that all three transactions do not receive error messages about constraint and unique-index violations, even though transactions B and C do not expect this behavior. For example, if transaction B issues an INSERT or UPDATE statement that violates a check constraint on table **tab1**, the database server does not issue a constraint violation error to transaction B. Instead, the database server filters the nonconforming row (also called a "bad row") to the violations table without notifying transaction B that a data-integrity violation occurred.

You can prevent this situation from arising in Dynamic Server by specifying WITH ERRORS when you specify the FILTERING mode in a SET Database Object Mode, CREATE TABLE, ALTER TABLE, or CREATE INDEX statement. When multiple transactions operate on a table and the WITH ERRORS option is in effect, any transaction that violates a constraint or unique-index requirement on a target table receives a data-integrity error message.

In Extended Parallel Server, once a transaction issues a START VIOLATIONS TABLE statement, you have no way to make the database server issue data-integrity violation messages to that transaction or to any other transactions that subsequently start on the same target table.

Stopping the Violations and Diagnostics Tables

After you use START VIOLATIONS TABLE to create an association between a target table and the violations and diagnostics tables, the only way to drop the association between the target table and the violations and diagnostics tables is to issue a STOP VIOLATIONS TABLE statement for the target table. For more information, see "STOP VIOLATIONS TABLE " on page 2-676.

USING Clause

You can use the USING clause to declare explicit names for the violations table and (for Dynamic Server) the diagnostics table.

If you omit the USING clause, the database server assigns names to the violation table and (for Dynamic Server) the diagnostics table. The system-assigned name of the violations table consists of the name of the target table followed by the string `_vio`. The name that Dynamic Server assigns to the diagnostics table consists of the name of the target table followed by the string `_dia`.

If you omit the USING clause, the maximum length of the name of the target table is 124 bytes.

Using the MAX ROWS Clause (IDS)

The MAX ROWS clause specifies the maximum number of rows that the database server can insert into the diagnostics table when a single statement is executed on the target table. If you omit the MAX ROWS clause, no upper limit is imposed on the number of rows that can be inserted into the diagnostics table when a single statement is executed on the target table.

Using the MAX VIOLATIONS Clause (XPS)

Use the MAX VIOLATIONS clause to specify the maximum number of rows that any single coserver can insert into the violations table when a single statement is executed on the target table. Each coserver where the violations table resides has this limit. The first coserver to reach this limit raises an error and causes the

statement to fail. If you omit the MAX VIOLATIONS clause, no upper limit exists on the number of rows that can be inserted into the violations table when a single statement is executed on the target table.

Specifying the Maximum Number of Rows in the Diagnostics Table (IDS)

The following statement starts violations and diagnostics tables for the target table named **orders**. The MAX ROWS clause specifies the maximum number of rows that can be inserted into the diagnostics table when a single statement, such as an INSERT statement, is executed on the target table.

```
START VIOLATIONS TABLE FOR orders MAX ROWS 50000;
```

Specifying the Maximum Number of Rows in the Violations Table (XPS)

The following statement starts a violations table for the target table named **orders**. The MAX VIOLATIONS clause specifies the maximum number of rows that any single coserver can insert into the violations table when a single statement, such as an INSERT statement, is executed on the target table.

```
START VIOLATIONS TABLE FOR orders MAX VIOLATIONS 50000;
```

Privileges Required for Starting Violations Tables

To start a violations and (for Dynamic Server) a diagnostics table for a target table, you must meet one of the following requirements:

- You must have the DBA privilege on the database.
- You must be the owner of the target table and have the Resource privilege on the database.
- You must have the Alter privilege on the target table and the Resource privilege on the database.

Structure of the Violations Table

When you issue START VIOLATIONS TABLE for a target table, the violations table that the statement creates has a predefined structure. This structure consists of the columns of the target table and three additional columns.

The following table shows the schema of the violations table.

Column Name	Data Type	Column Description
Same columns (in the same order) that appear in the target table	Same types as corresponding columns in the target table.	The violations table has the same schema as the target table, so that rows violating constraints or a unique-index during insert, update, and delete operations can be filtered to the violations table.
informix_tupleid	SERIAL	Unique serial code for the nonconforming row
informix_optype	CHAR(1)	The type of operation that caused this bad row. This column can have the following values:
		I = Insert
		D = Delete
		O = Update (with original values in this row)
		N = Update (with new values in this row)
		S = SET Database Object Mode (IDS only)

Column Name	Data Type	Column Description
<code>informix_recowner</code>	CHAR(32)	User who issued the statement that created this nonconforming row

If the target table of the START VIOLATIONS TABLE statement is protected by a security policy, Dynamic Server protects the violations table with same security policy. In this case, the schema of the violations table includes an IDSSECURITYLABEL column whose name and position among other columns corresponds to the IDSSECURITYLABEL column of the target table. When the violations table is created, any SECURED WITH *label* specifications that protect columns in the target table also protect the corresponding violations table columns.

Serial columns in the target table are converted to integer data types in the violations table.

Users can examine these nonconforming rows in the violations table, analyze the related rows that contain diagnostic information in the diagnostics table, and take corrective actions.

Examples of START VIOLATIONS TABLE Statements

The following examples show different ways to execute the START VIOLATIONS TABLE statement.

Violations and Diagnostics Tables with Default Names

The following statement starts violations and diagnostics tables for the target table named `cust_subset`. The violations table is named `cust_subset_vio` by default, and the diagnostics table is named `cust_subset_dia` by default.

```
START VIOLATIONS TABLE FOR cust_subset;
```

Violations and Diagnostics Tables with Explicit Names

The following statement starts a violations and diagnostics table for the target table named `items`. The USING clause assigns explicit names to the violations and diagnostics tables. The violations table is to be named `exceptions`, and the diagnostics table is to be named `reasons`.

```
START VIOLATIONS TABLE FOR items USING exceptions, reasons;
```

Relationships Among the Target, Violations, and Diagnostics Tables

Users can take advantage of the relationships among the target, violations, and diagnostics tables to obtain diagnostic information about rows that cause data-integrity violations during INSERT, DELETE, and UPDATE statements. Each row of the violations table has at least one corresponding row in the diagnostics table.

- One row in the violations table is a copy of any row in the target table for which a data-integrity violation was detected. A row in the diagnostics table contains information about the nature of the data-integrity violation caused by the nonconforming row in the violations table.
- One row in the violations table has a unique serial identifier in the `informix_tupleid` column. A row in the diagnostics table has the same serial identifier in its `informix_tupleid` column.

A given row in the violations table can have more than one corresponding row in the diagnostics table. The multiple rows in the diagnostics table all have the same

serial identifier in their **informix_tupleid** column so that they are all linked to the same row in the violations table. Multiple rows can exist in the diagnostics table for the same row in the violations table because a nonconforming row in the violations table can cause more than one data-integrity violation.

For example, the same nonconforming row can violate a unique index for one column, a not-NULL constraint for another column, and a check constraint for a third column. In this case, the diagnostics table contains three rows for the single nonconforming row in the violations table. Each of these diagnostic rows identifies a different data-integrity violation that the nonconforming row in the violations table caused.

By joining the violations and diagnostics tables, the DBA or target-table owner can obtain diagnostic information about any or all nonconforming rows in the violations table. SELECT statements can perform these joins interactively, or you can write a program to perform them within transactions.

Initial Privileges on the Violations Table

When you issue the START VIOLATIONS TABLE statement to create the violations table, the database server uses the set of privileges granted on the target table as a basis for granting privileges on the violations table. The database server follows different rules, however, when it grants each type of privilege.

The following table summarizes the circumstances under which the database server grants each type of privilege on the violations table.

Privilege

Condition for Granting the Privilege

Alter Alter privilege is not granted on the violations table. (Users cannot alter violations tables.)

Index User has Index privilege on the violations table if the user has the Index privilege on the target table.

The user cannot create a globally detached index on the violations table even if the user has the Index privilege on the violations table (XPS only).

Insert User has the Insert privilege on the violations table if the user has the Insert, Delete, or Update privilege on any column of the target table.

Delete User has the Delete privilege on the violations table if the user has the Insert, Delete, or Update privilege on any column of the target table.

Select User has the Select privilege on the **informix_tupleid**, **informix_optype**, and **informix_reowner** columns of the violations table if the user has the Select privilege on any column of the target table.

User has the Select privilege on any other column of the violations table if the user has the Select privilege on the same column in the target table.

Update

User has the Update privilege on the **informix_tupleid**, **informix_optype**, and **informix_reowner** columns of the violations table if the user has the Update privilege on any column of the target table.

(Even with the Update privilege on the **informix_tupleid** column, however, the user cannot update this SERIAL column.)

User has the Update privilege on any other violations table column if the user has the Update privilege on the same column in the target table.

References

The References privilege is not granted on the violations table. (Users cannot add referential constraints to violations tables.)

The following rules apply to ownership of the violations table and privileges on the violations table:

- When the violations table is created, the owner of the target table becomes the owner of the violations table.
- The owner of the violations table automatically receives all table-level privileges on the violations table, including the Alter and References privileges. The database server, however, prevents the owner of the violations table from altering the violations table or adding a referential constraint to the violations table.
- You can use the GRANT and REVOKE statements to modify the initial set of privileges on the violations table.
- When you issue an INSERT, DELETE, or UPDATE statement on a target table that has a filtering-mode unique index or constraint defined on it, you must have the Insert privilege on the violations and diagnostics tables.

If you do not have the Insert privilege on the violations and diagnostics tables, the database server executes the INSERT, DELETE, or UPDATE statement on the target table provided that you have the necessary privileges on the target table. The database server does not return an error concerning the lack of Insert privilege on the violations and diagnostics tables unless an integrity violation is detected during execution of the INSERT, DELETE, or UPDATE statement.

Similarly, when you issue a SET Database Object Mode statement to set a disabled constraint or disabled unique index to the enabled or filtering mode, and a violations table and diagnostics table exist for the target table, you must have the Insert privilege on the violations and diagnostics tables.

If you do not have the Insert privilege on the violations and diagnostics tables, the database server executes the SET Database Object Mode statement if you have the necessary privileges on the target table. The database server does not return an error concerning the lack of Insert privilege on the violations and diagnostics tables, unless an integrity violation is detected during the execution of the SET Database Object Mode statement.

- The grantor of the initial set of privileges on the violations table is the same as the grantor of the privileges on the target table.
For example, if user **henry** was granted the Insert privilege on the target table by both user **jill** and user **albert**, then the Insert privilege on the violations table is granted to **henry** both by **jill** and by **albert**.
- After the violations table is started, revoking a privilege on the target table from a user does not automatically revoke the same privilege on the violations table from that user. Instead, you must explicitly revoke the privilege on the violations table from the user.
- If you have fragment-level privileges on the target table, you have the corresponding fragment-level privileges on the violations table.

Example of Privileges on the Violations Table

The following example illustrates how the initial set of privileges on a violations table is derived from the current set of privileges on the target table. Assume that a table named **cust_subset** consists of the following columns: **ssn** (customer Social Security number), **fname** (customer first name), **lname** (customer last name), and **city** (city in which the customer lives).

The following set of privileges exists on the **cust_subset** table:

- User **barbara** has the Insert and Index privileges on the table. She also has the Select privilege on the **ssn** and **lname** columns.
- User **carrie** has the Update privilege on the **city** column. She also has the Select privilege on the **ssn** column.
- User **danny** has the Alter privilege on the table.

Now user **alvin** starts a violations table named **cust_subset_viol**s and a diagnostics table named **cust_subset_diag**s for the **cust_subset** table:

```
START VIOLATIONS TABLE FOR cust_subset
  USING cust_subset_viol, cust_subset_diag;
```

The database server grants the following set of initial privileges on the **cust_subset_viol**s violations table:

- User **alvin** is the owner of the violations table, so he has all table-level privileges on the table.
- User **barbara** has the Insert, Delete, and Index privileges on the table. User **barbara** has the Select privilege on five columns of the violations table: the **ssn**, the **lname**, the **informix_tupleid**, the **informix_optype**, and the **informix_reowner** columns.
- User **carrie** has Insert and Delete privileges on the violations table. User **carrie** has the Update privilege on four columns of the violations table: the **city**, the **informix_tupleid**, the **informix_optype**, and the **informix_reowner** columns. She cannot, however, update the **informix_tupleid** column (because this is a SERIAL column). User **carrie** has the Select privilege on four columns of the violations table: the **ssn** column, the **informix_tupleid** column, the **informix_optype** column, and the **informix_reowner** column.
- User **danny** has no privileges on the violations table.

Using the Violations Table

The following rules concern the structure and use of the violations table:

- Every pair of update rows in the violations table has the same value in the **informix_tupleid** column to indicate that both rows refer to the same row in the target table.
- If the target table has columns named **informix_tupleid**, **informix_optype**, or **informix_reowner**, the database server attempts to generate alternative names for these columns in the violations table by appending a digit to the end of the column name (for example, **informix_tupleid1**). If this fails, an error is returned, and no violations table is started for the target table.
- When a table functions as a violations table, it cannot have triggers or constraints defined on it.
- When a table functions as a violations table, users can create indexes on it, even though the existence of an index affects performance. Unique indexes on the violations table cannot be set to FILTERING database object mode.
- If a target table has a violations and diagnostics table associated with it, dropping the target table in cascade mode (the default mode) causes the violations and diagnostics tables to be dropped also. If the target table is dropped in the restricted mode, the DROP TABLE operation fails (because the violations and diagnostics tables exist).

- After a violations table is started for a target table, ALTER TABLE cannot add, modify, or drop columns of the violations, diagnostics, or target tables. Before you can alter any of these tables, you must issue a STOP VIOLATIONS TABLE statement for the target table.
- The database server does not clear out the contents of the violations table before or after it uses the violations table during an INSERT, UPDATE, DELETE, or SET Database Object Mode operation.
- If a target table has a filtering-mode constraint or unique index defined on it and a violations table associated with it, users cannot insert into the target table by selecting from the violations table. Before you insert rows into the target table by selecting from the violations table, you must take one of the following steps:
 - You can set the constraint or unique index to DISABLED mode.
 - You can issue STOP VIOLATIONS TABLE for the target table.

If it is inconvenient to take either of these steps, but you intend to copy records from the violations table into the target table, a third option is to select from the violations table into a temporary table and then insert the contents of the temporary table into the target table.
- If the target table that is specified in the START VIOLATIONS TABLE statement is fragmented, the violations table has the same fragmentation strategy as the target table. Each fragment of the violations table is stored in the same dbspace partition as the corresponding fragment of the target table.
- Once a violations table is started for a target table, you cannot use the ALTER FRAGMENT statement to alter the fragmentation strategy of the target table or the violations table.
- If the target table specified in the START VIOLATIONS TABLE statement is not fragmented, the database server places the violations table in the same dbspace as the target table.
- If the target table has BYTE or TEXT columns, BYTE or TEXT data values in the violations table are created in the same blob space that stores the BYTE or TEXT data in the target table.

Example of a Violations Table

To start a violations and diagnostics table for the target table named **customer** in the demonstration database, enter the following statement:

```
START VIOLATIONS TABLE FOR customer;
```

Because you include no USING clause, the violations table is named **customer_vio** by default. The **customer_vio** table includes these columns:

customer_num	company	city state		
fname lname	address1		zipcode	informix_tupleid
	address2		phone	informix_optype
				informix_recowner

The **customer_vio** table has the same table definition as the **customer** table except that the **customer_vio** table has three additional columns that contain information about the operation that caused the nonconforming row.

Structure of the Diagnostics Table (IDS)

When you issue a START VIOLATIONS TABLE statement for a target table, the diagnostics table that the statement creates has a predefined structure. This structure is independent of the structure of the target table.

The following table shows the schema of the diagnostics table.

Column Name	Data Type	Description
informix_tupleid	INTEGER	Implicitly refers to informix_tupleid column values in the violations table This relationship, however, is not declared as a foreign-key to primary-key relationship.
objtype	CHAR(1)	Identifies the type of violation This column can have the following values: <hr/> C = Constraint violation <hr/> I = Unique-index violation
objowner	CHAR(32)	Identifies the owner of the constraint or index for which an integrity violation was detected
objname	VARCHAR(128, 0)	Contains the name of the constraint or index for which an integrity violation was detected

Initial Privileges on the Diagnostics Table

When the `START VIOLATIONS TABLE` statement creates the diagnostics table, the set of privileges granted on the target table are a basis for granting privileges on the diagnostics table. The database server follows different rules, however, when it grants each type of privilege.

The following table explains the circumstances under which the database server grants each privilege on the diagnostics table.

Privilege

Condition for Granting the Privilege

Insert User has the Insert privilege on the diagnostics table if the user has the Insert, Delete, or Update privilege on any column of the target table.

Delete User has the Delete privilege on the diagnostics table if the user has the Insert, Delete, or Update privilege on any column of the target table.

Select User has the Select privilege on the diagnostics table if the user has the Select privilege on any column in the target table.

Update

User has the Update privilege on the diagnostics table if the user has the Update privilege on any column in the target table.

Index User has the Index privilege on the diagnostics table if the user has the Index privilege on the target table.

Alter Alter privilege is not granted on the diagnostics table.

(Users cannot alter diagnostics tables.)

References

References privilege is not granted on the diagnostics table.

(Users cannot add referential constraints to diagnostics tables.)

The following rules concern access privileges on the diagnostics table:

- When the diagnostics table is created, the owner of the target table becomes the owner of the diagnostics table.
- The owner of the diagnostics table automatically receives all table-level privileges on the diagnostics table, including the Alter and References privileges.

The database server, however, prevents the owner of the diagnostics table from altering the diagnostics table or adding a referential constraint to the diagnostics table.

- You can use the GRANT and REVOKE statements to modify the initial set of privileges on the diagnostics table.
- For INSERT, DELETE, or UPDATE operations on a target table that has a filtering-mode unique index or constraint defined on it, you must have the Insert privilege on the violations and diagnostics tables.

If you do not have the Insert privilege on the violations and diagnostics tables, the database server executes the INSERT, DELETE, or UPDATE statement on the target table, provided you have the necessary privileges on the target table. The database server does not return an error concerning the lack of Insert privilege on the violations and diagnostics tables unless an integrity violation is detected during execution of the INSERT, DELETE, or UPDATE statement.

Similarly, when you issue a SET Database Object Mode statement to set a disabled constraint or disabled unique index to the enabled or filtering mode, and a violations table and diagnostics table exist for the target table, you must have the Insert privilege on the violations and diagnostics tables.

If you do not have the Insert privilege on the violations and diagnostics tables, the database server executes the SET Database Object Mode statement, provided you have the necessary privileges on the target table. The database server does not return an error concerning the lack of Insert privilege on the violations and diagnostics tables unless an integrity violation is detected during the execution of the SET Database Object Mode statement.

- The grantor of the initial set of privileges on the diagnostics table is the same as the grantor of the privileges on the target table. For example, if the user **jenny** was granted the Insert privilege on the target table by both the user **wayne** and the user **laurie**, both user **wayne** and user **laurie** grant the Insert privilege on the diagnostics table to user **jenny**.
- Once a diagnostics table is started for a target table, revoking a privilege on the target table from a user does not automatically revoke the same privilege on the diagnostics table from that user. Instead you must explicitly revoke the privilege on the diagnostics table from the user.
- If you have fragment-level privileges on the target table, you have the corresponding table-level privileges on the diagnostics table.

The next example illustrates how the initial set of privileges on a diagnostics table is derived from the current privileges on the target table. Assume that you have a table called **cust_subset** that holds **customer data**. This table consists of the following columns: **ssn** (social security number), **fname** (first name), **lname** (last name), and **city** (city in which the customer lives). The following set of privileges exists on the **cust_subset** table:

- User **alvin** is the owner of the table.
- User **barbara** has the Insert and Index privileges on the table. She also has the Select privilege on the **ssn** and **lname** columns.
- User **danny** has the Alter privilege on the table.
- User **carrie** has the Update privilege on the **city** column. She also has the Select privilege on the **ssn** column.

Now user **alvin** starts a violations table named **cust_subset_viols** and a diagnostics table named **cust_subset_diags** for the **cust_subset** table:

```
START VIOLATIONS TABLE FOR cust_subset
  USING cust_subset_viols, cust_subset_diags;
```

The database server grants the following set of initial privileges on the **cust_subset_diags** diagnostics table:

- User **alvin** is the owner of the diagnostics table, so he has all table-level privileges on the table.
- User **barbara** has the Insert, Delete, Select, and Index privileges on the diagnostics table.
- User **carrie** has the Insert, Delete, Select, and Update privileges on the diagnostics table.
- User **danny** has no privileges on the diagnostics table.

Using the Diagnostics Table

For information on the relationship between the diagnostics table and the violations table, see “Relationships Among the Target, Violations, and Diagnostics Tables” on page 2-668.

The following issues concern the structure and use of the diagnostics table:

- The MAX ROWS clause of the START VIOLATIONS TABLE statement sets a limit on the number of rows that can be inserted into the diagnostics table when you execute a single statement, such as an INSERT or SET Database Object Mode statement, on the target table.
- The MAX ROWS clause limits the number of rows only for operations in which the table functions as a diagnostics table.
- When a table functions as a diagnostics table, it cannot have triggers or constraints defined on it.
- When a table functions as a diagnostics table, users can create indexes on the table, but the existence of an index affects performance. You cannot set unique indexes on a diagnostics table to FILTERING database object mode.
- If a target table has a violations and diagnostics table associated with it, dropping the target table in cascade mode (the default mode) causes the violations and diagnostics tables to be dropped also.
- If the target table is dropped in restricted mode, the DROP TABLE operation fails (because the violations and diagnostics tables exist).
- Once a violations table is started for a target table, you cannot use the ALTER TABLE statement to add, modify, or drop columns in the target table, violations table, or diagnostics table. Before you can alter any of these tables, you must issue a STOP VIOLATIONS TABLE statement for the target table.
- The database server does not clear out the contents of the diagnostics table before or after it uses the diagnostics table during an Insert, Update, Delete, or Set operation.
- If the target table that is specified in the START VIOLATIONS TABLE statement is fragmented, the diagnostics table is fragmented with a round-robin strategy over the same dbspaces in which the target table is fragmented.

To start a violations and diagnostics table for the target table named **stock** in the demonstration database, enter the following statement:

```
START VIOLATIONS TABLE FOR stock;
```

Because your START VIOLATIONS TABLE statement does not include a USING clause, the diagnostics table is named **stock_dia** by default. The **stock_dia** table includes the following columns:

`informix_tupleid objtype`

`objowner objname`

This list of columns shows an important difference between the diagnostics table and violations table for a target table. Whereas the violations table has a matching column for every column in the target table, the columns of the diagnostics table do not match any columns in the target table. The diagnostics table created by any START VIOLATIONS TABLE statement always has the same columns with the same column names and data types.

Related Information

Related statements: “CREATE INDEX ” on page 2-124, “CREATE TABLE ” on page 2-191, “SET Database Object Mode” on page 2-591, and “STOP VIOLATIONS TABLE ”

For a discussion of object modes and violation detection, see the *IBM Informix Guide to SQL: Tutorial*.

STOP VIOLATIONS TABLE

Use the STOP VIOLATIONS TABLE statement to drop the association between a target table and its violations table (and for Dynamic Server only, its diagnostics table). This statement is an extension to the ANSI/ISO standard for SQL.

Syntax

```
▶▶—STOP VIOLATIONS TABLE FOR [ 'owner' . ] table—▶▶
```

Element	Description	Restrictions	Syntax
<i>owner</i>	The owner of <i>table</i>	Must own <i>table</i>	“ Owner Name ” on page 5-45
<i>table</i>	Name of a target table whose association with the violations and diagnostics table is to be dropped. No default value exists.	Must be a local table that has an associated violations table and (for IDS) a diagnostics table	“Identifier ” on page 5-21

Usage

The STOP VIOLATIONS TABLE statement drops the association between the target table and the violations table (and for Dynamic Server, the diagnostics table). After you issue this statement, the former violations and diagnostics tables continue to exist, but they no longer function as violations and diagnostics tables for the target table. They now have the status of regular database tables instead of violations and diagnostics tables for the target table. You must issue the DROP TABLE statement to drop these two tables explicitly.

When DML operations (INSERT, DELETE, or UPDATE) cause data-integrity violations for rows of the target table, the nonconforming rows are no longer filtered to the former violations table, and diagnostic information about the data-integrity violations is not placed in the former diagnostics table.

In Extended Parallel Server, the diagnostics table does not exist. The STOP VIOLATIONS TABLE statement drops the association between the target table and the violations table.

Example of Stopping the Violations and Diagnostics Tables

Assume that a target table named `cust_subset` has an associated violations table named `cust_subset_vio` and (for Dynamic Server) an associated diagnostics table named `cust_subset_dia`. To drop the association between the target table and the violations and diagnostics tables, enter the following statement:

```
STOP VIOLATIONS TABLE FOR cust_subset;
```

This deletes the row in the `sysviolations` system catalog table that had registered the former association. Subsequent DML operations on the target `cust_subset` table will no longer cause the database server to insert information about nonconforming rows into its former violations table or diagnostics table.

Example of Dropping the Violations and Diagnostics Tables

After you execute the STOP VIOLATIONS TABLE statement in the preceding example, the `cust_subset_vio` and (for Dynamic Server) the `cust_subset_dia` tables continue to exist, but they are no longer associated with the `cust_subset` table. Instead they now have the status of regular database tables. To drop these two tables, enter the following statements:

```
DROP TABLE cust_subset_vio;  
DROP TABLE cust_subset_dia;
```

If you had previously issued the DROP TABLE statement without the RESTRICT keyword to successfully drop the `cust_subset` table, then the statements above would not be necessary, because dropping the target table in cascade mode implicitly drops any associated violations table and diagnostics table.

Privileges Required for Stopping a Violations Table

To stop a violations table (and for Dynamic Server, a diagnostics table) for a given target table, you must meet one of the following requirements:

- You must have the DBA privilege on the database.
- You must be the owner of the target table and have the Resource privilege on the database.
- You must have the Alter privilege on the target table and the Resource privilege on the database.

Related Information

Related statements: “SET Database Object Mode” on page 2-591 and “START VIOLATIONS TABLE” on page 2-663

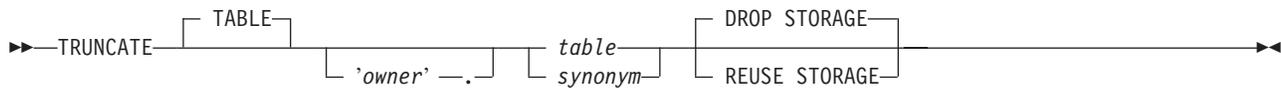
For a discussion of database object modes and violation detection, see the *IBM Informix Guide to SQL: Tutorial*.

TRUNCATE (IDS)

Use the TRUNCATE statement to quickly delete all rows from a local table and free the associated storage space. You can optionally reserve the space for the same table and its indexes. Only Dynamic Server supports this implementation of the TRUNCATE statement, which is an extension to the ANSI/ISO standard for SQL.

Extended Parallel Server supports a different implementation of TRUNCATE, as the section “TRUNCATE (IDS) ” on page 2-677 describes.

Syntax



Element	Description	Restrictions	Syntax
<i>owner</i>	Owner of <i>table</i> or <i>synonym</i>	See Usage notes.	“ Owner Name ” on page 5-45
<i>synonym</i>	Synonym for the table from which to remove all data	Must exist, and USETABLENAME must not be set	“Identifier ” on page 5-21
<i>table</i>	Name of table from which to remove all data and all B-tree structures of its indexes	Must exist in the database	“Identifier ” on page 5-21

Usage

The TRUNCATE statement rapidly deletes from a local table all active data rows and the B-tree structures of indexes on the table. You have the option of releasing the storage space that was occupied by the rows and index extents, or of reusing the same space when the table is subsequently repopulated with new rows.

To execute the TRUNCATE statement, you must be the owner of the table, or else hold DBA access privilege on the database. You must also hold Delete privilege on the table. If an enabled Delete trigger is defined on the table, the Alter privilege is also required, even though TRUNCATE does not activate triggers.

Although it requires the Delete privilege, TRUNCATE is a data definition language (DDL) statement. Like other DDL statements, TRUNCATE cannot operate on any table outside the database to which you are connected, nor on a table that a concurrent session is reading in Dirty Read isolation mode.

Dynamic Server always logs the TRUNCATE operation, even for a non-logging table. In databases that support transaction logging, only the COMMIT WORK or ROLLBACK WORK statement of SQL is valid after the TRUNCATE statement within the same transaction. Here the ROLLBACK statement must cancel the entire transaction that includes the TRUNCATE statement. Dynamic Server issues an error if ROLLBACK TO SAVEPOINT (or any other SQL statement except for COMMIT WORK or ROLLBACK WORK without the TO SAVEPOINT clause) immediately follows the TRUNCATE statement.

When you successfully rollback the TRUNCATE statement, no rows are removed from the table, and the storage extents that hold the rows and index partitions continue to be allocated to the table. Only databases with transaction logging can support the ROLLBACK WORK statement.

After the TRUNCATE statement successfully executes, Dynamic Server automatically updates the statistics and distributions for the table and for its indexes in the system catalog to show no rows in the table nor in its dbspace partitions. It is not necessary to run the UPDATE STATISTICS statement immediately after you commit the TRUNCATE statement.

If the table that the TRUNCATE statement specifies is a typed table, a successful TRUNCATE operation removes all the rows and B-tree structures from that table and from all its subtables within the table hierarchy.

The TRUNCATE statement does not reset the serial value of SERIAL, BIGSERIAL, or SERIAL8 columns. To reset the counter of a serial column, use the MODIFY clause of the ALTER TABLE statement, either before or after you execute the TRUNCATE statement.

The TABLE Keyword

The TABLE keyword has no effect on this statement, but it can be included to make your code more legible for human readers. Both of the following statements have the same effect, deleting all rows and any related index data from the **customer** table:

```
TRUNCATE TABLE customer;
```

```
TRUNCATE customer;
```

The Table Specification

You must specify the name or synonym of a table in the local database to which you are currently connected. If the **USETABLENAME** environment variable is set, you must use the name of the table, rather than a synonym. The table can be of type STANDARD, RAW, or TEMP, but you cannot specify the name or synonym of a view. (Categories of tables that are not valid with TRUNCATE are listed in the section “Restrictions” on page 2-681.)

In a database that is ANSI-compliant, you must specify the *owner* qualifier if you are not the owner of the *table* or *synonym*.

After the TRUNCATE statement begins execution, Dynamic Server attempts to place an exclusive lock on the specified table, to prevent other sessions from locking the table until the TRUNCATE statement is committed or rolled back. Exclusive locks are also applied to any dependent tables of the truncated table within a table hierarchy.

While concurrent sessions that use the Dirty Read isolation level are reading the table, however, the TRUNCATE statement fails with an RSAM -106 error. To reduce this risk, you can set the **IFX_DIRTY_WAIT** environment variable to specify that the TRUNCATE operation wait for a specified number of seconds for Dirty Read operations to commit or rollback.

The Storage Specification

By default, all of the partition extents that had been allocated to the specified table and to its indexes are released after TRUNCATE successfully executes, but you can include the DROP STORAGE keywords to achieve the same result, making the storage space available for other database objects. The first extent size remains the same as it was before the TRUNCATE operation released the storage.

Alternatively, if it is your intention to keep the same storage space allocated to the same table for subsequently loaded data, specify the REUSE STORAGE keywords to prevent the space from being deallocated. The REUSE STORAGE option of TRUNCATE can make storage management more efficient in applications where the same table is periodically emptied and reloaded with new rows.

This example truncates the **state** table and free all extents except the first extent:

```
TRUNCATE TABLE state DROP STORAGE;
```

The following example truncates the same table but only remove the actual data. All extents stay the same.

```
TRUNCATE TABLE state REUSE STORAGE;
```

Whether you specify `DROP STORAGE` or `REUSE STORAGE`, any out-of-row data values are released for all rows of the table when the `TRUNCATE` transaction is committed. Storage occupied by any BLOB or CLOB values that become unreferenced in the `TRUNCATE` transaction is also released.

The `AM_TRUNCATE` Purpose Function

Dynamic Server provides built-in `am_truncate` purpose functions for its primary access methods that support `TRUNCATE` operations on columns of permanent and temporary tables. It also provides a built-in `am_truncate` purpose function for its secondary access method for `TRUNCATE` operations on B-tree indexes.

For the `TRUNCATE` statement to work correctly in a virtual table interface (VTI) table requires a valid `am_truncate` purpose function in the primary access method for the data type of the VTI table. To register a new primary access method in the database, use the `CREATE PRIMARY ACCESS_METHOD` statement of SQL:

```
CREATE PRIMARY ACCESS_METHOD vti(  
  AM_GETNEXT = vti_getnext  
  AM_TRUNCATE = vti_truncate  
  ...);
```

You can also use the `ALTER ACCESS_METHOD` statement to add a valid `am_truncate` purpose function to an existing access method that has no `am_truncate` purpose function:

```
ALTER ACCESS_METHOD abc (ADD AM_TRUNCATE = abc_truncate);
```

In these examples, the `vti_truncate` and `abc_truncate` functions must be routines that support the functionality of the `AM_TRUNCATE` purpose option keyword, and that were previously registered in the database by the `CREATE FUNCTION` or `CREATE ROUTINE` statement.

Performance Advantages of `TRUNCATE`

The `TRUNCATE` statement is not equivalent to `DROP TABLE`. After `TRUNCATE` successfully executes, the specified *table* and all its columns and indexes are still registered in the database, but with no rows of data. In information management applications that require replacing all of the records in a table after some time interval, `TRUNCATE` requires fewer updates to the system catalog than the equivalent `DROP TABLE`, `CREATE TABLE`, and any additional DDL statements to redefine any synonyms, views, constraints, triggers, privileges, fragmentation schemes, and other attributes and associated database objects of the table.

In contexts where no existing rows of a table are needed, the `TRUNCATE` statement is typically far more efficient than using the `DELETE` statement with no `WHERE` clause to empty the table, because `TRUNCATE` requires fewer resources and less logging overhead than `DELETE`:

- `DELETE FROM table` deletes each row as a separately logged operation. If indexes exist on the table, each index must be updated when a row is deleted, and this update is also logged for each row. If an enabled Delete trigger is defined on the table, its triggered actions must also be executed and logged.

- `TRUNCATE table` performs the removal of all rows and of the B-tree structures of every index on the table as a single operation, and writes a single entry in the logical log when the transaction that includes `TRUNCATE` is committed or rolled back. The triggered action of any enabled trigger is ignored.

These performance advantages of `TRUNCATE` over `DELETE` are reduced when the table has one or more columns with the following attributes:

- Any simple large object data types stored in blobspaces
- Any `BLOB`, `CLOB`, complex, or user-defined types stored in sbspaces
- Any opaque types for which a **destroy** support function is defined.

Each of these features require the database server to read each row of the table, substantially reducing the speed of `TRUNCATE`.

If a table includes one or more UDTs for which you have registered an **am_truncate()** purpose function, then the performance difference between `TRUNCATE` and `DELETE` would reflect the relative costs of invoking the **am_truncate** interface once for `TRUNCATE` versus invoking the **destroy()** support function for each row.

As listed in the next section, certain conditions cause `TRUNCATE` to fail with an error. Some of these conditions have no effect on `DELETE` operations, so in those cases you can remove all rows more efficiently with a `DELETE` statement, as in the following operation on the **customer** table:

```
DELETE customer;
```

The `FROM` keyword that immediately follows `DELETE` can be omitted, as in this example, only if the **DELIMITED** environment variable is set.

Restrictions

The `TRUNCATE` statement fails if any of the following conditions exist:

- The user does not have Delete privilege on the table.
- The table has an enabled Delete trigger, but the user lacks the Alter privilege.
- The specified table or synonym does not exist in the local database.
- The specified synonym does not reference a table in the local database.
- The statement specifies a synonym for a local table, but the **USETABLENAME** environment variable is set.
- The statement specifies the name of a view or a synonym for a view.
- The table is a system catalog table or a system-monitoring interface (SMI) table.
- An R-tree index is defined on the table.
- The table is a virtual table (or has a virtual-index interface) for which no valid **am_truncate** access method exists in the database.
- An Enterprise Replication replicate that is not a master replicate is defined on the table. (For more information about replicates, see the *IBM Informix Dynamic Server Enterprise Replication Guide*.)
- A shared or exclusive lock on the table already exists.
- One or more cursors are open on the table.
- A concurrent session with Dirty Read isolation level is reading the table.
- Another table, with at least one row, has an enabled foreign-key constraint on the specified table. (An enabled foreign key constraint of another table that has no rows, however, has no effect on a `TRUNCATE` operation.)

Related Information

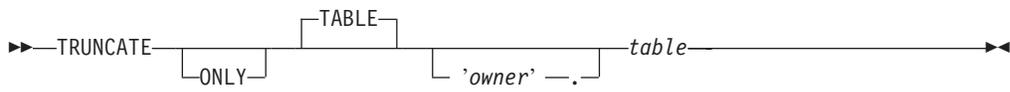
“DELETE ” on page 2-303, “DROP TABLE ” on page 2-342

TRUNCATE (XPS)

Use the TRUNCATE statement for quick removal of all rows from a table and all corresponding index data. Only Extended Parallel Server supports this implementation of the TRUNCATE statement, which is an extension to the ANSI/ISO standard for SQL.

Dynamic Server supports a different implementation of TRUNCATE, as described in the section “TRUNCATE (IDS) ” on page 2-677.

Syntax



Element	Description	Restrictions	Syntax
<i>table</i>	Name of table from which to remove all data	Must exist	“Identifier ” on page 5-21

Usage

You must be the owner of the table or have DBA privilege to use this statement.

The TRUNCATE statement does not automatically reset the serial value of a column. To reset the serial value of a column, you must do so explicitly, either before or after you run the TRUNCATE statement.

TRUNCATE is not equivalent to DROP TABLE. After TRUNCATE successfully executes, the specified *table* (and all its columns, and any synonyms, views, constraints, indexes, triggers, and access privileges) still exists in the database schema, but with no rows of data.

Restrictions

The TRUNCATE statement fails if any of the following conditions exist:

- One or more cursors are open on the table.
- Referential constraints exist on the table and any of the referencing tables has at least one row.
- A shared or exclusive lock on the table already exists.
- The statement references a view.
- The statement references any of the following types of tables:
 - An external table
 - A system catalog table
 - A violations table
- The statement is issued within a transaction.

Using the ONLY and TABLE Keywords

For Extended Parallel Server, the TABLE keyword and the ONLY keyword have no effect on this statement, but they can be included to make your code more legible for human readers. All of the following statements have the same effect, deleting all rows and any related index data from the **customer** table:

```
TRUNCATE ONLY TABLE customer
```

```
TRUNCATE TABLE customer
```

```
TRUNCATE ONLY customer
```

```
TRUNCATE customer
```

After the TRUNCATE Statement Executes

Information about the success of this statement appears in the logical-log files. For information about logical-log files, see your *IBM Informix Administrator's Guide*.

If the table was fragmented, after the statement executes, each fragment has a space allocated for it that is the same size as that of the first extent size. The fragment size of any indexes also corresponds to the size of the first extents.

Because the TRUNCATE statement does not alter the schema, the database server does not automatically update statistics. After you use this statement, you might want to issue an UPDATE STATISTICS statement on the table.

When You Might Use the TRUNCATE Statement

TRUNCATE performs similar operations to those that you can perform with the DELETE statement or with a combination of DROP TABLE and CREATE TABLE.

Using TRUNCATE can be faster than removing all rows from a table with the DELETE statement, because it does not activate any DELETE triggers. In addition, when you use TRUNCATE, the database server creates one log entry for the entire TRUNCATE statement, rather than one for each deleted row.

You might also use this statement instead of dropping a table and then re-creating it. When you drop and re-create a table, you must re-grant any access privileges that you want to preserve on the table. In addition, you must re-create any indexes, constraints, and triggers that were defined on the table. The TRUNCATE statement leaves these database objects and access privileges intact.

Related Information

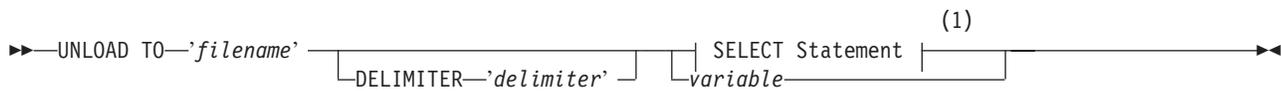
Related statements: "DELETE " on page 2-303 and "DROP TABLE " on page 2-342

UNLOAD

Use the UNLOAD statement to write the rows retrieved by a SELECT statement to an operating-system file. Use UNLOAD only with DB–Access

UNLOAD is an extension to the ANSI/ISO standard for SQL.

Syntax



Notes:

1 See “SELECT” on page 2-528

Element	Description	Restrictions	Syntax
<i>delimiter</i>	Quoted string to specify the field delimiter character in <i>filename</i> file	See “DELIMITER Clause” on page 2-687	“Quoted String” on page 4-189
<i>filename</i>	Operating-system file to receive the rows. Default pathname is the current directory.	See “UNLOAD TO File.”	“Quoted String” on page 4-189
<i>variable</i>	Host variable that contains the text of a valid SELECT statement	Must have been declared as a character data type	Language- specific

Usage

UNLOAD copies to a file the rows retrieved by a query. You must have the Select privilege on all columns specified in the SELECT statement. For information on database-level and table-level privileges, see “GRANT ” on page 2-395.

You can specify a literal SELECT statement or a character variable that contains the text of a SELECT statement. (See “SELECT” on page 2-528.)

The following example unloads rows whose value of **customer.customer_num** is greater than or equal to 138, and writes them to a file named **cust_file**:

```

UNLOAD TO 'cust_file' DELIMITER '!'
  SELECT * FROM customer WHERE customer_num >= 138;

```

The resulting output file, **cust_file**, contains two rows of data values:

```

138!Jeffery!Padgett!Wheel Thrills!3450 El Camino!Suite 10!Palo Alto!CA!94306!!
139!Linda!Lane!Palo Alto Bicycles!2344 University!!Palo Alto!CA!94301!
(415)323-5400

```

UNLOAD TO File

The UNLOAD TO file, as specified by the *filename* parameter, receives the retrieved rows. You can use an UNLOAD TO file as input to a LOAD statement.

In the default locale, data values have these formats in the UNLOAD TO file.

Data Type	Output Format
BOOLEAN	BOOLEAN values appear as either t for TRUE or f for FALSE.
Character	If a character field contains the delimiter, IBM Informix products automatically escape it with a backslash (\) to prevent interpretation as a special character. (If you use a LOAD statement to insert the rows into a table, backslash escape characters are automatically stripped.)
Collections	A collection is unloaded with its values between braces ({ }) and a delimiter between each element. For more information, see “ Unloading Complex Types (IDS)” on page 2-687.
DATE	DATE values are represented as <i>mm/dd/yyyy</i> (or the default format for the database locale), where <i>mm</i> is the month (January = 1, and so on), <i>dd</i> is the day, and <i>yyyy</i> is the year. If you have set the GL_DATE or DBDATE environment variable, the UNLOAD statement uses the specified date format for DATE values.

Data Type	Output Format
DATETIME, INTERVAL	Literal DATETIME and INTERVAL values appear as digits and delimiters, without keyword qualifiers, in the default format <i>yyyy-mm-dd hh:mi:ss.fff</i> . Time units outside the declared precision are omitted. If the GL_DATETIME or DBTIME environment variable is set, DATETIME values appear in the specified format.
DECIMAL, MONEY	Values are unloaded with no leading currency symbol. In the default locale, comma (,) is the <i>thousands</i> separator and period (.) is the decimal separator. If DBMONEY is set, UNLOAD uses its specified separators and currency format for MONEY values.
NULL	NULL appears as two delimiters with no characters between them.
Number	Values appear as literals, with no leading blanks. BIGINT, INTEGER, INT8, and SMALLINT zero appear as 0, and MONEY, FLOAT, SMALLFLOAT, and DECIMAL zero appear as 0.0.
ROW types (named and unnamed)	A ROW type is unloaded with its values enclosed between parentheses and a field delimiter separating each element. For more information, see “Unloading Complex Types (IDS)” on page 2-687.
Simple large objects (TEXT, BYTE)	TEXT and BYTE columns are unloaded directly into the UNLOAD TO file. BYTE values appear in ASCII hexadecimal form, with no added white space or newline characters. For more information, see “Unloading Simple Large Objects.”
Smart large objects (CLOB, BLOB)	CLOB and BLOB columns are unloaded into a separate operating-system file (for each column) on the client computer. The CLOB or BLOB field in the UNLOAD TO file contains the name of this file. For more information, see “Unloading Smart Large Objects (IDS)” on page 2-686.
User-defined data types (opaque types)	Opaque types must have an export() support function defined. They need special processing to copy data from the internal format of the opaque data type to the UNLOAD TO file format. An export binary support function might also be required for data in binary format. The data in the UNLOAD TO file would correspond to the format that the export() or exportbinary() support function returns.

For more information on **DB*** environment variables, refer to the *IBM Informix Guide to SQL: Reference*. For more information on **GL*** environment variables, refer to the *IBM Informix GLS User's Guide*.

In a nondefault locale, DATE, DATETIME, MONEY, and numeric column values have formats that the locale supports for these data types. For more information, see the *IBM Informix GLS User's Guide*.

Unloading Character Columns

In unloading files that contain VARCHAR or NVARCHAR columns, trailing blanks are retained in VARCHAR, LVARCHAR, or NVARCHAR fields. Trailing blanks are discarded when CHAR or NCHAR columns are unloaded.

For CHAR, VARCHAR, NCHAR, and NVARCHAR columns, an empty string (a data string of zero length, containing no characters) appears in the UNLOAD TO file as the four bytes “\ | ” (delimiter, backslash, blank space, delimiter).

Some earlier releases of Informix database servers used “||” (consecutive delimiters) to represent the empty string in LOAD and UNLOAD operations. In this release, however, “||” only represents NULL values in CHAR, VARCHAR, LVARCHAR, NCHAR, and NVARCHAR columns.

Unloading Simple Large Objects

The database server writes BYTE and TEXT values directly into the UNLOAD TO file. BYTE values are written in hexadecimal dump format with no added blank spaces or new line characters. The logical length of an UNLOAD TO file containing BYTE data can therefore be long and difficult to print or edit.

If you are unloading files that contain simple-large-object data types, do not use characters that can appear in BYTE or TEXT values as delimiters in the UNLOAD TO file. See also the section “DELIMITER Clause” on page 2-687.

The database server handles any required code-set conversions for TEXT data. For more information, see the *IBM Informix GLS User's Guide*.

If you are unloading files that contain simple-large-object data types, objects smaller than 10 kilobytes are stored temporarily in memory. You can adjust the 10-kilobyte setting to a larger setting with the **DBBLOBBUF** environment variable. BYTE or TEXT values larger than the default or the **DBBLOBBUF** setting are stored in a temporary file. For additional information about **DBBLOBBUF**, see the *IBM Informix Guide to SQL: Reference*.

Unloading Smart Large Objects (IDS)

The database server unloads smart large objects (from BLOB or CLOB columns) into a separate operating-system file for each column, in the same directory on the client computer as the UNLOAD TO file. All the smart blobs in the same column are stored in a single file. The filename has one of these formats:

- For a BLOB value: blob#####
- For a CLOB value: clob#####

In the preceding formats, the pound (#) symbols represent the digits of the unique hexadecimal smart-large-object identifier of the first smart large object in the file. The maximum number of digits for a smart-large-object identifier is 17. Most smart large objects, however, would have an identifier with fewer digits.

When the database server unloads the first smart large object, it creates the appropriate BLOB or CLOB client file with the hexadecimal identifier of the smart large object. If additional smart-large-object values are present in the same column, the database server writes their values to the same file, and lists in the UNLOAD TO file (*.unl) the sbpace, chunk, and page numbers, and smart large object identifier, for each BLOB or CLOB value in the file.

The following example shows an UNLOAD TO file entry for two smart large object values from the same column:

```
Object # 1  
Space Chunk Page = [5,6,3] Object ID = 1192071051
```

```
Object #2  
Space Chunk Page = [5,6,4] Object ID = 1192071050
```

both rows unloaded

In an UNLOAD TO file, a BLOB or CLOB column value appears in this format:

start_off,length,client_path

Here *start_off* is the starting offset (in hexadecimal format) of the smart-large-object value within the client file, *length* is the length (in hexadecimal) of the BLOB or CLOB value, and *client_path* is the pathname for the client file. No blank spaces can appear between these values. If a CLOB value is 512 bytes long and is at offset 256 in the **/usr/apps/clob9ce7.318** file, for example, then the CLOB value appears as follows in the UNLOAD TO file:

```
|100,200,/usr/apps/clob9ce7.318|
```

If a BLOB or CLOB column value occupies an entire client file, the CLOB or BLOB column value appears as follows in the UNLOAD TO file:

```
client_path
```

For example, if a CLOB value occupies the entire file `/usr/apps/clob9ce7.318`, the CLOB value appears as follows in the UNLOAD TO file:

```
|/usr/apps/clob9ce7.318|
```

For locales that support multibyte code sets, be sure that the declared size (in bytes) of any column that receives character data is large enough to store the entire data string. For some locales, this can require up to 4 times the number of logical characters in the longest data string.

The database server handles any required code-set conversions for CLOB data. For more information, see the *IBM Informix GLS User's Guide*.

Unloading Complex Types (IDS)

In an UNLOAD TO file, values of complex data types appear as follows:

- Collections are introduced with the appropriate constructor (MULTISET, LIST, SET), with their comma-separated elements enclosed in braces ({ }):

```
constructor{val1 , val2 , ... }
```

For example, to unload the SET values {1, 3, 4} from a column of the SET (INTEGER NOT NULL) data type, the corresponding field of the UNLOAD TO file appears as follows:

```
|SET{1 , 3 , 4}|
```

- ROW types (named and unnamed) are introduced by the ROW constructor and have their fields enclosed between parentheses and comma-separated:

```
ROW(val1 , val2 , ... )
```

For example, to unload the ROW values (1, 'abc'), the corresponding field of the UNLOAD TO file appears as follows:

```
|ROW(1 , abc)|
```

DELIMITER Clause

Use the DELIMITER clause to specify the delimiter that separates the data contained in each column in a row in the output file. If you omit this clause, then DB–Access checks the setting of the **DBDELIMITER** environment variable. If **DBDELIMITER** has not been set, the default delimiter is the pipe (|) symbol. You can specify TAB (CTRL-I) or a blank space (ASCII 32) as the delimiter symbol, but the following characters are not valid in any locale as delimiter symbols:

- Backslash (\)
- Newline character (CTRL-J)
- Hexadecimal digits (0 to 9, a to f, A to F)

The backslash (\) is not a valid field separator or record delimiter because it is the default escape character, indicating that the next character is a literal character in the data, rather than a special character.

The following UNLOAD statement specifies the semicolon (;) as the delimiter:

```
UNLOAD TO 'cust.out' DELIMITER ';'
  SELECT fname, lname, company, city FROM customer;
```

Related Information

Related statements: "LOAD " on page 2-443 and "SELECT" on page 2-528

For information about how to set environment variables, see the *IBM Informix Guide to SQL: Reference*.

For a discussion of the GLS aspects of the UNLOAD statement, see the *IBM Informix GLS User's Guide*.

For a task-oriented discussion of the UNLOAD statement and other utilities for moving data, see the *IBM Informix Migration Guide*.

UNLOCK TABLE

Use the UNLOCK TABLE statement in a database that does not use implicit transactions to unlock a table that you previously locked with the LOCK TABLE statement. The UNLOCK TABLE statement is not valid within a transaction.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax

►► UNLOCK TABLE table
synonym ◀◀

Element	Description	Restrictions	Syntax
<i>synonym</i>	Synonym for a table to unlock	The synonym and the table to which it points must exist	"Database Object Name " on page 5-16
<i>table</i>	Table to unlock	Must be in a database without transactions and must be a table that you previously locked	"Database Object Name " on page 5-16

Usage

You can lock a table if you own the table or if you have the Select privilege on the table, either from a direct grant to your user ID or from a grant to PUBLIC. You can only unlock a table that you locked. You cannot unlock a table that another process locked. Only one lock can apply to a table at a time.

You must specify the name or synonym of the table that you are unlocking. Do not specify the name of a view, or a synonym for a view.

To change the lock mode of a table in a database without transactions, use the UNLOCK TABLE statement to unlock the table, then issue a new LOCK TABLE statement. The following example shows how to change the lock mode of a table in a database that was created without transactions:

```
LOCK TABLE items IN EXCLUSIVE MODE;
...
UNLOCK TABLE items;
...
LOCK TABLE items IN SHARE MODE;
```

The UNLOCK TABLE statement fails if it is issued within a transaction. Table locks set within a transaction are released automatically when the transaction completes.

If you are using an ANSI-compliant database, do not issue an UNLOCK TABLE statement. The UNLOCK TABLE statement fails if it is issued within a transaction, and a transaction is always in effect in an ANSI-compliant database.

Related Information

Related statements: "BEGIN WORK " on page 2-74, "COMMIT WORK " on page 2-81, "LOCK TABLE " on page 2-449, and "ROLLBACK WORK" on page 2-520

For a discussion of concurrency and locks, see the *IBM Informix Guide to SQL: Tutorial*.

UPDATE

Use the UPDATE statement to change the values in one or more columns of one or more existing rows in a table or view.

With Dynamic Server, you can also use this statement to change the values in one or more elements in an Informix ESQL/C collection variable.

Syntax

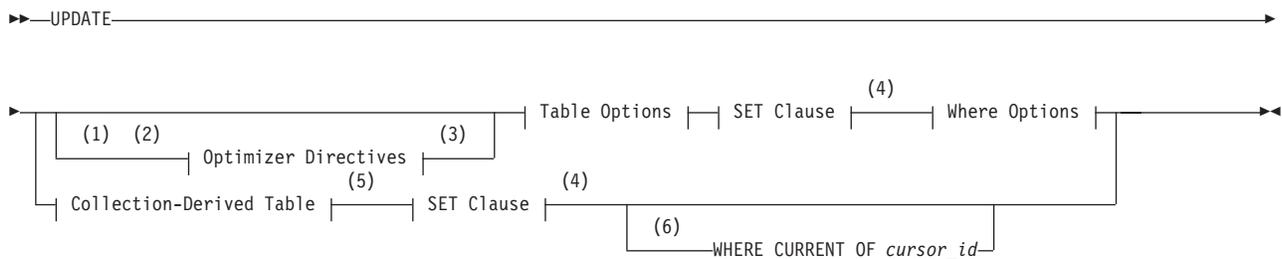
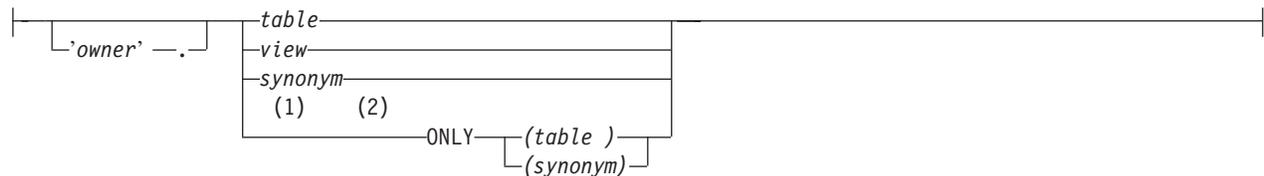
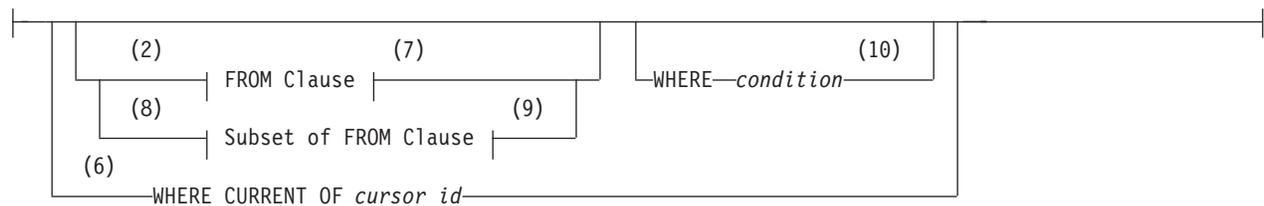


Table Options:



Where Options:



Notes:

- 1 Informix extension
- 2 Dynamic Server only
- 3 See “Optimizer Directives ” on page 5-35
- 4 See “SET Clause” on page 2-693
- 5 See “Collection-Derived Table” on page 5-4
- 6 ESQL/C or SPL only
- 7 See “FROM Clause” on page 2-544
- 8 Extended Parallel Server only
- 9 See “Subset of FROM Clause (XPS)” on page 2-699
- 10 See “WHERE Clause of UPDATE” on page 2-699

Element	Description	Restrictions	Syntax
<i>condition</i>	Logical criteria that updated rows must satisfy	Cannot be a UDR nor a correlated subquery	“Condition ” on page 4-5
<i>cursor_id</i>	Name of a cursor whose current row is to be updated	You cannot update a row that includes aggregates	“Identifier ” on page 5-21
<i>synonym, table, view</i>	Synonym, table, or view that contains the rows to be updated	The <i>synonym</i> and the table or view to which it points must exist	“Database Object Name ” on page 5-16

Usage

Use the UPDATE statement to update any of the following types of objects:

- A row in a table: a single row, a group of rows, or all rows in a table
- For Dynamic Server, an element in a collection variable
- An Informix ESQL/C **row** variable: a field or all fields

For information on how to update elements of a collection variable, see “Collection-Derived Table” on page 5-4. Sections that follow in this description of the UPDATE statement describe how to update a row in a table.

You must either own the table or have the Update privilege for the table; see “GRANT ” on page 2-395. To update data in a view, you must have the Update privilege, and the view must meet the requirements that are explained in “Updating Rows Through a View” on page 2-691.

The cursor (as defined in the SELECT ... FOR UPDATE portion of a DECLARE statement) can contain only column names. If you omit the WHERE clause, all rows of the target table are updated.

If you are using effective checking and the checking mode is set to IMMEDIATE, all specified constraints are checked at the end of each UPDATE statement. If the checking mode is set to DEFERRED, all specified constraints are *not* checked until the transaction is committed.

In Extended Parallel Server, if UPDATE is constructed in such a way that a single row might be updated more than once, the database server returns an error. If the

new value is the same in every update, however, the database server allows the update to take place without reporting an error.

In DB-Access, if you omit the WHERE clause and are in interactive mode, DB-Access does not run the UPDATE statement until you confirm that you want to change all rows. If the statement is in a command file, however, and you are running at the command line, the statement executes immediately.

Using the ONLY Keyword (IDS)

If you use the UPDATE statement to update rows of a supertable, rows from its subtables can also be updated. To update rows from the supertable only, use the ONLY keyword prior to the table name, as this example shows:

```
UPDATE ONLY(am_studies_super)
  WHERE advisor = "johnson"
  SET advisor = "camarillo";
```

Note: If you use the UPDATE statement on a supertable without the ONLY keyword and without a WHERE clause, all rows of the supertable and its subtables are updated. You cannot use the ONLY keyword if you plan to use the WHERE CURRENT OF clause to update the current row of the active set of a cursor.

Updating Rows Through a View

You can update data through a *single-table* view if you have the Update privilege on the view (see “GRANT ” on page 2-395). For a view to be updatable, the query that defines the view must not contain any of the following items:

- Columns in the projection list that are aggregate values
- Columns in the projection list that use the UNIQUE or DISTINCT keyword
- A GROUP BY clause
- A UNION operator

In addition, if a view is built on a table that has a derived value for a column, that column cannot be updated through the view. Other columns in the view, however, can be updated. In an updatable view, you can update the values in the underlying table by inserting values into the view.

```
CREATE VIEW cust_view AS SELECT * FROM customer;
UPDATE cust_view SET customer_num=10001 WHERE customer_cum=101;
```

The following statements define a view that includes all the rows in the **customer** table and changes the **customer_num** value to 10001 in any row where the value of that column is 101:

```
CREATE VIEW cust_view AS SELECT * FROM customer;
UPDATE cust_view SET customer_num=10001 WHERE customer_num=101;
```

You can use data-integrity constraints to prevent users from updating values in the underlying table when the update values do not fit the SELECT statement that defined the view. For more information, see “WITH CHECK OPTION Keywords” on page 2-277.

Because duplicate rows can occur in a view even if its base table has unique rows, be careful when you update a table through a view. For example, if a view is defined on the **items** table and contains only the **order_num** and **total_price** columns, and if two items from the same order have the same total price, the view contains duplicate rows. In this case, if you update one of the two duplicate **total_price** values, you have no way to know which item price is updated.

Important: If you are using a view with a check option, you cannot update rows in a remote table.

For Dynamic Server, an alternative to directly modifying data values in a view with the UPDATE statement is to create an INSTEAD OF trigger on the view. For more information, see “INSTEAD OF Triggers on Views (IDS)” on page 2-270.

Updating Rows in a Database Without Transactions

If you are updating rows in a database without transactions, you must take explicit action to restore updated rows. For example, if the UPDATE statement fails after updating some rows, the successfully updated rows remain in the table. You cannot automatically recover from a failed update.

Updating Rows in a Database with Transactions

If you are updating rows in a database with transactions, and you are using transactions, you can undo the update using the ROLLBACK WORK statement. If you do not execute a BEGIN WORK statement before the update, and the update fails, the database server automatically rolls back any database modifications made since the beginning of the update.

You can create temporary tables with the WITH NO LOG option. These tables are not logged and are not recoverable.

Tables that you create with the RAW logging type are not logged. Thus, RAW tables are not recoverable, even if the database uses logging. For information about RAW tables, refer to the *IBM Informix Guide to SQL: Reference*.

In an ANSI-compliant database, transactions are implicit, and all database modifications take place within a transaction. In this case, if an UPDATE statement fails, you can use ROLLBACK WORK to undo the update.

If you are within an explicit transaction, and the update fails, the database server automatically undoes the effects of the update.

Locking Considerations

When a row is selected with the intent to update, the update process acquires an update lock. Update locks permit other processes to read, or *share*, a row that is about to be updated, but they do not allow those processes to update or delete it. Just before the update occurs, the update process *promotes* the shared lock to an exclusive lock. An exclusive lock prevents other processes from reading or modifying the contents of the row until the lock is released.

An update process can acquire an update lock on a row or on a page that has a shared lock from another process, but you cannot promote the update lock from shared to exclusive (and the update cannot occur) until the other process releases its lock.

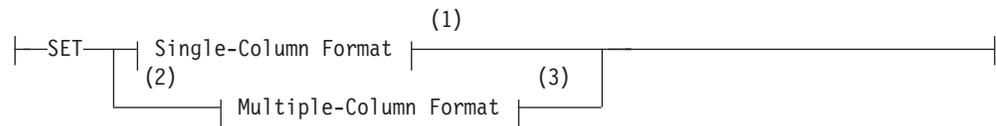
If the number of rows that a single update affects is large, you can exceed the limits placed on the maximum number of simultaneous locks. If this occurs, you can reduce the number of transactions per UPDATE statement, or you can lock the page or the entire table before you execute the statement.

SET Clause

Use the SET clause to identify the columns to update and assign values to each column. The clause supports the following formats:

- A single-column format, which pairs each column with a single expression
- A multiple-column format, which associates a list of multiple columns with the values returned by one or more expressions

SET Clause:



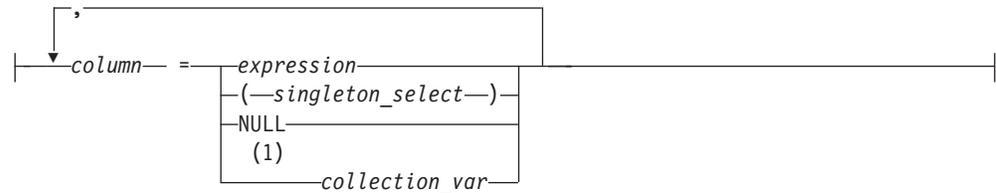
Notes:

- 1 See "Single-Column Format"
- 2 Informix extension
- 3 See "Multiple-Column Format" on page 2-694

Single-Column Format

Use the single-column format to pair one column with a single expression.

Single-Column Format:



Notes:

- 1 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>column</i>	Column to be updated	Cannot be a serial data type	"Identifier" on page 5-21
<i>collection_var</i>	Host or program variable	Must be declared as a collection data type	Language specific
<i>expression</i>	Returns a value for <i>column</i>	Cannot contain aggregate functions	"Expression" on page 4-40
<i>singleton_select</i>	Subquery that returns exactly one row	Returned subquery values must have a 1-to-1 correspondence with <i>column</i> list	"SELECT" on page 2-528

You can use this syntax to update a column that has a ROW data type.

You can include any number of "single *column* = single *expression*" terms. The *expression* can be an SQL subquery (enclosed between parentheses) that returns a single row, provided that the corresponding *column* is of a data type that can store the value (or the set of values) from the row that the subquery returns.

To specify values of a ROW-type column in a SET clause, see “Updating ROW-Type Columns (IDS)” on page 2-696. The following examples illustrate the single-column format of the SET clause.

```
UPDATE customer
  SET address1 = '1111 Alder Court', city = 'Palo Alto',
      zipcode = '94301' WHERE customer_num = 103;
```

```
UPDATE stock
  SET unit_price = unit_price * 1.07;
```

Using a Subquery to Update a Single Column

You can update the column specified in the SET clause with the value that a subquery returns.

```
UPDATE orders
  SET ship_charge =
    (SELECT SUM(total_price) * .07 FROM items
     WHERE orders.order_num = items.order_num)
  WHERE orders.order_num = 1001;
```

In Dynamic Server, if you are updating a supertable in a table hierarchy, the SET clause cannot include a subquery that references a subtable. If you are updating a subtable in a table hierarchy, a subquery in the SET clause can reference the supertable if it references only the supertable. That is, the subquery must use the SELECT ... FROM ONLY (*supertable*) syntax.

Updating a Column to NULL

Use the NULL keyword to modify a column value when you use the UPDATE statement. For example, for a customer whose previous address required two address lines but now requires only one, you would use the following entry:

```
UPDATE customer
  SET address1 = '123 New Street',
      address2 = null,
      city = 'Palo Alto',
      zipcode = '94303'
  WHERE customer_num = 134;
```

Updating the Same Column Twice

You can specify the same column more than once in the SET clause. If you do so, the column is set to the last value that you specified for the column. In the next example, the **fname** column appears twice in the SET clause. For the row where the customer number is 101, the user sets **fname** first to gary and then to harry. After the UPDATE statement executes, the value of **fname** is harry.

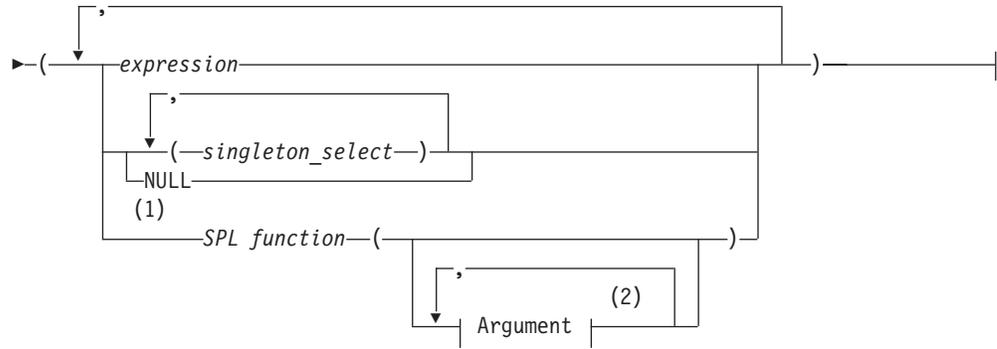
```
UPDATE customer
  SET fname = "gary", fname = "harry"
  WHERE customer_num = 101;
```

Multiple-Column Format

Use the multiple-column format of the SET clause to list multiple columns and set them equal to corresponding expressions.

Multiple-Column Format:





Notes:

- 1 Extended Parallel Server only
- 2 See "Arguments " on page 5-1

Element	Description	Restrictions	Syntax
<i>column</i>	Name of a column to be updated	Cannot have a serial or ROW type. The number of <i>column</i> names must equal the number of values returned to the right of the = sign.	"Identifier " on page 5-21
<i>expression</i>	Expression that returns a value for a <i>column</i>	Cannot include aggregate functions	"Expression" on page 4-40
<i>singleton_select</i>	Subquery that returns exactly one row	Values that the subquery returns must correspond to columns in the <i>column</i> list	"SELECT" on page 2-528
<i>SPL function</i>	SPL routine that returns one or more values	Returned values must have a 1-to-1 correspondence to columns in <i>the column</i> list	"Identifier " on page 5-21

The multiple-column format of the SET clause offers the following options for listing a set of columns that you intend to update:

- Explicitly list each column, placing commas between columns and enclosing the set of columns between parentheses.
- Implicitly list all columns in the table by using an asterisk (*).

You must list each expression explicitly, placing comma (,) separators between expressions and enclosing the set of expressions between parentheses. The number of columns must equal the number of values returned by the expression list, unless the expression list includes an SQL subquery.

The following examples show the multiple-column format of the SET clause:

```
UPDATE customer
  SET (fname, lname) = ('John', 'Doe') WHERE customer_num = 101;
```

```
UPDATE manufact
  SET * = ('HNT', 'Hunter') WHERE manu_code = 'ANZ';
```

Using a Subquery to Update Multiple Column Values

The expression list can include one or more subqueries. Each must return a single row containing one or more values. The number of columns that the SET clause explicitly or implicitly specifies must equal the number of values returned by the expression (or expression list) that follows the equal (=) sign in the multiple-column SET clause.

The subquery must be enclosed between parentheses. These parentheses are nested within the parentheses that immediately follow the equal (=) sign. If the expression list includes multiple subqueries, each subquery must be enclosed between parentheses, with a comma (,) separating successive subqueries:

```
UPDATE ... SET ... = ((subqueryA),(subqueryB), ... (subqueryN))
```

The following examples show the use of subqueries in the SET clause:

```
UPDATE items
  SET (stock_num, manu_code, quantity) =
      ( (SELECT stock_num, manu_code FROM stock
        WHERE description = 'baseball'), 2)
  WHERE item_num = 1 AND order_num = 1001;

UPDATE table1
  SET (col1, col2, col3) =
      ((SELECT MIN (ship_charge), MAX (ship_charge) FROM orders), '07/01/2007')
  WHERE col4 = 1001;
```

In Dynamic Server, if you are updating a supertable in a table hierarchy, the SET clause cannot include a subquery that references one of its subtables. If you are updating a subtable in a table hierarchy, a subquery in the SET clause can reference the supertable if it references only the supertable. That is, the subquery must use the SELECT... FROM ONLY (*supertable*) syntax.

Using an SPL Function to Update Column Values (XPS)

When you use an SPL function to update column values, the returned values of the function must have a one-to-one correspondence with the listed columns. That is, each value that the SPL function returns must be of the data type expected by the corresponding column in the column list.

If the called SPL routine contains certain SQL statements, a runtime error occurs. For information on which SQL statements cannot be used in an SPL routine that is called in a data-manipulation statement, see “Restrictions on SPL Routines in Data-Manipulation Statements” on page 5-78. In the next example, the SPL function `p2()` updates the `i2` and `c2` columns of the `t2` table:

```
CREATE PROCEDURE p2() RETURNING int, char(20);
  RETURN 3, 'three';
END PROCEDURE;
UPDATE t2 SET (i2, c2) = (p2()) WHERE i2 = 2;
```

In Extended Parallel Server, you create an SPL function with the CREATE PROCEDURE statement. The CREATE FUNCTION statement is not available.

Updating ROW-Type Columns (IDS)

Use the SET clause to update a named or unnamed ROW-type column. For example, suppose you define the following named ROW type and a table that contains columns of both named and unnamed ROW types:

```
CREATE ROW TYPE address_t
(
  street CHAR(20), city CHAR(15), state CHAR(2)
);
CREATE TABLE empinfo
(
  emp_id INT
  name ROW ( fname CHAR(20), lname CHAR(20)),
  address address_t
);
```

To update an unnamed ROW type, specify the ROW constructor before the parenthesized list of field values.

The following statement updates the **name** column (an unnamed ROW type) of the **empinfo** table:

```
UPDATE empinfo SET name = ROW('John','Williams') WHERE emp_id =455;
```

To update a named ROW type, specify the ROW constructor before the list (in parentheses) of field values, and use the cast (::) operator to cast the ROW value as a named ROW type. The following statement updates the **address** column (a named ROW type) of the **empinfo** table:

```
UPDATE empinfo
  SET address = ROW('103 Baker St','Tracy','CA')::address_t
  WHERE emp_id = 3568;
```

For more information on the syntax for ROW constructors, see “Constructor Expressions (IDS)” on page 4-82. See also “Literal Row” on page 4-186.

The ROW-column SET clause can only support literal values for fields. To use an ESQL/C variable to specify a field value, you must select the ROW data into a **row** variable, use host variables for the individual field values, then update the ROW column with the **row** variable. For more information, see “Updating a Row Variable (IDS, ESQL/C)” on page 2-703.

You can use Informix ESQL/C host variables to insert *non-literal* values as:

- An entire row type into a column
Use a **row** variable as a variable name in the SET clause to update all fields in a ROW column at one time.
- Individual fields of a ROW type
To insert non-literal values into a ROW-type column, you can first update the elements in a **row** variable and then specify the **collection** variable in the SET clause of an UPDATE statement.

When you use a **row** variable in the SET clause, the **row** variable must contain values for each field value. For information on how to insert values into a **row** variable, see “Updating a Row Variable (IDS, ESQL/C)” on page 2-703.

You can use the UPDATE statement to modify only some of the fields in a row:

- Specify the field names with field projection for all fields whose values remain unchanged.

For example, the following UPDATE statement changes only the **street** and **city** fields of the **address** column of the **empinfo** table:

```
UPDATE empinfo
  SET address = ROW('23 Elm St', 'Sacramento', address.state)
  WHERE emp_id = 433;
```

The **address.state** field remains unchanged.

- Select the row into an ESQL/C **row** variable and update the desired fields.
For more information, see “Updating a Row Variable (IDS, ESQL/C)” on page 2-703.

Updating Collection Columns (IDS)

You can use the SET clause to update values in a collection column. For more information, see “Collection Constructors” on page 4-83.

A collection variable can update a collection-type column. With a collection variable, you can insert one or more individual elements of a collection. For more information, see “Collection-Derived Table” on page 5-4.

For example, suppose you define the **tab1** table as follows:

```
CREATE TABLE tab1
(
  int1 INTEGER,
  list1 LIST(ROW(a INTEGER, b CHAR(5)) NOT NULL),
  dec1 DECIMAL(5,2)
);
```

The following UPDATE statement updates a row in **tab1**:

```
UPDATE tab1
SET list1 = LIST{ROW(2, 'zyxwv'),
  ROW(POW(2,6), '=64'),
  ROW(ROUND(ROOT(146)), '=12')},
WHERE int1 = 10;
```

Collection column **list1** in this example has three elements. Each element is an unnamed ROW type with an INTEGER field and a CHAR(5) field. The first element includes two literal values: an integer (2) and a quoted string ('zyxwv').

The second and third elements also use a quoted string to indicate the value for the second field. They each designate the value for the first field with an expression, however, rather than with a literal value.

Updating Values in Opaque-Type Columns (IDS)

Some opaque data types require special processing when they are updated. For example, if an opaque data type contains spatial or multirepresentational data, it might provide a choice of how to store the data: inside the internal structure or, for large objects, in a smart large object.

This processing is accomplished by calling a user-defined support function called **assign()**. When you execute UPDATE on a table whose rows contain one of these opaque types, the database server automatically invokes the **assign()** function for the type. This function can make the decision of how to store the data. For more information about the **assign()** support function, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

Data Types in Distributed UPDATE Operations (IDS)

In Dynamic Server, an UPDATE (or any other SQL data-manipulation language statement) that accesses a database of another Dynamic Server instance can reference only the following data types:

- Built-in data types that are not opaque
- BOOLEAN
- LVARCHAR
- DISTINCT of the built-in data types that are not opaque
- DISTINCT of BOOLEAN
- DISTINCT of LVARCHAR
- DISTINCT of any DISTINCT data type that appears in this list.

Cross-server distributed UPDATE operations can support these DISTINCT types only if the DISTINCT types are cast explicitly to built-in types, and all of the DISTINCT types, their data type hierarchies, and their casts are defined exactly the same way in each participating database. For additional information about the data

types that Dynamic Server supports in cross-server DML operations, see “Data Types in Cross-Server Transactions” on page 2-537.

Distributed operations that access other databases of the local Dynamic Server instance, however, can access the cross-server data types in the preceding list, and also the following data types:

- Most of the *built-in opaque data types*, as listed in “Data Types in Cross-Database Transactions” on page 2-536
- DISTINCT of the built-in types that are referenced in the preceding line
- DISTINCT of any of the data types that are listed in either of the two preceding lines
- Opaque user-defined data types (UDTs) that are explicitly cast to built-in data types.

Cross-database distributed UPDATE operations can support these DISTINCT types and opaque UDTs only if all the opaque UDTs and DISTINCT types are cast explicitly to built-in types, and all of the opaque UDTs, DISTINCT types, data type hierarchies, and casts are defined exactly the same way in each of the participating databases.

Updates cannot access the database of another database server unless both servers define TCP/IP connections in DBSERVERNAME or DBSERVERALIASES configuration parameters. This applies to any communication between Dynamic Server instances, even if both database servers reside on the same computer.

Subset of FROM Clause (XPS)

For Extended Parallel Server, UPDATE supports only a subset of the syntax listed in “FROM Clause” on page 2-544. You cannot include the LOCAL or the SAMPLES OF keywords if you are using Extended Parallel Server.

For Dynamic Server, the UPDATE statement supports the FROM clause syntax of the SELECT statement.

WHERE Clause of UPDATE

The WHERE clause lets you specify search criteria to limit the rows to be updated. If you omit the WHERE clause, every row in the table is updated. For more information, see the “WHERE Clause of SELECT” on page 2-560.

You can use a join to determine which column values to update by specifying a FROM clause. Columns from any table in the FROM clause can appear in the WHERE clause to provide values for the columns and rows to update. For example, in the following UPDATE statement, a FROM clause introduces tables to be joined in the WHERE clause:

```
UPDATE tab1 SET tab1.a = tab2.a FROM tab1, tab2, tab3
WHERE tab1.b = tab2.b AND tab2.c =tab3.c;
```

The next example uses WHERE and FROM clauses to update three columns (**state**, **zipcode**, and **phone**) in each row of the **customer** table that has a corresponding entry in a table of new addresses called **new_address**:

```
UPDATE customer
SET (state, zipcode, phone) =
((SELECT state, zipcode, phone FROM new_address N
WHERE N.cust_num = customer.customer_num))
WHERE customer_num IN
(SELECT cust_num FROM new_address);
```

SQLSTATE Values When Updating an ANSI-Compliant Database

If you update a table in an ANSI-compliant database with an UPDATE statement that contains the WHERE clause and no rows are found, the database server issues a warning.

You can detect this warning condition in either of the following ways:

- The GET DIAGNOSTICS statement sets the **RETURNED_SQLSTATE** field to the value 02000. In an SQL API application, the **SQLSTATE** variable contains this same value.
- In an SQL API application, the **sqlca.sqlcode** and **SQLCODE** variables contain the value 100.

The database server also sets **SQLSTATE** and **SQLCODE** to these values if the UPDATE ... WHERE statement is part of a multistatement PREPARE and the database server returns no rows.

SQLSTATE Values When Updating a Non-ANSI Database

In a database that is not ANSI compliant, the database server does not return a warning when it finds no matching rows for the WHERE clause of an UPDATE statement. The **SQLSTATE** code is 00000 and the **SQLCODE** code is zero (0). If the UPDATE ... WHERE statement is part of a multistatement PREPARE, however, and no rows are returned, the database server issues a warning, and sets **SQLSTATE** to 02000 and sets **SQLCODE** to 100.

Client-server communication protocols of Dynamic Server, such as SQLI and DRDA, support **SQLSTATE** code values. For a list of these codes, and for information about how to get the message text, see "Using the SQLSTATE Error Status Code" on page 2-387.

Subqueries in the WHERE Clause of UPDATE

The FROM clause of a subquery in the WHERE clause of the UPDATE statement can specify as a data source the same table or view that the Table Options clause of the UPDATE statement specifies. UPDATE operations with subqueries that reference the same table object are supported only if all of the following conditions are true:

- The subquery either returns a single row, or else has no correlated column references.
- The subquery is in the UPDATE statement WHERE clause, using Condition with Subquery syntax.
- No SPL routine in the subquery can reference the same table that UPDATE is modifying.

Unless all of these conditions are satisfied, UPDATE statements that include subqueries that reference the same table or view that the UPDATE statement modifies return error -360.

The following example updates the **stock** table by reducing the **unit_price** value by 5% for a subset of prices. The WHERE clause specifies which prices to reduce by applying the IN operator to the rows returned by a subquery that selects only the rows of the **stock** table where the **unit_price** value is greater than 50:

```
UPDATE stock SET unit_price = unit_price * 0.95
WHERE unit_price IN
(SELECT unit_price FROM stock WHERE unit_price > 50);
```

This subquery includes only uncorrelated column references, because its only referenced column is in a table specified in its FROM clause. The requirements listed above are in effect, because the data source of the subquery is the same **stock** table that the Table Options clause of the outer UPDATE statement specifies.

The previous example produces the same results as issuing two separate DML statements:

- The SELECT statement, to return a temporary table, **tmp1**, that contains the same rows from the **stock** table that the subquery returned.
- The UPDATE statement, to issue a subquery of the temporary table as a predicate in its WHERE clause to modify every row of the **stock** table where the **unit_price** matches a value in the temporary table:

```
SELECT unit_price FROM stock WHERE unit_price > 50 INTO TEMP tmp1;
UPDATE stock SET unit_price = unit_price * 0.95
WHERE unit_price IN ( SELECT * FROM tmp1 );
```

Here is an example of a more complex UPDATE statement that includes multiple uncorrelated subqueries in its WHERE clause:

```
UPDATE t1 SET a = a + 10
WHERE a > ALL (SELECT a FROM t1 WHERE a > 1) AND
a > ANY (SELECT a FROM t1 WHERE a > 10) AND
EXISTS (SELECT a FROM t1 WHERE a > 5);;
```

If an enabled Select trigger is defined on a table that is the data source of a subquery in the WHERE clause of an UPDATE statement that modifies the same table, executing that subquery within the UPDATE statement does not activate the Select trigger. Consider the following program fragment:

```
CREATE TRIGGER sel11 SELECT ON t1 BEFORE
(UPDATE d1
SET (c1, c2, c3, c4, c5) =
(c1 + 1, c2 + 1, c3 + 1, c4 + 1, c5 + 1));

UPDATE t2 SET c1 = c1 + 1
WHERE c1 IN
(SELECT t1.c1 from t1 WHERE t1.c1 > 10 );
```

In the example above, trigger **sel11** is not activated as part of the UPDATE operation on table **t2**.

A subquery in the WHERE clause of the UPDATE statement can include the UNION or the UNION ALL operator, as in the following example.

```
UPDATE t1 SET a = a + 10 WHERE a in (SELECT a FROM t1 WHERE a > 1
UNION SELECT a FROM t1, t2 WHERE a < b);
```

If the table that the outer UPDATE statement modifies a typed table within a table hierarchy, Dynamic Server supports all of the following operations that use valid subqueries in the WHERE clause of UPDATE:

- UPDATE on target parent table with subquery (SELECT from parent table)
- UPDATE on target parent table with subquery (SELECT from child table)
- UPDATE on target child table with subquery (SELECT from parent table)
- UPDATE on target child table with subquery (SELECT from child table).

The following program fragment illustrates UPDATE operations with subqueries on typed tables:

```

CREATE ROW TYPE r1 (c1 INT, c2 INT);
CREATE ROW TYPE r2 UNDER r1;
CREATE TABLE t1 OF TYPE r1; -- parent table
CREATE TABLE t2 OF TYPE r2 UNDER t1; -- child table

UPDATE t1 SET c1 = c1 + 1 WHERE c1 IN
  ( SELECT t1.c1 FROM t1 WHERE t1.c1 > 10);

UPDATE t1 SET c1 = c1 + 1 WHERE c1 IN
  ( SELECT t2.c1 FROM t2 WHERE t2.c1 > 10);

UPDATE t2 SET c1 = c1 + 1 WHERE c1 IN
  ( SELECT t2.c1 FROM t2 WHERE t2.c1 > 10);

UPDATE t2 SET c1 = c1 + 1 WHERE c1 IN
  ( SELECT t1.c1 FROM t1 WHERE t1.c1 > 10);

```

See the “Condition with Subquery” on page 4-15 topic for more information about how to use subqueries that return multiple rows as predicates in the WHERE clause of the UPDATE statement.

Using the WHERE CURRENT OF Clause (ESQL/C, SPL)

Use the WHERE CURRENT OF clause to update the current row of the active set of a cursor in the current element of a Collection cursor.

The UPDATE statement does not advance the cursor to the next row, so the current row position remains unchanged.

For table hierarchies of Dynamic Server, you cannot use this clause if you are selecting from only one table in a table hierarchy. That is, you cannot use this option if you use the ONLY keyword.

In ESQL/C routines, to include the WHERE CURRENT OF keywords, you must have previously used the DECLARE statement to define the *cursor* with the FOR UPDATE option. If the DECLARE statement that created the cursor specified one or more columns in the FOR UPDATE clause, you are restricted to updating only those columns in a subsequent UPDATE ... WHERE CURRENT OF statement. The advantage to specifying columns in the FOR UPDATE clause of a DECLARE statement is speed. The database server can usually perform updates more quickly if columns are specified in the DECLARE statement.

In SPL routines, you can specify a cursor after the WHERE CURRENT OF keywords in an UPDATE statement only if you declared the *cursor_id* in the FOREACH statement of SPL. You cannot declare an Update cursor in the DECLARE statement of an SPL routine.

Note: An Update cursor can perform updates that are not possible with the UPDATE statement.

The following Informix ESQL/C example illustrates the CURRENT OF form of the WHERE clause. In this example, updates are performed on a range of customers who receive 10-percent discounts (assume that a new column, **discount**, is added to the **customer** table). The UPDATE statement is prepared outside the WHILE loop to ensure that parsing is done only once.

```

char answer [1] = 'y';
EXEC SQL BEGIN DECLARE SECTION;
  char fname[32],lname[32];
  int low,high;
EXEC SQL END DECLARE SECTION;

```

```

main()
{
    EXEC SQL connect to 'stores_demo';
    EXEC SQL prepare sel_stmt from
        'select fname, lname from customer
        where cust_num between ? and ? for update';
EXEC SQL declare x cursor for sel_stmt;
printf("\nEnter lower limit customer number: ");
scanf("%d", &low);
printf("\nEnter upper limit customer number: ");
scanf("%d", &high);
EXEC SQL open x using :low, :high;
EXEC SQL prepare u from
    'update customer set discount = 0.1 where current of x';
while (1)
{
    EXEC SQL fetch x into :fname, :lname;
    if ( SQLCODE == SQLNOTFOUND) break;
}
printf("\nUpdate %.10s %.10s (y/n)?", fname, lname);
if (answer = getch() == 'y')
    EXEC SQL execute u;
EXEC SQL close x;
}

```

Updating a Row Variable (IDS, ESQL/C)

The UPDATE statement with the Collection-Derived Table segment allows you to update fields in a **row** variable. The Collection-Derived Table segment identifies the **row** variable in which to update the fields. For more information, see “Collection-Derived Table” on page 5-4.

To update fields

1. Create a **row** variable in your Informix ESQL/C program.
2. Optionally, select a ROW-type column into the **row** variable with the SELECT statement (without the Collection-Derived Table segment).
3. Update fields of the **row** variable with the UPDATE statement and the Collection-Derived Table segment.
4. After the **row** variable contains the correct fields, you then use the UPDATE or INSERT statement on a table or view name to save the **row** variable in the ROW column (named or unnamed).

The UPDATE statement and the Collection-Derived Table segment allow you to update a field or a group of fields in the **row** variable. Specify the new field values in the SET clause. For example, the following UPDATE changes the **x** and **y** fields in the **myrect** Informix ESQL/C **row** variable:

```

EXEC SQL BEGIN DECLARE SECTION;
    row (x int, y int, length float, width float) myrect;
EXEC SQL END DECLARE SECTION;
. . .
EXEC SQL select into :myrect from rectangles where area = 64;
EXEC SQL update table(:myrect) set x=3, y=4;

```

Suppose that after the SELECT statement, the **myrect2** variable has the values **x=0**, **y=0**, **length=8**, and **width=8**. After the UPDATE statement, the **myrect2** variable has field values of **x=3**, **y=4**, **length=8**, and **width=8**. You cannot use a **row** variable in the Collection-Derived Table segment of an INSERT statement.

You can, however, use the UPDATE statement and the Collection-Derived Table segment to insert new field values into a **row** host variable, if you specify a value for every field in the row.

For example, the following code fragment inserts new field values into the **row** variable **myrect** and then inserts this **row** variable into the database:

```
EXEC SQL update table(:myrect)
    set x=3, y=4, length=12, width=6;
EXEC SQL insert into rectangles
    values (72, :myrect);
```

If the **row** variable is an untyped variable, you must use a SELECT statement *before* the UPDATE so that Informix ESQL/C can determine the data types of the fields. An UPDATE of fields in a **row** variable cannot include a WHERE clause.

The **row** variable can store the field values of the row, but it has no intrinsic connection with a database column. Once the **row** variable contains the correct field values, you must then save the variable into the ROW column with one of the following SQL statements:

- To update the ROW column in the table with contents of the **row** variable, use an UPDATE statement on a table or view name and specify the **row** variable in the SET clause. (For more information, see “Updating ROW-Type Columns (IDS)” on page 2-696.)
- To insert a **row** into a column, use the INSERT statement on a table or view name and specify the **row** variable in the VALUES clause. (For more information, see “Inserting Values into ROW-Type Columns (IDS)” on page 2-437.)

For examples of SPL ROW variables, see the *IBM Informix Guide to SQL: Tutorial*. For more information on using Informix ESQL/C **row** variables, see the discussion of complex data types in the *IBM Informix ESQL/C Programmer's Manual*.

Examples

```
CREATE VIEW cust_view AS SELECT * FROM customer;
UPDATE cust_view SET customer_num=10001 WHERE customer_num=101;
```

Related Information

Related statements: “DECLARE ” on page 2-286, “INSERT ” on page 2-430, “OPEN ” on page 2-459, “SELECT” on page 2-528, and “FOREACH ” on page 3-22

For a task-oriented discussion of the UPDATE statement, see the *IBM Informix Guide to SQL: Tutorial*.

For a discussion of the GLS aspects of the UPDATE statement, see the *IBM Informix GLS User's Guide*.

For information on how to access row and collections with Informix ESQL/C host variables, see the discussion of complex data types in the *IBM Informix ESQL/C Programmer's Manual*.

UPDATE STATISTICS

Use the UPDATE STATISTICS statement to perform any of the following tasks:

- Calculate the distribution of column values.
- Update system catalog tables that the database server uses to optimize queries.
- Force reoptimization of SPL routines.
- Convert existing indexes when you upgrade the database server.

This statement is an extension to the ANSI/ISO standard for SQL.

Syntax

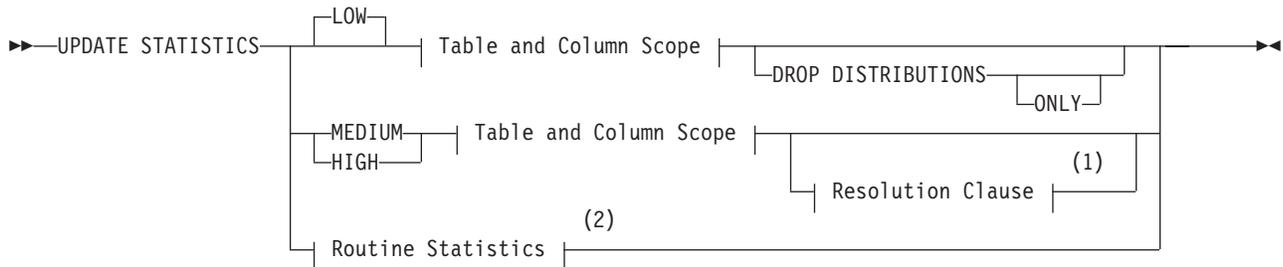
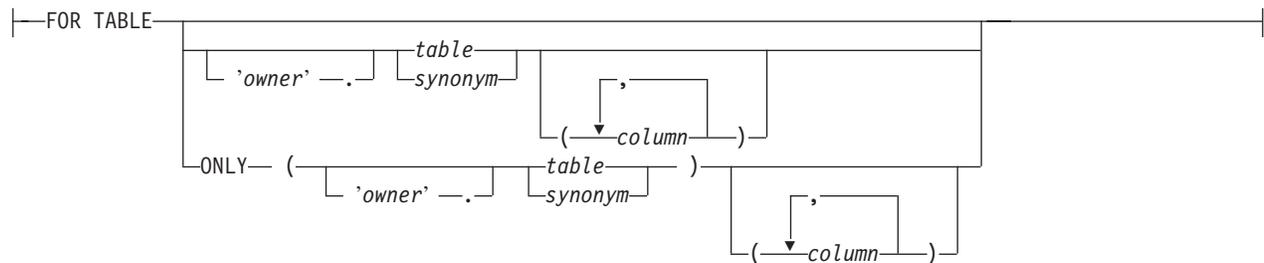


Table and Column Scope:



Notes:

- 1 See "Resolution Clause" on page 2-711
- 2 See "Routine Statistics" on page 2-714

Element	Description	Restrictions	Syntax
<i>column</i>	A column in <i>table</i> or <i>synonym</i>	Must exist. With MEDIUM or HIGH keywords, <i>column</i> cannot be of BYTE or TEXT data type	"Identifier " on page 5-21
<i>owner</i>	The <i>owner</i> of <i>table</i> or <i>synonym</i>	Must be the owner of <i>table</i> or <i>synonym</i>	" Owner Name " on page 5-45
<i>synonym</i>	Synonym for a table whose statistics are to be updated	The <i>synonym</i> and the table to which it points must exist in the current database	"Identifier " on page 5-21
<i>table</i>	Table for which statistics are to be updated	Must exist in the current database or be a temporary table created in the current session	"Identifier " on page 5-21

Usage

Use the UPDATE STATISTICS statement to update system catalog information that the query optimizer uses for operations on objects in the local database.

You cannot, however, update the statistics for a table or the query plan of a UDR in any database except the current database. That is, the database server ignores remote database objects when executing the UPDATE STATISTICS statement.

Scope of UPDATE STATISTICS

The scope of UPDATE STATISTICS is restricted by whatever tables, columns, or SPL routines follow the FOR TABLE keywords or the FOR PROCEDURE keywords.

- If you include no FOR PROCEDURE specification, no Table and Column Scope clause, and no Resolution clause, then statistics are updated for every table and SPL routine in the current database, including the system catalog tables.
- If you use the FOR TABLE keywords without also specifying the name or synonym of a table, the database server recalculates distributions on all of the tables in the current database, and on all of the temporary tables in your session.
- If you specify a table after the FOR TABLE keywords without also specifying a list of columns, the database server recalculates distributions on all of the columns of the specified table.
- If you include the FOR PROCEDURE keywords, but do not specify the name of any SPL routine, the database server reoptimizes the query plans of all SPL routines in the current database.

In databases of Extended Parallel Server, the UPDATE STATISTICS statement does not update, maintain, or collect statistics on indexes, and it does not update the **syscolumns** or **sysindexes** tables. References to indexes or to the **syscolumns** or **sysindexes** system catalog tables in sections that follow do not apply to Extended Parallel Server.

Updating Statistics for Tables

Although a change to the database might make information in the **sysables**, **syscolumns**, **sysindexes**, and **sysdistrib** system catalog tables obsolete, the database server does not automatically update those tables. Issue an UPDATE STATISTICS statement in the following situations to ensure that the stored distribution information reflects the state of the database:

- You perform extensive modifications to a table.
- An application changes the distribution of column values.
UPDATE STATISTICS reoptimizes queries on the modified objects.
- You upgrade a database for use with a newer database server.

The UPDATE STATISTICS statement converts the old indexes to conform to the newer database server index format and implicitly drops the old indexes.

You can convert the indexes table by table or for the entire database at one time. Follow the conversion guidelines in the *IBM Informix Migration Guide*.

If your application makes many modifications to the data in a particular table, update the system catalog for that table routinely with UPDATE STATISTICS to improve query efficiency. The term *many modifications* is relative to the resolution of the distributions. If the data modifications have little effect on the distribution of column values, you do not need to execute UPDATE STATISTICS.

Automated Table Statistics Maintenance

To simplify the complex and repetitive task of the DBA in maintaining current table statistics from which the query plan optimizer can design efficient query plans, Dynamic Server provides a table statistics maintenance system, called Auto Update Statistics (AUS). This can automate the identification of tables whose

statistics are stale, and can automate the construction and execution of the corresponding UPDATE STATISTICS statements to recalculate their column distributions. The AUS system is provided with built-in criteria for when table statistics should be updated, but the DBA can modify these criteria to reflect current requirements and workloads.

For more information, see the description of the Auto Update Statistics maintenance system in the *IBM Informix Performance Guide*. See also the description of the Scheduler, which the DBA can use to specify the policies, resources, and frequency with which the AUS system recalculates table statistics, in the *IBM Informix Dynamic Server Administrator's Guide*.

The AUS maintenance system for table statistics is also available in the OpenAdmin Tool for IDS (OAT). Refer to the OAT online help for detailed information on how to configure the AUS maintenance system to provide current table statistics automatically. OAT is available as an open source download from the iiug.org website and from IBM websites.

Using the FOR TABLE ONLY Keywords

Use the FOR TABLE ONLY keywords to collect data for a single table within a hierarchy of typed tables. If you do not include the ONLY keyword immediately after FOR TABLE when the specified *table* has subtables, Dynamic Server creates distributions for that table and also for every subtable under it in the hierarchy.

For example, suppose your database has the typed table hierarchy that appears in Figure 2-2, which shows a supertable named **employee** that has a subtable named **sales_rep**. The **sales_rep** table, in turn, has a subtable named **us_sales_rep**.

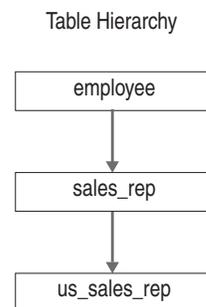


Figure 2-2. Example of Typed Table Hierarchy

When the following statement executes, the database server generates statistics on both the **sales_rep** and **us_sales_rep** tables:

```
UPDATE STATISTICS FOR TABLE sales_rep;
```

In contrast, the following example generates statistical data for each column in table **sales_rep** but does not act on tables **employee** or **us_sales_rep**:

```
UPDATE STATISTICS FOR TABLE ONLY (sales_rep);
```

If you specify FOR TABLE ONLY, as in this example, the identifier of the table (or *owner.table*) must be enclosed between parentheses.

Because neither of the previous examples specified the level at which to update the statistical data, the database server uses the LOW mode by default.

Updating Statistics for Columns

The Table and Column Scope specification can also include the names of one or more columns for which you want distributions calculated. For example, this statement calculates the distributions for three columns of the **orders** table:

```
UPDATE STATISTICS FOR TABLE orders (order_num, customer_num, ship_date);
```

If you include no *column* name in the FOR TABLE clause, then distributions are calculated for all columns of the specified *table*, using the LOW, MEDIUM, or HIGH mode and the number of bins implied by the specified or default Resolution clause *percentage* (for MEDIUM or HIGH mode) that you request.

Distributions are not calculated for BYTE or TEXT columns. See also “Updating Statistics for Columns of User-Defined Types (IDS)” for UPDATE STATISTICS restrictions on columns that store UDTs.

Examining Index Pages

In Dynamic Server, when you execute the UPDATE STATISTICS statement in any mode, the database server reads through index pages to:

- Compute statistics for the query optimizer
- Locate pages that have the delete flag marked as 1

If pages are found with the delete flag marked as 1, the corresponding keys are removed from the B-tree cleaner list.

This operation is particularly useful if a system failure causes the B-tree cleaner list (which exists in shared memory) to be lost. To remove the B-tree items that have been marked as deleted but are not yet removed from the B-tree, run the UPDATE STATISTICS statement. For information on the B-tree cleaner list, see your *IBM Informix Administrator's Guide*.

Updating Statistics for Columns of User-Defined Types (IDS)

To collect statistics for a column of a user-defined data type, you must specify either MEDIUM or HIGH mode. The UPDATE STATISTICS statement does not collect values for the **colmin** and **colmax** columns of the **syscolumns** system catalog table for columns that hold user-defined data types.

To drop statistics for a column that holds one of these data types, you must execute UPDATE STATISTICS in LOW mode with the DROP DISTRIBUTIONS option. When you use this option, the database server deletes the row in the **sysdistrib** system catalog table that corresponds to the **tableid** and **colno** values for the column. In addition, the database server removes any large objects that might have been created to store statistics for the specified opaque column.

Requirements for Statistics on Opaque Columns

UPDATE STATISTICS can collect statistics for columns of user-defined opaque data types only if support routines for **statcollect()**, **statprint()**, and the selectivity functions are defined for the UDTs. You must also hold Usage privilege on these routines.

In some cases, UPDATE STATISTICS also requires an sbspace as specified by the SYSSBSPACENAME configuration parameter. For information about how to provide statistical data for a column whose data type is a UDT, refer to the *IBM Informix Datablade API Programmer's Guide*. For information about SYSSBSPACENAME, refer to your *IBM Informix Administrator's Reference*.

Using the LOW Mode Option

Use the `LOW` option to generate and update some of the relevant statistical data regarding table, row, and page count statistics in the `systables` system catalog table. If you do not specify any mode, the `LOW` mode is the default.

In Dynamic Server, the `LOW` mode also generates and updates some index and column statistics for specified columns in the `syscolumns` and the `sysindexes` system catalog tables.

The `LOW` mode generates the least amount of information about the column. If you want the `UPDATE STATISTICS` statement to do minimal work, specify a column that is not part of an index. The `colmax` and `colmin` values in `syscolumns` are not updated unless there is an index on the column.

The following example updates statistics on the `customer_num` column of the `customer` table:

```
UPDATE STATISTICS LOW FOR TABLE customer (customer_num);
```

Because the `LOW` mode option does not update data in the `sysdistrib` system catalog table, all distributions associated with the `customer` table remain intact, even those that already exist on the `customer_num` column.

Using the DROP DISTRIBUTIONS Option

Use the `DROP DISTRIBUTIONS` option to force the removal of distribution information from the `sysdistrib` system catalog table.

When you specify the `DROP DISTRIBUTIONS` option, the database server removes the existing distribution data for the column or columns that you specify. If you do not specify any columns, the database server removes all the distribution data for that table.

You must have the `DBA` privilege or be owner of the table to use this option.

The following example shows how to remove distributions for the `customer_num` column in the `customer` table:

```
UPDATE STATISTICS LOW
    FOR TABLE customer (customer_num) DROP DISTRIBUTIONS;
```

As the example shows, you drop the distribution data at the same time you update the statistical data that the `LOW` mode option generates.

Using the DROP DISTRIBUTIONS ONLY Option

Use the `DROP DISTRIBUTIONS ONLY` option to remove distribution information from the `sysdistrib` table and update the `systables.version` column in the system catalog for those tables whose distributions were dropped, without gathering any `LOW` mode table and index statistics.

If you specify both the `DROP DISTRIBUTIONS ONLY` option and the `FOR TABLE` clause, Dynamic Server removes the existing distribution data for the set of columns of the `table` that the `FOR TABLE` clause specifies (or for all columns, if you provide no `column` specification), but does not gather any `LOW` mode table and index statistics.

You must have the `DBA` privilege or be owner of the table to use this option.

The following example removes distributions for the **customer_num** column in the **customer** table:

```
UPDATE STATISTICS LOW
  FOR TABLE customer (customer_num) DROP DISTRIBUTIONS ONLY;
```

This drops the **customer.customer_num** distribution data without updating the statistical information that the LOW mode option generates when the ONLY keyword does not follow the DROP DISTRIBUTIONS keywords. This example deletes from the system catalog any row describing **customer.customer_num** from the **sysdistrib** table, and updates the **version** number for **customer** in the **systables** table. None of the other LOW mode updates are performed on **systables**, so the **nrow** and **npused** column values are unchanged by this example, and the **syscolumns**, **sysfragments** and **sysindexes** tables of the system catalog are not updated. The LOW keyword has no effect in this example, but the DROP DISTRIBUTIONS ONLY option is not available in MEDIUM or HIGH mode.

Because it specifies no FOR TABLE clause, the next example drops all rows from the **sysdistrib** table and updates the **systables.version** column in the system catalog for all tables in the database.

```
UPDATE STATISTICS DROP DISTRIBUTIONS ONLY;
```

Using the MEDIUM Mode Option

Use the MEDIUM mode option to update the same statistics that you can perform with the LOW mode option and also generate statistics about the distribution of data values for each specified column. The database server places distribution information in the **sysdistrib** system catalog table.

If you use the MEDIUM mode option, the database server scans tables at least once and takes longer to execute on a given table than the LOW mode option.

When you use the MEDIUM mode option, the data distributions are obtained by sampling a percentage of data rows, using a statistical confidence level that you specify, or else a default confidence level of 95 percent. You can also specify an explicit minimum sampling size in the Resolution clause. Because the MEDIUM sample size is usually much smaller than the actual number of rows, this mode executes more quickly than the HIGH mode.

In distributions obtained by sampling, the results can vary, because different samples of rows can have different sampling errors. If the results vary significantly, you can use the Resolution clause to increase the sampling size, or to lower the *percent*, or to increase the *confidence* level to obtain more consistent results.

If the Resolution clause specifies no *percent* of sampled rows per bin, the default average percentage of the sample in each bin is 2.5, which divides the range into 40 intervals. If you do not specify a value for *confidence* level, the default level is 0.95. This value can be roughly interpreted to mean that 95 times out of 100, the difference between the MEDIUM estimate and the exact value from HIGH distributions is not statistically significant.

You must have the DBA privilege or be the owner of the table to create MEDIUM mode distributions. For more information on the MEDIUM and HIGH mode options, see the "Resolution Clause" on page 2-711.

Using the HIGH Mode Option

Use the HIGH mode option to update the same statistics that you can perform with the LOW mode option and also generate statistics about the distribution of data values for each specified column. The database server places distribution information in the **sysdistrib** system catalog table.

If you do not specify a Resolution clause, the default percentage of data distributed to every bin is 0.5, a value that partitions the range of values for each column into 200 intervals.

The constructed distribution is exact. Because more information is gathered, this mode executes more slowly than LOW or MEDIUM modes. If you use the HIGH mode option of UPDATE STATISTICS, the database server can take considerable time to gather the information across the database, particularly a database with large tables. The HIGH mode might scan each table several times for each column. To minimize processing time, specify a table name and column names within that table, rather than accept the default scope of all tables.

You must have the DBA privilege or be the owner of the table to create HIGH mode distributions. For more information on the MEDIUM and HIGH mode options, see the "Resolution Clause."

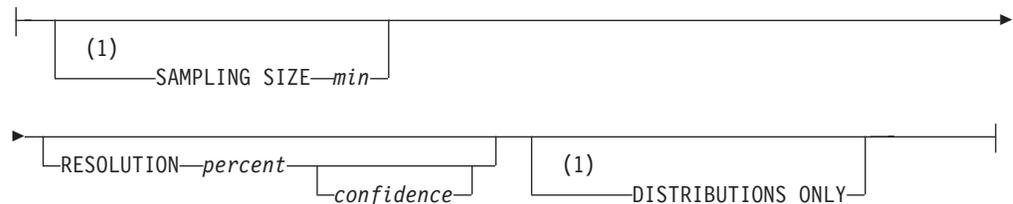
Resolution Clause

Use the Resolution clause in MEDIUM or HIGH mode to adjust the size of the distribution bins and to avoid calculating data on indexes. In MEDIUM mode only, you can also use the Resolution clause to specify a lower limit to the sampling size and to adjust the confidence level.

Resolution Clause:



Resolution Clause for MEDIUM Mode:



Resolution Clause for HIGH Mode:



Notes:

- 1 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>confidence</i>	Estimated likelihood that sampling in MEDIUM mode produces results equivalent to the exact HIGH mode. Default level is 0.95.	Must be within the range from 0.80 (minimum) to 0.99 (maximum)	"Literal Number " on page 4-185
<i>percent</i>	Average percentage of the sample in each distribution bin. Default is 2.5 for MEDIUM and 0.5 for HIGH.	Minimum value is 1/ <i>nrows</i> , for <i>nrows</i> the number of rows in the table	"Literal Number " on page 4-185
<i>min</i>	The minimum integer number of randomly selected rows on which to generate the data distributions	Must be greater than zero but cannot exceed <i>nrows</i>	"Literal Number " on page 4-185

A *distribution* is a mapping of the data in a column into a set of column values, ordered by magnitude or by collation. The range of these sample values is partitioned into disjunct intervals, called *bins*, each containing an approximately equal portion of the sample of column values. For example, if one bin holds 2 percent of the data, 50 such intervals hold the entire sample.

Some statistical texts call these bins *equivalence categories*. Each contains a disjunct subset of the range of the data values that are sampled from the column.

If you include the RESOLUTION keyword, it must be followed by a literal number, specifying the *percent* of values in each bin. In MEDIUM mode, it can be followed by either one or two literal numbers, with the optional second number specifying the *confidence* level, as in this example:

```
UPDATE STATISTICS MEDIUM FOR TABLE orders
RESOLUTION 4 0.90 DISTRIBUTIONS ONLY;
```

This specifies 4% of the data per bin, implying 25 bins, and a confidence level of 90%, and no examination of index data. If the 0.90 value were omitted, then the default level of confidence would have been in effect. If the RESOLUTION keyword and both numeric values were omitted, then default values for *percent* (2.5%) and for *confidence* (0.95) would be used.

The query optimizer estimates the selectivity of a WHERE clause by examining, for each column included in the WHERE clause, the proportional occurrence of the data values contained in the column.

You cannot create distributions for BYTE or TEXT columns. If you include a BYTE or TEXT column in an UPDATE STATISTICS statement that specifies MEDIUM or HIGH distributions, no distributions are created for those columns. Distributions are constructed for other columns in the list, however, and the statement does not return an error.

Columns of the VARCHAR data type do not use overflow bins, even when multiple bins are being used for duplicate values.

You can use the first two parameters of the DBUPSPACE environment variable to constrain the disk space and memory resources that the UPDATE STATISTICS statement can use to sort data when it constructs column distributions. These settings affect performance, because they determine how many times the database server scans the specified table to construct each distribution. (A third DBUPSPACE parameter can control whether UPDATE STATISTICS sorts with

indexes when calculating column distributions, and whether the `sqexplain.out` file stores the plan by which the column distributions are calculated.)

Specifying the SAMPLING SIZE (IDS)

In MEDIUM mode, you can optionally use the SAMPLING SIZE keywords to specify the minimum number of rows to sample for calculating column distribution statistics. If the Resolution clause omits the RESOLUTION keyword and specifies no *confidence* level and no *percent* value, then the number of rows that Dynamic Server samples will be the larger of the following two values:

- The *min* value that you specify immediately after the SAMPLING SIZE keywords
- The sampling size that is required for the default *percent* of rows in each bin (2.5%) and for the minimum *confidence* level (0.80).

If a sampling size is specified in a Resolution clause that includes explicit values for both the average *percent* of sampled rows per bin and for the *confidence* level, then the number of sampled rows will be the larger of these two values:

- The *min* value that you specify immediately after the SAMPLING SIZE keywords
- The sampling size that is required for the specified *percent* of rows and for the specified *confidence* level.

If a sampling size is specified in a Resolution clause that includes an average *percentage* value but sets no *confidence* level, then the minimum *confidence* value of 0.80 is used to calculate the actual sampling size for Dynamic Server to use if the specified *size* is smaller.

For example, the following statement calculates statistics for two columns of the **customer** table, without updating index information. At least 200 rows will be sampled, but the actual size of the sample might be larger than 200 if more rows are required to provide the default 0.80 confidence level for a sample distribution that uses 50 equivalence categories, with an average percentage of 2% of the sampled values in each bin.

```
UPDATE STATISTICS MEDIUM FOR TABLE customer (city, state)
  SAMPLING SIZE 200 RESOLUTION 2 DISTRIBUTIONS ONLY;
```

Whether or not you include an explicit SAMPLING SIZE specification in the Resolution clause, Dynamic Server records in the system catalog the actual sampling size (as a percentage of the total number of rows in the table) at the time of MEDIUM mode UPDATE STATISTICS creation.

Using the DISTRIBUTIONS ONLY Option to Suppress Index Information

In Dynamic Server, when you specify the DISTRIBUTIONS ONLY option, you do not update index information. This option does not affect existing index information.

Use this option to avoid the examination of index information that can consume considerable processing time.

This option does not affect the recalculation of information on tables, such as the number of pages used, the number of rows, and fragment information. UPDATE STATISTICS needs this information to construct accurate column distributions and requires little time and system resources to collect it.

Do not confuse this DISTRIBUTIONS ONLY option with the DROP DISTRIBUTIONS ONLY option of LOW mode, whose syntax and semantics are not supported in MEDIUM or HIGH mode. For information on how to suppress the collection of column distributions, see “Using the DROP DISTRIBUTIONS ONLY Option” on page 2-709.

Using DBUPSPACE Settings to Suppress Index Information

You can also prevent indexes from being used by UPDATE STATISTICS operations in sorting rows by setting the third parameter of the **DBUPSPACE** environment variable to a value of 1. Refer to Chapter 3 of the *IBM Informix Guide to SQL: Reference* for information about the **DBSPACETEMP** and **DBUPSPACE** environment variables, which can restrict the system resources that are available for UPDATE STATISTICS operations. (The database server uses the storage locations that **DBSPACETEMP** specifies only when you use the HIGH option of UPDATE STATISTICS.)

Output for UPDATE STATISTICS from the SET EXPLAIN Statement

The SET EXPLAIN statement can display the plan that UPDATE STATISTICS uses to generate column distributions. The following output is based on the default **DBUPSPACE** value of 15 megabytes of sort memory, which in this example requires two passes to sort the 21.9 megabytes of data:

```
UPDATE STATISTICS:
=====

Table: zelaine.t1
Mode: HIGH
Number of Bins: 267 Bin size 2505
Sort data 21.9 MB Sort memory granted 15.0 MB
Estimated number of table scans 2
PASS #1 b
PASS #2 a
Scan 9 Sort 1 Build 2 Insert 0 Close 0 Total 12
Completed pass 1 in 0 minutes 12 seconds
Scan 5 Sort 2 Build 1 Insert 0 Close 0 Total 8
Completed pass 2 in 0 minutes 8 seconds
```

Routine Statistics

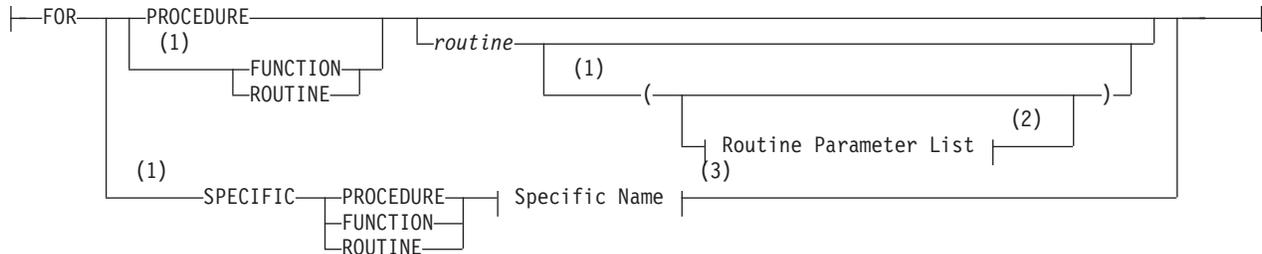
Before the database server executes a new SPL routine the first time, it optimizes the statements in the SPL routine. Optimization makes the code depend on the structure of tables referenced by the routine. If a DDL operation modifies the schema of a referenced table after the routine is optimized, but before it is executed, the routine can fail with an error.

This failure typically does not occur, however, if an index is added or dropped while automatic recompilation is enabled for routines referencing tables that ALTER TABLE, CREATE INDEX, or DROP INDEX operations have modified. This is the default behavior of Dynamic Server. For more information about enabling or disabling automatic reoptimization after changes to the schema of a table, see the description of the IFX_AUTO_REPREPARE option to the SET ENVIRONMENT statement.

When the AUTO_REPREPARE configuration parameter and the IFX_AUTO_REPREPARE session environment variable are set to disable recompilation of SPL routines that reference tables whose schema has been modified, however, adding or dropping an index to a table that an SPL routine references indirectly can cause the routine to return error -710. To avoid this error after DDL operations, or to reoptimize SPL routines after table distributions might

have been modified by DML operations, use the Routine Statistics segment of UPDATE STATISTICS to update the execution plans of any SPL routines that reference the table.

Routine Statistics:



Notes:

- 1 Dynamic Server only
- 2 See "Routine Parameter List" on page 5-67
- 3 See "Specific Name " on page 5-73

Element	Description	Restrictions	Syntax
<i>routine</i>	Name declared for a SPL routine in a CREATE FUNCTION or CREATE PROCEDURE statement	Must reside in the database. In ANSI-compliant databases, qualify <i>routine</i> with <i>owner</i> if you are not <i>owner</i> .	"Identifier " on page 5-21

The following table explains the keywords of the Routine Statistics segment.

Keyword

Which Execution Plan is Reoptimized

SPECIFIC

The plan for the SPL routine called *specific name*

FUNCTION

The plan for any SPL function with the specified name (and with parameter types that match *routine parameter list*, if supplied). If you specify the FUNCTION keyword, the UPDATE STATISTICS statement fails with an error unless the specified routine returns a value or values, with or without the WITH RESUME option.

PROCEDURE

The plan for any SPL procedure with the specified name (and parameter types that match *routine parameter list*, if supplied)

ROUTINE

The plan for SPL functions and procedures with the specified name (and parameter types that match *routine parameter list*, if supplied)

The parentheses symbols are optional if you include no argument list.

If you specify no *routine*, the execution plans are reoptimized for all SPL routines in the current database.

The database server keeps a list of tables that the SPL routine references explicitly. Whenever an explicitly referenced table is modified, the database server reoptimizes the procedure the next time the procedure is executed.

The **sysprocplan** system catalog table stores execution plans for SPL routines. Two actions can update the **sysprocplan** system catalog table:

- Execution of an SPL routine that uses a modified table
- The UPDATE STATISTICS statement

If you change a table that an SPL routine references, you can run UPDATE STATISTICS to reoptimize the procedures that reference the table, rather than waiting until the next time an SPL routine that uses the table executes. (If a table that an SPL routine references is dropped, however, running UPDATE STATISTICS cannot prevent the SPL routine from failing with an error.)

Altered Tables that are Referenced Indirectly in SPL Routines

If the SPL routine depends on a table that is referenced only indirectly, however, the database server cannot detect the need to reoptimize the procedure after that table is modified. For example, a table can be referenced indirectly if the SPL routine invokes a trigger. If the schema of a table that is referenced by the trigger (but not directly by the SPL routine) is changed, the database server does not know that it should reoptimize the SPL routine before running it. When the procedure is run after the table has been changed, error -710 can occur.

Each SPL routine is optimized the first time that it is run (not when it is created). This behavior means that an SPL routine might succeed the first time it is run but fail later under virtually identical circumstances, if the schema of an indirectly referenced table has been changed. The failure of an SPL routine can also be intermittent, because failure during one execution forces an internal warning to reoptimize the procedure before the next execution.

You can use either of two methods to recover from this error:

- Issue UPDATE STATISTICS to force reoptimization of the routine.
- Rerun the routine.

To prevent this error, you can force reoptimization of the SPL routine. To force reoptimization, execute the following statement:

```
UPDATE STATISTICS FOR PROCEDURE routine;
```

You can add this statement to your program in either of the following ways:

- Issue UPDATE STATISTICS after each statement that changes the mode of an object.
- Issue UPDATE STATISTICS before each invocation of the SPL routine.

For efficiency, you can put the UPDATE STATISTICS statement with the action that occurs less frequently in the program (change of object mode or execution of the procedure). In most cases, the action that occurs less frequently in the program is the change of object mode.

When you follow this method of recovering from this error, you must execute UPDATE STATISTICS for each procedure that indirectly references the altered tables unless the procedure also references the tables explicitly.

You can also recover from error -710 after an indirectly referenced table is altered simply by re-executing the SPL routine. The first time that the stored procedure fails, the database server marks the procedure as in need of reoptimization. The next time that you run the procedure, the database server reoptimizes the procedure before running it. Running the SPL routine twice, however, might be

neither practical nor safe. A safer choice is to use the UPDATE STATISTICS statement to force reoptimization of the procedure.

Updating Statistics When You Upgrade the Database Server (IDS)

When you upgrade a database to use with a newer database server, you can use the UPDATE STATISTICS statement to convert the indexes to the form that the newer database server uses. You can choose to convert the indexes one table at a time or for the entire database at one time. Follow the conversion guidelines that are outlined in the *IBM Informix Migration Guide*.

When you use the UPDATE STATISTICS statement to convert the indexes to use with a newer database server, the indexes are implicitly dropped and re-created. The only time that an UPDATE STATISTICS statement causes table indexes to be implicitly dropped and re-created is when you upgrade a database for use with a newer database server.

Performance

The more specific you make the list of objects that UPDATE STATISTICS examines, the faster it completes execution. Limiting the number of columns distributed speeds the update. Similarly, precision affects the speed of the update. If all other keywords are the same, LOW works fastest, but HIGH examines the most data.

Examples

```
UPDATE STATISTICS MEDIUM;  
UPDATE STATISTICS MEDIUM RESOLUTION 10;  
UPDATE STATISTICS MEDIUM RESOLUTION 10 .95;  
{ RESOLUTION 10, CONFIDENCE .95}  
UPDATE STATISTICS MEDIUM RESOLUTION 10 DISTRIBUTIONS ONLY;  
UPDATE STATISTICS MEDIUM RESOLUTION 10 .95 DISTRIBUTIONS ONLY;
```

```
UPDATE STATISTICS HIGH;  
UPDATE STATISTICS HIGH RESOLUTION 10;  
UPDATE STATISTICS HIGH RESOLUTION 10 DISTRIBUTIONS ONLY;
```

Resolution must be greater than 0.005 and less than or equal to 10.0. Confidence must be in the range [0.80, 0.99] (inclusive).

Examples that follow are based on the **company_proc** procedure and **square_w_default** function, as defined below:

```
CREATE PROCEDURE company_proc ( no_of_items INT,  
                               itm_quantity SMALLINT, sale_amount MONEY,  
                               customer VARCHAR(50), sales_person VARCHAR(30) )  
SPECIFIC spec_cmpy  
  
DEFINE salesperson_proc VARCHAR(60);  
  
-- Update the company table  
INSERT INTO company_tbl VALUES (no_of_items, itm_quantity,  
                                sale_amount, customer, sales_person);  
  
-- Generate the procedure name for the variable salesperson_proc  
LET salesperson_proc = sales_person || "." || "tbl" ||  
                      month(current) || "_" || year(current) || "_proc" ;  
  
-- Execute the SPL procedure that the salesperson_proc  
-- variable specifies  
EXECUTE PROCEDURE salesperson_proc (no_of_items,  
                                    itm_quantity, sale_amount, customer);
```

```

END PROCEDURE;

CREATE FUNCTION square_w_default
  (i INT DEFAULT 0) {Specifies default value of i}
RETURNING INT {Specifies return of INT value}
SPECIFIC spec_square
  DEFINE j INT; {Defines routine variable j}
  LET j = i * i; {Finds square of i and assigns it to j}
  RETURN j; {Returns value of j to calling module}
END FUNCTION;

```

The UPDATE STATISTICS examples that follow reference the **company_proc** procedure and **square_w_default** function:

```

UPDATE STATISTICS FOR PROCEDURE;
UPDATE STATISTICS FOR PROCEDURE company_proc1;
UPDATE STATISTICS FOR PROCEDURE
  company_proc1(INT,SMALLINT,MONEY,VARCHAR(50), VARCHAR(30));
UPDATE STATISTICS FOR SPECIFIC PROCEDURE spec_cmpy;

UPDATE STATISTICS FOR FUNCTION;
UPDATE STATISTICS FOR FUNCTION square_w_default;
UPDATE STATISTICS FOR FUNCTION square_w_default(INT);
UPDATE STATISTICS FOR SPECIFIC FUNCTION spec_square;

```

Related Statements

Related statements: “SET EXPLAIN ” on page 2-623 and “SET OPTIMIZATION ” on page 2-641

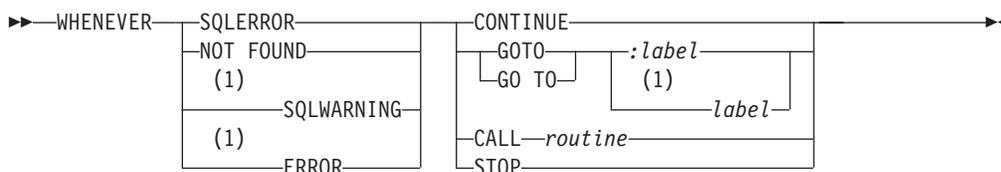
For a discussion of the performance implications of UPDATE STATISTICS, see your *IBM Informix Performance Guide*.

For a discussion of how to use the **dbschema** utility to view distributions created with UPDATE STATISTICS, see the *IBM Informix Migration Guide*.

WHENEVER

Use the WHENEVER statement to trap exceptions that occur during the execution of SQL statements. Use this statement only with Informix ESQ/C.

Syntax



Notes:

- 1 Informix extension

Element	Description	Restrictions	Syntax
<i>label</i>	Statement label to which program control transfers when an exception occurs	Must exist in the same source-code module.	Language-specific

Element	Description	Restrictions	Syntax
<i>routine</i>	Name of a user-defined routine (UDR) to be invoked when an exception occurs	No arguments; UDR must exist at compile time.	"Identifier " on page 5-21

Usage

The WHENEVER statement is equivalent to placing an exception-checking routine after every SQL statement. The following table summarizes the types of exceptions for which you can check with the WHENEVER statement.

Type of Exception	WHENEVER Keyword	For More Information
Errors	SQLERROR <i>or</i> ERROR	"SQLERROR Keyword" on page 2-720
Warnings	"SQLWARNING Keyword" on page 2-720	
Not Found <i>or</i> End of Data	"NOT FOUND Keywords" on page 2-721	

Programs that do not use the WHENEVER statement do not automatically abort when an exception occurs. Such programs must explicitly check for exceptions and take whatever corrective action their logic specifies. If you do not check for exceptions, the program simply continues running. If errors occur, however, the program might not perform its intended purpose.

The first keyword that follows WHENEVER specifies some type of exceptional condition; the last part of the statement specifies some action to take when the exception is encountered (or no action, if CONTINUE is specified). The following table summarizes possible actions that WHENEVER can specify.

Type of Action	WHENEVER Keyword	For More Information
Continue program execution	"CONTINUE Keyword" on page 2-721	
Stop program execution	"STOP Keyword" on page 2-721	
Transfer control to a specified label	GOTO GO TO	"GOTO Keyword" on page 2-721
Transfer control to a UDR	"CALL Clause" on page 2-722	

The Scope of WHENEVER

WHENEVER is a preprocessor directive, rather than an executable statement. The Informix ESQL/C preprocessor, not the database server, handles the interpretation of the WHENEVER statement. When the preprocessor encounters a WHENEVER statement in an Informix ESQL/C source file, it inserts appropriate code into the preprocessed code after each SQL statement, based on the exception and the action that WHENEVER specifies. The scope of the WHENEVER statement begins where the statement appears in the source module and remains in effect until the preprocessor encounters one or the other of the following things while sequentially processing the source module:

- The next WHENEVER statement with the same condition (SQLERROR, SQLWARNING, or NOT FOUND) in the same source module
- The end of the source module

The following Informix ESQL/C example program has three WHENEVER statements, two of which are WHENEVER SQLERROR statements. Line 4 uses STOP with SQLERROR to override the default CONTINUE action for errors.

Line 8 specifies the CONTINUE keyword to return the handling of errors to the default behavior. For all SQL statements between lines 4 and 8, the preprocessor inserts code that checks for errors and halts program execution if an error occurs. Therefore, any errors that the INSERT statement on line 6 generates cause the program to stop.

After line 8, the preprocessor does not insert code to check for errors after SQL statements. Therefore, any errors that the INSERT statement (line 10), the SELECT statement (line 11), and DISCONNECT statement (line 12) generate are ignored. The SELECT statement, however, does not stop program execution if it does not locate any rows; the WHENEVER statement on line 7 tells the program to continue if such an exception occurs:

```

1  main()
2  {
3  EXEC SQL connect to 'test';
4  EXEC SQL WHENEVER SQLERROR STOP;
5  printf("\n\nGoing to try first insert\n\n");
6  EXEC SQL insert into test_color values ('green');
7  EXEC SQL WHENEVER NOT FOUND CONTINUE;
8  EXEC SQL WHENEVER SQLERROR CONTINUE;
9  printf("\n\nGoing to try second insert\n\n");
10 EXEC SQL insert into test_color values ('blue');
11 EXEC SQL select paint_type from paint where color='red';
12 EXEC SQL disconnect all;
13 printf("\n\nProgram over\n\n");
14 }

```

SQLERROR Keyword

If you use the SQLERROR keyword, any SQL statement that encounters an error is handled as the WHENEVER SQLERROR statement directs. If an error occurs, the `sqlcode` variable (`sqlca.sqlcode`, `SQLCODE`) is set to a value less than zero (0) and the `SQLSTATE` variable is set to a class code with a value greater than 02.

The next example terminates program execution if an SQL error is detected:

```
WHENEVER SQLERROR STOP
```

If you do not include any WHENEVER SQLERROR statements in a program, the default action for WHENEVER SQLERROR is CONTINUE.

ERROR Keyword

Within the WHENEVER statement (and only in this context), the keyword ERROR is a synonym for the SQLERROR keyword.

SQLWARNING Keyword

If you use the SQLWARNING keyword, any SQL statement that generates a warning is handled as the WHENEVER SQLWARNING statement directs. If a

warning occurs, the first field (**sqlca.sqlwarn.sqlwarn0**) of the warning structure in SQLCA is set to W, and the SQLSTATE variable is set to a class code of 01.

Besides the first field of the warning structure, a warning also sets an additional field to W. The field that is set indicates what type of warning occurred.

The next statement causes execution to stop if a warning condition exists:

```
WHENEVER SQLWARNING STOP
```

If you do not use any WHENEVER SQLWARNING statements in a program, the default action for WHENEVER SQLWARNING is CONTINUE.

NOT FOUND Keywords

If you use the NOT FOUND keywords, exception handling for SELECT and FETCH statements (including implicit SELECT and FETCH statements in FOREACH and UNLOAD statements) is treated differently from other SQL statements. The NOT FOUND keyword checks for the following cases:

- The **End of Data** condition: a FETCH statement that attempts to get a row beyond the first or last row in the active set
- The **Not Found** condition: a SELECT statement that returns no rows

In each case, the **sqlcode** variable is set to 100, and the SQLSTATE variable has a class code of 02. For the name of the **sqlcode** variable in each IBM Informix product, see the table in “SQLERROR Keyword” on page 2-720.

The following statement calls the **no_rows()** function each time the NOT FOUND condition exists:

```
WHENEVER NOT FOUND CALL no_rows
```

If you do not use any WHENEVER NOT FOUND statements in a program, the default action for WHENEVER NOT FOUND is CONTINUE.

CONTINUE Keyword

Use the CONTINUE keyword to instruct the program to ignore the exception and to continue execution at the next statement after the SQL statement. The default action for all exceptions is CONTINUE. You can use this keyword to turn off a previously specified action for an exceptional condition.

STOP Keyword

Use the STOP keyword to instruct the program to stop execution when the specified exception occurs. The following statement halts execution of an Informix ESQL/C program each time that an SQL statement generates a warning:

```
EXEC SQL WHENEVER SQLWARNING STOP;
```

GOTO Keyword

Use the GOTO clause to transfer control to the statement that the label identifies when a specified exception occurs. The GOTO and GO TO keywords are ANSI-compliant syntax for this feature of embedded SQL languages like ESQL/C.

The following Informix ESQL/C code fragment shows a WHENEVER statement that transfers control to the label **missing** each time that the NOT FOUND condition occurs:

```
query_data()
...
EXEC SQL WHENEVER NOT FOUND GO TO missing;
...
EXEC SQL fetch lname into :lname;
...
missing:
    printf("No Customers Found\n");
```

Within the scope of the WHENEVER GOTO statement, you must define the labeled statement in *each* routine that contains SQL statements. If your program contains more than one user-defined function, you might need to include the labeled statement and its code in *each* function.

If the preprocessor encounters an SQL statement within the scope of a WHENEVER ... GOTO statement, but within a routine that does not have the specified label, the preprocessor tries to insert the code associated with the labeled statement, but generates an error when it cannot find the label.

To correct this error, either put a labeled statement with the same label name in each UDR, or issue another WHENEVER statement to reset the error condition, or use the CALL clause to call a separate function.

CALL Clause

Use the CALL clause to transfer program control to the specified UDR when the specified type of exception occurs. Do not include parentheses after the UDR name. The following WHENEVER statement causes the program to call the **error_recovery()** function if the program detects an error:

```
EXEC SQL WHENEVER SQLERROR CALL error_recovery;
```

When the UDR returns, execution resumes at the next statement after the line that is causing the error. If you want to halt execution when an error occurs, include statements that terminate the program as part of the specified UDR.

Observe the following restrictions on the specified routine:

- The UDR cannot accept arguments nor can it return values. If it needs external information, use global variables or the WHENEVER ... GOTO option to transfer program control to a label that calls the UDR.
- You cannot specify the name of an SPL routine in the CALL clause. To call an SPL routine, use the CALL clause to invoke a UDR that contains the EXECUTE FUNCTION (or EXECUTE PROCEDURE) statement.
- Make sure that all functions within the scope of WHENEVER ... CALL statements can find a declaration of the specified function.

Related Statements

Related statements: "EXECUTE FUNCTION " on page 2-356, "EXECUTE PROCEDURE " on page 2-364, and "FETCH " on page 2-367

For discussions on exception handling and error checking, see the *IBM Informix ESQL/C Programmer's Manual*.

Chapter 3. SPL Statements

In This Chapter

You can use Stored Procedure Language (SPL) statements to write SPL routines (formerly referred to as *stored procedures*), and you can store these routines in the database as user-defined routines (UDRs). SPL routines are effective tools for controlling SQL activity. This chapter contains descriptions of the SPL statements. The description of each statement includes the following information:

- A brief introduction that explains the effect of the statement
- A syntax diagram that shows how to enter the statement correctly
- A syntax table that explains each input parameter in the syntax diagram
- Rules of usage, including examples that illustrate these rules

If a statement is composed of multiple clauses, the statement description provides information about each clause.

For an overview of the SPL language and task-oriented information about creating and using SPL routines, see the *IBM Informix Guide to SQL: Tutorial*.

In Extended Parallel Server, only the CREATE PROCEDURE statement or the CREATE PROCEDURE FROM statements can create SPL routines. Extended Parallel Server does not support the CREATE FUNCTION or the CREATE FUNCTION FROM statement.

Dynamic Server can create an SPL function with the CREATE PROCEDURE or CREATE PROCEDURE FROM statement, but requires the CREATE FUNCTION or CREATE FUNCTION FROM statement for external functions. It is recommended, however, that you use the CREATE FUNCTION or CREATE FUNCTION FROM statement to create new user-defined functions.

<< Label >>

Use the <<*label*>> statement of SPL to declare a statement label or a loop label.

- A *statement label* is an SQL identifier, delimited by double angle-brackets, immediately preceding a statement within a statement block to which the GOTO statement of SPL can transfer control of program execution.
- A *loop label* is an SQL identifier, delimited by double angle-brackets, immediately preceding a loop statement of SPL. The same label, without double angle-bracket delimiters, can follow the END LOOP keywords, or END FOR keywords, or END WHILE keywords that terminate the labeled loop statement. The EXIT *label* statement can pass control of execution to whatever statement follows the uncommitted statement label.

Note that *label* is not a keyword of the <<*label*>> statement, but is a placeholder for some specific user-defined identifier of the statement label or loop label that the <<*label*>> statement declares.

Only Dynamic Server supports the <<*label*>> statement.

Syntax

►<<—*label*—>>◄

Element	Description	Restrictions	Syntax
<i>label</i>	Name that you declare here for a statement label or for a loop label	Must be unique among the identifiers of statement labels and of loop labels within the SPL routine	“Identifier ” on page 5-21

Usage

You can use the <<*Label*>> statement in two ways:

- To declare a *statement label* before an executable statement to which the GOTO statement of SPL can transfer control of execution. The SPL statement that immediately follows the statement label declaration is called a *labeled statement*.
- To declare a *loop label* immediately before a LOOP, FOR, or WHILE statement of SPL. The LOOP, FOR, or WHILE statement that immediately follows the loop label declaration is called a *labeled loop*.

The EXIT *label* or EXIT *label* WHEN (*condition*) statement can exit from the labeled loop, passing control of execution to the statement immediately following an END LOOP *label* statement. The *label* specified in the EXIT statement can match the label identifier of the labeled loop of the EXIT statement, or if loops are nested, this *label* can match the label of an outer labeled loop. In either case, the EXIT *label* statement passes control to a statement that follows an END LOOP *label* statement that specifies the same loop label. This EXIT *label* behavior differs from that of the GOTO *label* statement, which passes control to the statement that follows the declaration of the specified statement label.

The following restrictions apply to labels in SPL routines:

- The name of a statement label must be within the scope of reference of the GOTO statement.
- The GOTO option of the WHENEVER statement of SQL cannot reference an SPL statement label, because the WHENEVER statement is valid only in ESQL/C applications.
- The GOTO statement of SPL cannot reference a loop label.
- The GOTO statement cannot reference a statement label within an ON EXCEPTION statement block.
- A statement label cannot be declared within an ON EXCEPTION statement block.
- The label name must be unique among statement labels and loop labels within the SPL routine.

Examples of Labels

The following example illustrates a statement label called **increment_x** within an SPL routine:

```
DEFINE x INT;  
LET x = 0;  
BEGIN
```

```

    <<increment_x>>
    BEGIN
        LET x = x + 1;
    END;
    IF x < 10 THEN
        GOTO increment_x;
    END IF;
END;
END PROCEDURE;

```

The following program fragment shows an example of a labeled FOR loop:

```

<<lb_for>>
FOR i IN 1..5
    i := i + 1 ;
END FOR lb_for;

```

The following program fragment illustrates a labeled loop from which an EXIT *label* statement can exit:

```

<<outer>>
LOOP
...
LOOP
...
EXIT outer WHEN ... -- exit from both loops
END LOOP;
...
END LOOP outer;

```

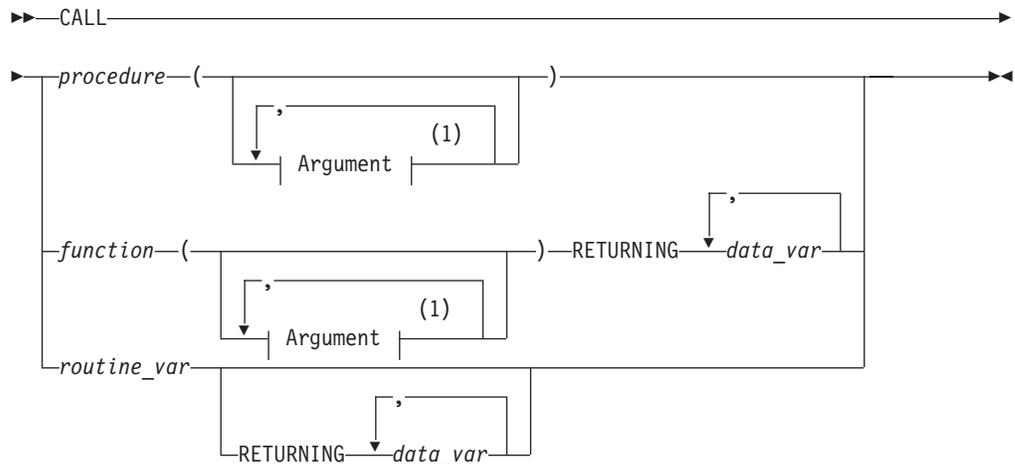
Related Statements

“EXIT ” on page 3-17, “FOR ” on page 3-19, “GOTO” on page 3-27, “LOOP ” on page 3-34, “WHILE ” on page 3-49

CALL

Use the CALL statement to execute a user-defined routine (UDR) from within an SPL routine.

Syntax



Notes:

- 1 See "Arguments " on page 5-1

Element	Description	Restrictions	Syntax
<i>data_var</i>	Variable to receive the values <i>function</i> returns	The data type of <i>data_var</i> must be appropriate for the returned value	"Identifier " on page 5-21
<i>function, procedure</i>	User-defined function or procedure	The function or procedure must exist	"Identifier " on page 5-21
<i>routine_var</i>	Variable that contains the name of a UDR	Must be a character data type that contains the non-NULL name of an existing UDR	"Identifier " on page 5-21

Usage

The CALL statement invokes a UDR. The CALL statement is identical in behavior to the EXECUTE PROCEDURE and EXECUTE FUNCTION statements, but you can only use CALL from within an SPL routine.

You can use CALL in an Informix ESQL/C program or with DB–Access, but only if the statement is in an SPL routine that the program or DB–Access executed.

If you CALL a user-defined *function*, you must specify a RETURNING clause.

The CALL statement cannot invoke a TABLE function within a subquery in the FROM clause of the SELECT statement.

Receiving Input from the Called UDR

The RETURNING clause specifies the variable that receives values that a called function returns.

The following example shows two UDR calls:

```
CREATE PROCEDURE not_much()  
  DEFINE i, j, k INT;  
  CALL no_args (10,20);  
  CALL yes_args (5) RETURNING i, j, k;  
END PROCEDURE;
```

The first routine call (**no_args**) expects no returned values. The second routine call is to a function (**yes_args**), which expects three returned values. The **not_much()** procedure declares three integer variables (**i**, **j**, and **k**) to receive the returned values from **yes_args**.

Receiving Input from the Called UDR

The RETURNING clause specifies the variable that receives values that a called function returns.

The following example shows two UDR calls:

```
CREATE PROCEDURE not_much()  
  DEFINE i, j, k INT;  
  CALL no_args (10,20);  
  CALL yes_args (5) RETURNING i, j, k;  
END PROCEDURE;
```

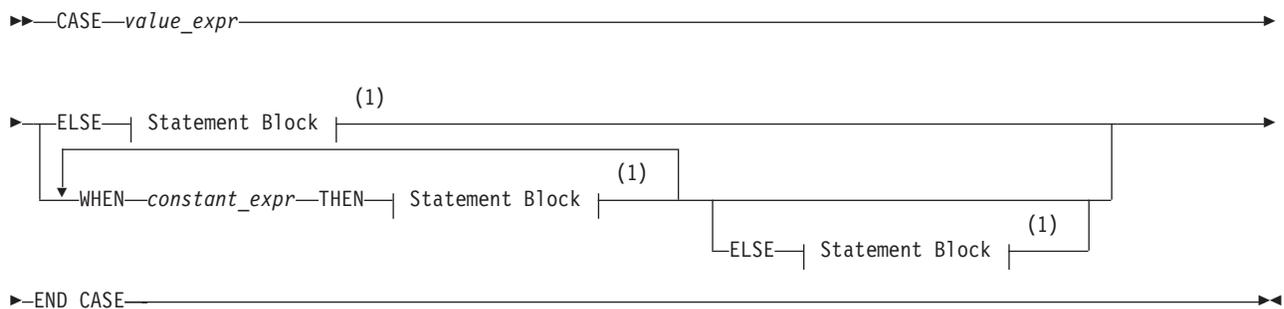
The first routine call (**no_args**) expects no returned values. The second routine call is to a function (**yes_args**), which expects three returned values. The **not_much()** procedure declares three integer variables (**i**, **j**, and **k**) to receive the returned values from **yes_args**.

CASE

Use the CASE statement when you need to take one of many branches depending on the value of an SPL variable or a simple expression. The CASE statement is a fast alternative to the IF statement.

Only Extended Parallel Server supports the CASE statement.

Syntax



Notes:

1 See "Statement Block " on page 5-74

Element	Description	Restrictions	Syntax
<i>constant_expr</i>	Expression that specifies a literal value	Can be a literal number, quoted string, literal datetime, or literal interval. The data type must be compatible with the data type of <i>value_expr</i> .	"Constant Expressions" on page 4-70
<i>value_expr</i>	Expression that returns a value	An SPL variable or any other type of expression that returns a value	"Expression" on page 4-40

Usage

You can use the CASE statement to create a set of conditional branches within an SPL routine. The WHEN and the ELSE clauses are optional, but you must include at least one of them. If you specify no WHEN and no ELSE clause, you receive a syntax error.

Do not confuse the CASE statement with CASE expressions of SQL. (For Extended Parallel Server, CASE expressions support the same keywords as the CASE statement, but use different syntax and semantics to evaluate *conditions* that you specify and return a single value or NULL, as described in "CASE Expressions" on page 4-66.)

How the Database Server Executes a CASE Statement

The database server executes the CASE statement in the following way:

- The database server evaluates the *value_expr* parameter.

- If the resulting value matches a literal value specified in the *constant_expr* parameter of a WHEN clause, the database server executes the statement block that follows the THEN keyword in that WHEN clause.
- If the value resulting from the evaluation of the *value_expr* parameter matches the *constant_expr* parameter in more than one WHEN clause, the database server executes the statement block that follows the THEN keyword in the first matching WHEN clause in the CASE statement.
- After the database server executes the statement block that follows the THEN keyword, it executes the statement that follows the CASE statement in the SPL routine.
- If the value of the *value_expr* parameter does not match the literal value specified in the *constant_expr* parameter of any WHEN clause, and if the CASE statement includes an ELSE clause, the database server executes the statement block that follows the ELSE keyword.
- If the value of the *value_expr* parameter does not match the literal value specified in the *constant_expr* parameter of any WHEN clause, and if the CASE statement does not include an ELSE clause, the database server executes the statement that follows the CASE statement in the SPL routine.
- If the CASE statement includes an ELSE clause but not a WHEN clause, the database server executes the statement block that follows the ELSE keyword.

The statement block that follows the THEN or ELSE keywords can include any SQL statement or SPL statement that is valid in the statement block of an SPL routine. For more information, see “Statement Block ” on page 5-74.

Computation of the Value Expression in CASE

The database server computes the value of the *value_expr* parameter only one time. It computes this value at the start of execution of the CASE statement. If the value expression specified in the *value_expr* parameter contains SPL variables and the values of these variables change subsequently in one of the statement blocks within the CASE statement, the database server does not recompute the value of the *value_expr* parameter. So a change in the value of any variables contained in the *value_expr* parameter has no effect on the branch taken by the CASE statement.

Example of CASE Statement

In the following example, the CASE statement initializes one of a set of SPL variables (named *j*, *k*, *l*, and *m*) to the value of an SPL variable named *x*, depending on the value of another SPL variable named *i*:

```

CASE i
  WHEN 1 THEN LET j = x;
  WHEN 2 THEN LET k = x;
  WHEN 3 THEN LET l = x;
  WHEN 4 THEN LET m = x;
  ELSE
    RAISE EXCEPTION 100; --invalid value
END CASE;

```

Related Statements

“IF ” on page 3-28

CONTINUE

Use the CONTINUE statement to start the next iteration of the innermost FOR, LOOP, WHILE, or FOREACH loop.

Syntax



Usage

When control of execution passes to a CONTINUE statement, the SPL routine skips the rest of the statements in the innermost loop of the specified type. Execution continues at the top of the loop with the next iteration.

In the following example, the `loop_skip` function inserts values 3 through 15 into the table `testtable`. The function also returns values 3 through 9 and 13 through 15 in the process. The function does not return the value 11 because it encounters the CONTINUE FOR statement. The CONTINUE FOR statement causes the function to skip the RETURN WITH RESUME statement:

```
CREATE FUNCTION loop_skip()
  RETURNING INT;
  DEFINE i INT;
  ...
  FOR i IN (3 TO 15 STEP 2)
    INSERT INTO testtable values(i, null, null);
    IF i = 11
      CONTINUE FOR;
    END IF;
    RETURN i WITH RESUME;
  END FOR;

END FUNCTION;
```

Just as with the EXIT statement, (“EXIT ” on page 3-17), in FOREACH statements and in FOR or WHILE statements that do not include the LOOP keyword, the FOR, WHILE, or FOREACH keyword must immediately follow the CONTINUE keyword to specify the type of loop. Errors are generated if the specified type of loop does not match the context in which the CONTINUE statement is issued.

In the LOOP, FOR LOOP, and WHILE LOOP statements, whether labeled or unlabeled, a keyword indicating the type of loop is optional after the CONTINUE keyword, but Dynamic Server issues an error if you specify a keyword that does not correspond to the type of loop.

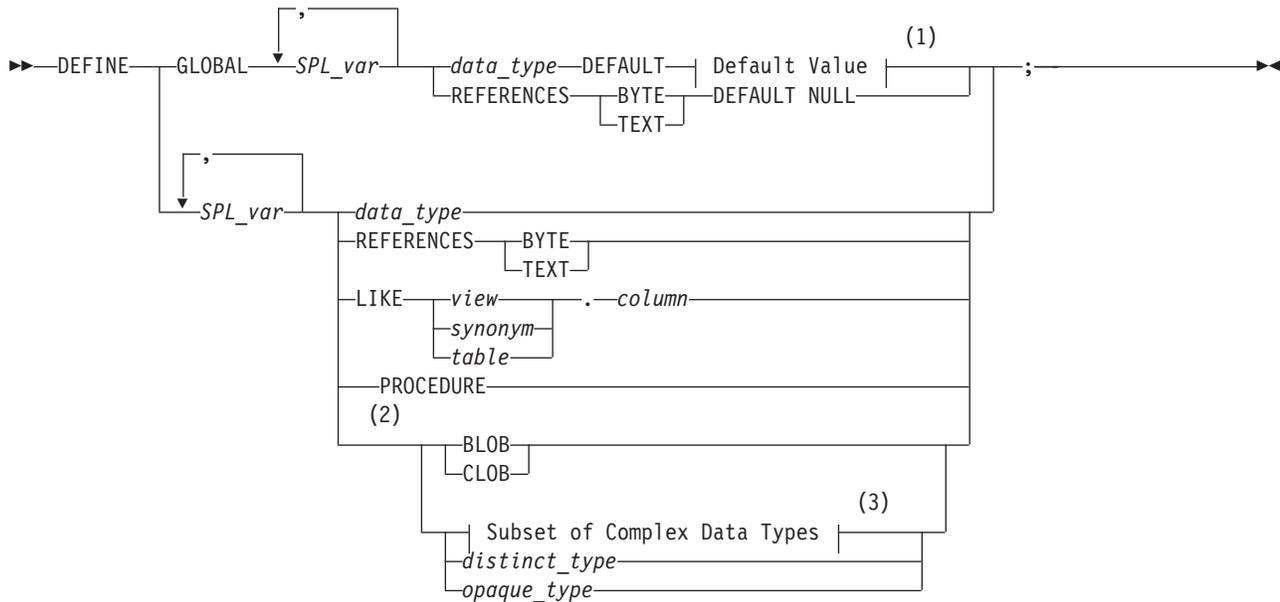
Related Statements

“FOR ” on page 3-19, “FOREACH ” on page 3-22, “LOOP ” on page 3-34, “WHILE ” on page 3-49

DEFINE

Use the DEFINE statement to declare local variables that an SPL routine uses, or to declare global variables that can be shared by several SPL routines.

Syntax



Notes:

- 1 See "Default Value" on page 3-10
- 2 Dynamic Server only
- 3 See "Subset of Complex Data Types (IDS)" on page 3-13

Element	Description	Restrictions	Syntax
<i>column</i>	Column name	Must already exist in the <i>table</i> or <i>view</i>	"Identifier " on page 5-21
<i>data_type</i>	Type of <i>SPL_var</i>	See "Declaring Global Variables" on page 3-9	"Data Type " on page 4-21
<i>distinct_type</i>	A distinct type	Must already be defined in the database	"Data Type " on page 4-21
<i>opaque_type</i>	An opaque type	Must already be defined in the database	"Data Type " on page 4-21
<i>SPL_var</i>	New SPL variable	Must be unique within statement block	"Identifier " on page 5-21
<i>synonym, table, view</i>	Name of a table, view, or synonym	Synonym and the table or view to which it points must exist when DEFINE is issued	"Identifier " on page 5-21

Usage

The DEFINE statement is not an executable statement. The DEFINE statement must appear after the routine header and before any other statements. If you declare a local variable (by using DEFINE without the GLOBAL keyword), its

scope of reference is the statement block in which it is defined. You can use the variable within the statement block. Another variable outside the statement block with a different definition can have the same name.

A variable with the GLOBAL keyword is global in scope and is available outside the statement block and to other SPL routines. Global variables can be any built-in data type except BIGSERIAL, BLOB, BYTE, CLOB, SERIAL, SERIAL8, or TEXT. Local variables can be any built-in data type except BIGSERIAL, BYTE, SERIAL, SERIAL8, or TEXT. If *column* is of the BIGSERIAL, SERIAL, or SERIAL8 data type, declare a BIGINT, INT, or INT8 variable to store its value.

Referencing TEXT and BYTE Variables

The REFERENCES keyword lets you use BYTE and TEXT variables. These do not contain the actual data but are pointers to the data. The REFERENCES keyword indicates that the SPL variable is just a pointer. You can use BYTE and TEXT variables exactly as you would use any other variable in SPL.

Redeclaration or Redefinition

If you define the same variable twice in the same statement block, you receive an error. You can redefine a variable within a nested block, in which case it temporarily hides the outer declaration. This example produces an error:

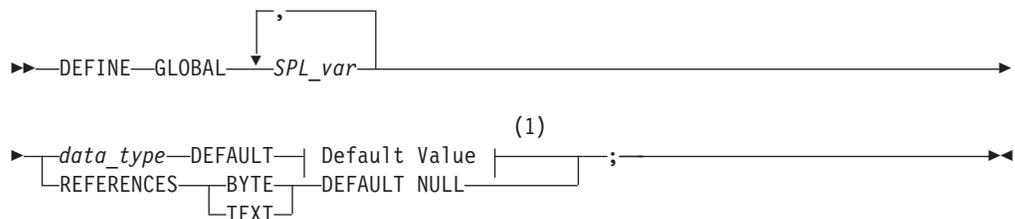
```
CREATE PROCEDURE example1()
  DEFINE n INT; DEFINE j INT;
  DEFINE n CHAR (1);    -- redefinition produces an error
```

Redeclaration is valid in the following example. Within the nested statement block, **n** is a character variable. Outside the block, **n** is an integer variable.

```
CREATE PROCEDURE example2()
  DEFINE n INT; DEFINE j INT;
  ...
  BEGIN
  DEFINE n CHAR (1);    -- character n masks global integer variable
  ...
END;
```

Declaring Global Variables

Use the following syntax for declaring global variables:



Notes:

- 1 See "Default Value" on page 3-10

Element	Description	Restrictions	Syntax
<i>data_type</i>	Type of <i>SPL_var</i>	See "Declaring Global Variables."	"Data Type " on page 4-21

Element	Description	Restrictions	Syntax
<i>SPL_var</i>	New SPL variable	Must be unique within statement block	"Identifier " on page 5-21

The GLOBAL keyword indicates that the variables that follow have a scope of reference that includes all SPL routines that run in a given DB-Access or SQL API session. The data types of these variables must match the data types of variables in the *global environment*. The global environment is the memory that is used by all the SPL routines that run in a given DB-Access or SQL API session. The values of global variables are stored in memory.

SPL routines that are running in the current session share global variables. Because the database server does not save global variables in the database, the global variables do not remain when the current session closes.

The first declaration of a global variable establishes the variable in the global environment; subsequent global declarations simply bind the variable to the global environment and establish the value of the variable at that point.

The following example shows two SPL procedures, **proc1** and **proc2**; each has defined the global variable **gl_out**:

- SPL procedure **proc1**

```
CREATE PROCEDURE proc1()
...
  DEFINE GLOBAL gl_out INT DEFAULT 13;
...
  LET gl_out = gl_out + 1;
END PROCEDURE;
```
- SPL procedure **proc2**

```
CREATE PROCEDURE proc2()
...
  DEFINE GLOBAL gl_out INT DEFAULT 23;
  DEFINE tmp INT;
...
  LET tmp = gl_out
END PROCEDURE;
```

If **proc1** is called first, **gl_out** is set to 13 and then incremented to 14. If **proc2** is then called, it sees that **gl_out** is already defined, so the default value of 23 is not applied. Then, **proc2** assigns the existing value of 14 to **tmp**. If **proc2** had been called first, **gl_out** would have been set to 23, and 23 would have been assigned to **tmp**. Later calls to **proc1** would not apply the default of 13.

Databases of different database server instances do not share global variables, but all the databases of the same database server instance can share global SPL variables in a single session. The database server and any application development tools, however, do not share global variables.

Default Value

Global variables can have literal, NULL, or system constant default values.

Default Value:

TODAY

If you use **TODAY** as the default, the variable must be a DATE value. (See “Constant Expressions” on page 4-70 for descriptions of **TODAY** and of the other system constants that can appear in the Default Value clause.)

BYTE and TEXT

The only default value valid for a BYTE or TEXT variable is NULL. The following example defines a TEXT global variable that is called **l_blob**:

```
CREATE PROCEDURE use_text()  
  DEFINE i INT;  
  DEFINE GLOBAL l_blob REFERENCES TEXT DEFAULT NULL;  
  ...  
END PROCEDURE
```

Here the REFERENCES keyword is required, because the DEFINE statement cannot declare a BYTE or TEXT data type directly; the **l_blob** variable is a pointer to a TEXT value that is stored in the global environment.

SITENAME or DBSERVERNAME

If you use the SITENAME or DBSERVERNAME keyword as the default, the variable must be a CHAR, VARCHAR, NCHAR, NVARCHAR, or LVARCHAR data type. Its default value is the name of the database server at runtime. It is recommended that the size of the variable be at least 128 bytes long. You risk getting an error message during INSERT and ALTER TABLE operations if the length of the variable is too small to store the default value.

The following example uses the SITENAME keyword to specify a default value. This example also initializes a global BYTE variable to NULL:

```
CREATE PROCEDURE g1_def()  
  DEFINE GLOBAL g1_site CHAR(200) DEFAULT SITENAME;  
  DEFINE GLOBAL g1_byte REFERENCES BYTE DEFAULT NULL;  
  ...  
END PROCEDURE
```

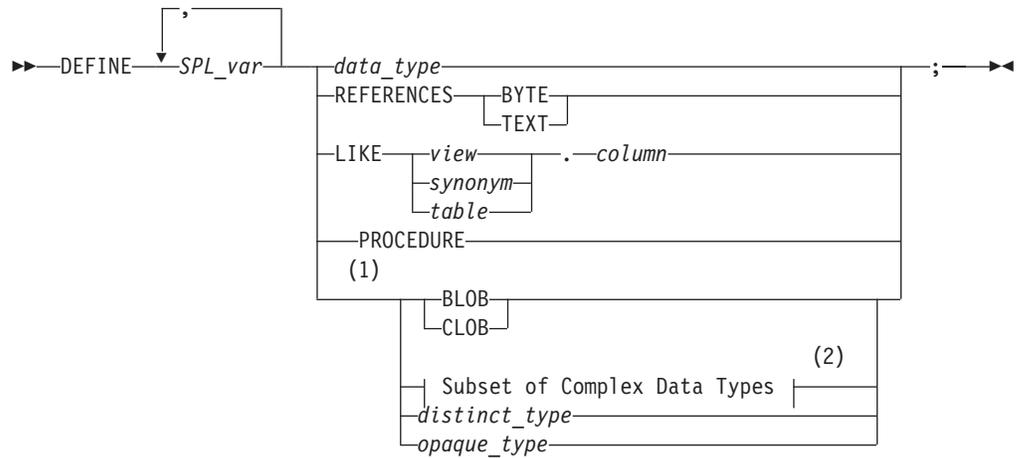
Declaring Local Variables

A *local variable* has as its scope of reference the routine in which it is declared. If you omit the GLOBAL keyword, any variables declared in the DEFINE statement are local variables, and are not visible in other SPL routines.

For this reason, different SPL routines that declare local variables of the same name can run without conflict in the same DB-Access or SQL API session.

If a local variable and a global variable have the same name, the global variable is not visible within the SPL routine where the local variable is declared. (In all other SPL routines, only the global variable is in scope.)

The following DEFINE statement syntax is for declaring local variables:



Notes:

- 1 Dynamic Server only
- 2 See " Subset of Complex Data Types (IDS)"

Element	Description	Restrictions	Syntax
<i>column</i>	Column name	Must already exist in the <i>table</i> or <i>view</i>	"Identifier " on page 5-21;
<i>data_type</i>	Type of <i>SPL_var</i>	Cannot be BIGSERIAL, BYTE, SERIAL, SERIAL8, or TEXT	"Data Type " on page 4-21
<i>distinct_type</i>	A distinct type	Must already be defined in the database	"Identifier " on page 5-21
<i>opaque_type</i>	An opaque type	Must already be defined in the database	"Identifier " on page 5-21
<i>SPL_var</i>	New SPL variable	Must be unique within statement block	"Identifier " on page 5-21;
<i>synonym, table, view</i>	Name of a table, view, or synonym	Synonym and the table or view to which it points must already exist when the statement is issued	"Database Object Name " on page 5-16

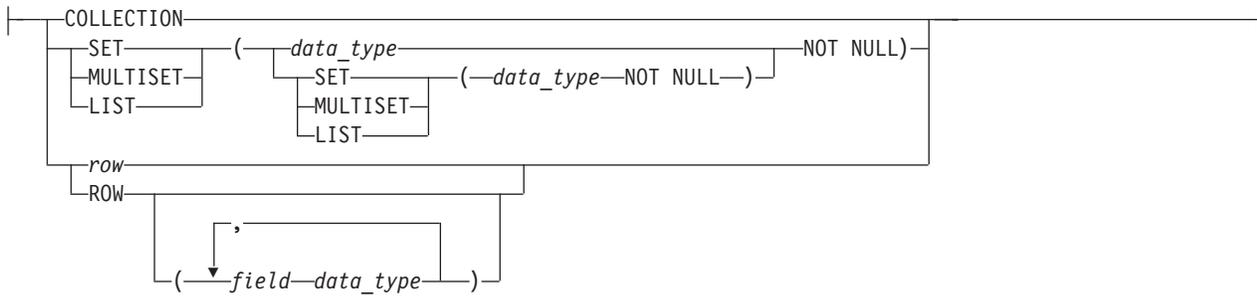
Local variables do not support default values. The following example shows some typical definitions of local variables:

```
CREATE PROCEDURE def_ex()
  DEFINE i INT;
  DEFINE word CHAR(15);
  DEFINE b_day DATE;
  DEFINE c_name LIKE customer.fname;
  DEFINE b_text REFERENCES TEXT;
END PROCEDURE
```

Subset of Complex Data Types (IDS)

You can use the following syntax to declare an SPL variable as a typed or generic collection, or as a named, unnamed, or generic ROW data type.

Complex Data Types (Subset):



Element	Description	Restrictions	Syntax
<i>data_type</i>	Type of elements of a collection or of fields of an unnamed ROW type	Must match the data type of the values that the variable will store. Cannot be BIGSERIAL, BLOB, BYTE, CLOB, SERIAL, SERIAL8, or TEXT.	"Data Type " on page 4-21
<i>field</i>	Field of unnamed ROW	Must exist in the database	"Identifier " on page 5-21
<i>row</i>	Named ROW data type	Must exist in the database	"Identifier " on page 5-21

Declaring Collection Variables (IDS)

A local variable of type COLLECTION, SET, MULTISET, or LIST can hold a collection of values fetched from the database. You cannot define a collection variable as global (with the GLOBAL keyword) or with a default value.

A variable declared with the keyword COLLECTION is an untyped (or generic) collection variable that can hold a collection of any data type.

A variable declared as type SET, MULTISET, or LIST is a *typed collection variable*. It can hold a collection of its specified data type only.

You must use the NOT NULL keywords when you define the elements of a typed collection variable, as in the following examples:

```
DEFINE a SET ( INT NOT NULL );

DEFINE b MULTISET ( ROW ( b1 INT,
                        b2 CHAR(50)
                      ) NOT NULL );

DEFINE c LIST( SET( INTEGER NOT NULL ) NOT NULL );
```

With variable *c*, both the INTEGER values in the SET and the SET values in the LIST are defined as NOT NULL.

You can define collection variables with nested complex types to hold matching nested complex type data. Any type or depth of nesting is allowed. You can nest ROW types within collection types, collection types within ROW types, collection types within collection types, ROW types within collection and ROW types, and so on.

If you declare a variable as COLLECTION type, the variable acquires varying data type declarations if it is reassigned within the same statement block, as in the following example:

```

DEFINE a COLLECTION;
LET a = setB;
...
LET a = listC;

```

In this example, **varA** is a generic collection variable that changes its data type to the data type of the currently assigned collection. The first LET statement makes **varA** a SET variable. The second LET statement makes **varA** a LIST variable.

Declaring ROW Variables (IDS)

ROW variables hold data from named or unnamed ROW types. You can define a generic ROW variable, a named ROW variable, or an unnamed ROW variable.

A generic ROW variable, defined with the ROW keyword, can hold data from any ROW type. A named ROW variable holds data from the named ROW type that you specified in the declaration of the variable.

The following statements show examples of generic ROW variables and named ROW variables:

```

DEFINE d ROW;           -- generic ROW variable

DEFINE rectv rectangle_t; -- named ROW variable

```

A named ROW variable holds named ROW types of the same type in the declaration of the variable.

To define a variable that will hold data stored in an unnamed ROW type, use the ROW keyword followed by the fields of the ROW type, as in:

```

DEFINE area ROW ( x int, y char(10) );

```

Unnamed ROW types are type-checked only by structural equivalence. Two unnamed ROW types are considered equivalent if they have the same number of fields, and if the fields have the same type definitions. Therefore, you could fetch either of the following ROW types into the variable **area** defined above:

```

ROW ( a int, b char(10) )
ROW ( area int, name char(10) )

```

ROW variables can have fields, just as ROW types have fields. To assign a value to a field of a ROW variable, use the qualifier notation *variableName.fieldName*, followed by an expression, as in the following example:

```

CREATE ROW TYPE rectangle_t (start point_t, length real, width real);

DEFINE r rectangle_t;
    -- Define a variable of a named ROW type
LET r.length = 45.5;
    -- Assign a value to a field of the variable

```

When you assign a value to a ROW variable, you can use any valid expression.

Declaring Opaque-Type Variables (IDS)

Opaque-type variables hold data retrieved from opaque data types, which you create with the CREATE OPAQUE TYPE statement. An opaque-type variable can only hold data of the same opaque type on which it is defined. The following example defines a variable of the opaque type **point**, which holds the x and y coordinates of a two-dimensional point:

```

DEFINE b point;

```

Declaring Variables LIKE Columns

If you use the LIKE clause, the database server assigns the variable the same data type as a specified column in a table, synonym, or view.

The data types of variables that are defined as database columns are resolved at runtime; therefore, *column* and *table* do not need to exist at compile time.

You can use the LIKE keyword to declare that a variable is like a serial column. This declares an INTEGER variable if the column is of the SERIAL data type, or an INT8 variable if the column is of the SERIAL8 data type. For example, if the column **serialcol** in the **mytab** table has the SERIAL data type, you can create the following SPL function:

```
CREATE FUNCTION func1()
DEFINE local_var LIKE mytab.serialcol;
RETURN;
END FUNCTION;
```

The variable **local_var** is treated as an INTEGER variable.

Declaring Variables as the PROCEDURE Type

The PROCEDURE keyword indicates that in the current scope, the variable is a call to a UDR.

The DEFINE statement does not support a FUNCTION keyword. Use the PROCEDURE keyword, whether you are calling a user-defined procedure or a user-defined function.

Declaring a variable as PROCEDURE type indicates that in the current statement scope, the variable is not a call to a built-in function. For example, the following statement defines **length** as an SPL routine, not as the built-in LENGTH function:

```
DEFINE length PROCEDURE;
...
LET x = length (a,b,c)
```

This definition disables the built-in LENGTH function within the scope of the statement block. You would use such a definition if you had already created a user-defined routine with the name **length**.

If you create an SPL routine with the same name as an aggregate function (SUM, MAX, MIN, AVG, COUNT) or with the name **extend**, you must qualify the routine name with the owner name.

Declaring Variables for BYTE and TEXT Data

The keyword REFERENCES indicates that the variable does not contain a BYTE or TEXT value but is a pointer to the BYTE or TEXT value. Use the variable as though it holds the data.

The following example defines a local BYTE variable:

```
CREATE PROCEDURE use_byte()
  DEFINE i INT;
  DEFINE l_byte REFERENCES BYTE;
END PROCEDURE --use_byte
```

If you pass a variable of BYTE or TEXT data type to an SPL routine, the data is passed to the database server and stored in the root dbspace or dbspaces that the **DBSPACETEMP** environment variable specifies, if it is set. You do not need to

know the location or name of the file that holds the data. BYTE or TEXT manipulation requires only the name of the BYTE or TEXT variable as it is defined in the routine.

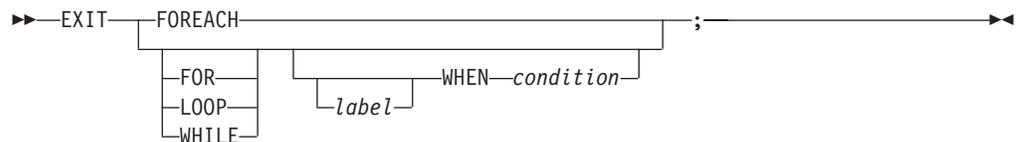
Related Information

Declaring the names of SQL keywords or the identifiers of other database objects as SPL variables can produce errors or unexpected results in some contexts. See the sections “Variables that Conflict with Column Names” on page 5-32, “Declaring Keywords or Routine Names as SPL Variables” on page 5-31, “Using NULL and SELECT in a Condition” on page 5-31 and “Using CURRENT, DATETIME, INTERVAL, and NULL in INSERT” on page 5-31 for discussions of some potential problems of name conflicts that involve SPL variables.

EXIT

The EXIT statement can terminate FOR, FOREACH, LOOP, or WHILE statements.

Syntax



Element	Description	Restrictions	Syntax
<i>condition</i>	Loop terminates when this evaluates to TRUE.	If <i>condition</i> evaluates to FALSE, the loop continues.	“Condition ” on page 4-5
<i>label</i>	Label of a loop from which to exit	Must be the label of a loop statement that includes the EXIT statement	“Identifier ” on page 5-21

Usage

The EXIT statement transfers control of execution from an iterative statement, causing the innermost loop of the enclosing statement type (FOR, FOREACH, LOOP, or WHILE) to terminate. If no loop label or WHEN condition is specified, execution resumes at the first statement that follows the current FOR, FOREACH, LOOP, or WHILE statement.

EXIT From FOREACH Statements

If the EXIT statement has the FOREACH statement as its innermost enclosing statement, the FOREACH keyword must immediately follow the EXIT keyword. The EXIT FOREACH statement unconditionally terminates the FOREACH statement, or else returns an error, if no FOREACH statement encloses the EXIT FOREACH statement.

The following program fragment includes the EXIT FOREACH statement:

```
FOREACH cursor1 FOR
  SELECT * INTO a FROM TABLE(b);
  IF a = 4 THEN
    DELETE FROM TABLE(b)
```

```

        WHERE CURRENT OF cursor1;4
    EXIT FOREACH;
END IF;
END FOREACH;

```

EXIT From FOR, LOOP, and WHILE Loops

If the EXIT statement is issued outside the FOREACH statement, it returns an error unless it is issued from the FOR, FOR LOOP, LOOP, WHILE LOOP, or WHILE statement as its innermost enclosing statement. In FOR or WHILE statements that do not include the LOOP keyword, the corresponding FOR or WHILE keyword is required after the EXIT keyword. Execution resumes at the first executable statement that follows the innermost loop from which the EXIT statement was issued.

The EXIT statement requires no other keyword when it is issued from the FOR LOOP, LOOP, or WHILE LOOP statement, with or without a loop label, but if you include the FOR, LOOP, or WHILE keyword after the EXIT keyword, that keyword must correspond to the type of loop from which the EXIT statement is issued.

If the EXIT keyword is followed by the identifier of a loop label, and no *condition* is specified, execution resumes at the first executable statement that follows the FOR, FOR LOOP, LOOP, WHILE LOOP, or WHILE statement whose label is specified. This enables the EXIT statement to exit from nested loops, if an outer loop is labeled.

If a WHEN *condition* follows the EXIT or EXIT *label* specification, EXIT has no effect unless the *condition* is true. If the condition is true, execution resumes after the labeled loop, or after the innermost loop, if no *label* is specified.

If the database server cannot find the specified loop or loop label, the EXIT statement fails. If EXIT is issued outside any FOR, FOREACH, LOOP, or WHILE statement, it generates errors.

The following example uses an EXIT FOR statement. In the FOR loop, when *j* becomes 6, the IF condition *i* = 5 in the WHILE loop is true. The FOR loop stops executing, and the SPL procedure continues at the next statement outside the FOR loop (in this case, the END PROCEDURE statement). In this example, the procedure ends when *j* equals 6:

```

CREATE PROCEDURE ex_cont_ex()
  DEFINE i,s,j, INT;
  FOR j = 1 TO 20
    IF j > 10 THEN
      CONTINUE FOR;
    END IF
    LET i,s = j,0;
    WHILE i > 0
      LET i = i -1;
      IF i = 5 THEN
        EXIT FOR;
      END IF
    END WHILE
  END FOR
END PROCEDURE;

```

The following program fragment shows two conditional EXIT statements in a labeled WHILE LOOP statement that is nested within another labeled LOOP statement:

```

<<outer>>
LOOP
LET x = x+1;
  <<inner>>
  WHILE ( i >10 ) LOOP
    LET x = x+1;
    EXIT inner WHEN x = 2;
    EXIT outer WHEN x > 3;
    END LOOP inner;
  LET x = x+1;
END LOOP outer;

```

When the x=2 condition is true, the EXIT inner statement transfers control to the LET statement that follows the loop whose label is **inner**. When the x>3 condition is true, the EXIT outer statement terminates execution of the **outer** loop.

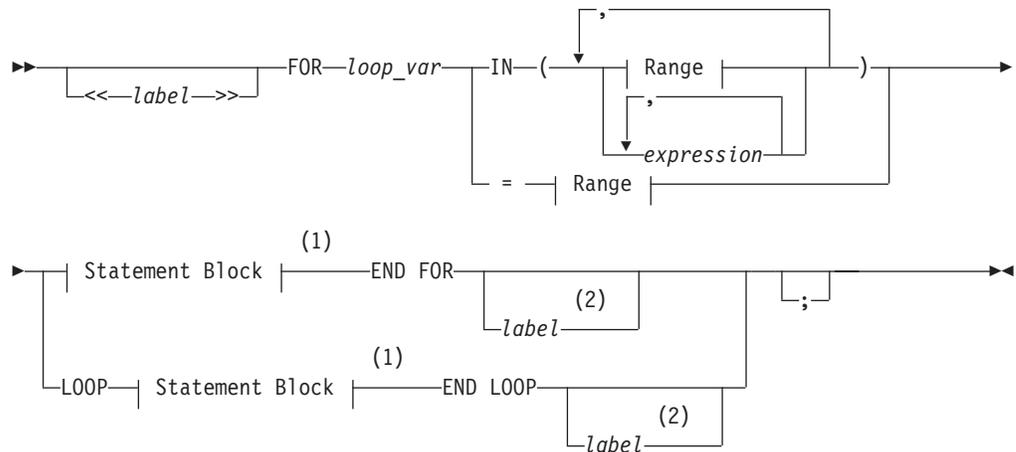
Related Statements

"<< Label >>" on page 3-1, "FOR," "FOREACH" on page 3-22, "LOOP" on page 3-34, "WHILE" on page 3-49

FOR

Use the FOR statement to initiate a controlled (definite) loop when you want to guarantee termination of the loop. The FOR statement uses expressions or range operators to specify a finite number of iterations for a loop.

Syntax



Range:

```

| left_expression TO right_expression |
| STEP increment_expr |

```

Notes:

- 1 See "Statement Block" on page 5-74
- 2 Valid only if <<label>> precedes the first FOR keyword

Element	Description	Restrictions	Syntax
<i>expression</i>	Value to compare with <i>loop_var</i>	Must match <i>loop_var</i> data type	"Expression" on page 4-40
<i>increment_expr</i>	Positive or negative value by which <i>loop_var</i> is incremented	Must return an integer. Cannot return 0.	"Expression" on page 4-40
<i>label</i>	Name of the loop label for this loop	Must be unique among labels in this SPL routine	"Identifier " on page 5-21
<i>left_expression</i>	Starting expression of a range	Value must match SMALLINT or INT data type of <i>loop_var</i>	"Expression" on page 4-40
<i>loop_var</i>	Variable that determines how many times the loop executes	Must be defined and in scope within this statement block	"Identifier " on page 5-21
<i>right_expression</i>	Ending expression in the range	Same as for <i>left_expression</i>	"Expression" on page 4-40

Usage

The database server evaluates all expressions before the FOR statement executes. If one or more of the expressions are variables whose values change during the loop, the change has no effect on the iterations of the loop.

You can use the output from a SELECT statement as the *expression*.

The FOR loop terminates when *loop_var* is equal to the values of each element in the expression list or range in succession, or when it encounters an EXIT FOR statement. An error is issued, however, if an assignment within the body of the FOR statement attempts to modify the value of *loop_var*.

The size of *right_expression* relative to *left_expression* determines if the range is stepped through by positive or negative increments.

Using the TO Keyword to Define a Range

The TO keyword implies a range operator. The range is defined by *left_expression* and *right_expression*, and the STEP *increment_expr* option implicitly sets the number of increments. If you use the TO keyword, *loop_var* must be an INT or SMALLINT data type.

The next examples shows two equivalent FOR statements. Each uses the TO keyword to define a range. The first uses the IN keyword, and the second uses an equal sign (=). Each statement causes the loop to execute five times:

```
FOR index_var IN (12 TO 21 STEP 2)
  -- statement block
END FOR;

FOR index_var = 12 TO 21 STEP 2
  -- statement block
END FOR;
```

If you omit the STEP option, the database server gives *increment_expr* the value of -1 if *right_expression* is less than *left_expression*, or +1 if *right_expression* is more than *left_expression*. If *increment_expr* is specified, it must be negative if *right_expression* is less than *left_expression*, or positive if *right_expression* is more than *left_expression*.

The two statements in the following example are equivalent. In the first statement, the STEP increment is explicit. In the second statement, the STEP increment is implicitly 1:

```
FOR index IN (12 TO 21 STEP 1)
  -- statement block
END FOR;
```

```
FOR index = 12 TO 21
  -- statement block
END FOR;
```

The database server initializes the value of *loop_var* to the value of *left_expression*. In subsequent iterations, the server adds *increment_expr* to the value of *loop_var* and checks *increment_expr* to determine whether the value of *loop_var* is still between *left_expression* and *right_expression*. If so, the next iteration occurs. Otherwise, an exit from the loop takes place. Or, if you specify another range, the variable takes on the value of the first element in the next range.

Specifying Two or More Ranges in a Single FOR Statement

The following example shows a statement that traverses a loop forward and backward and uses different increment values for each direction:

```
FOR index_var IN (15 to 21 STEP 2, 21 to 15 STEP -3)
  -- statement body
END FOR;
```

Using an Expression List as the Range

The database server initializes the value of *loop_var* to the value of the first expression specified. In subsequent iterations, *loop_var* takes on the value of the next expression. When the database server has evaluated the last expression in the list and used it, the loop stops.

The expressions in the IN list do not need to be numeric values, as long as you do not use range operators in the IN list. The following example uses a character expression list:

```
FOR c IN ('hello', (SELECT name FROM t), 'world', v1, v2)
  INSERT INTO t VALUES (c);
END FOR;
```

The following FOR statement shows the use of a numeric expression list:

```
FOR index IN (15,16,17,18,19,20,21)
  -- statement block
END FOR;
```

Mixing Range and Expression Lists in the Same FOR Statement

If *loop_var* is an INT or SMALLINT value, you can mix ranges and expression lists in the same FOR statement. The following example shows a mixture that uses an integer variable. Values in the expression list include the value that is returned from a SELECT statement, a sum of an integer variable and a constant, the values that are returned from an SPL function named `p_get_int`, and integer constants:

```
CREATE PROCEDURE for_ex ()
  DEFINE i, j INT;
  LET j = 10;
  FOR i IN (1 TO 20, (SELECT c1 FROM tab WHERE id = 1),
```

```

        j+20 to j-20, p_get_int(99),98,90 to 80 step -2)
    INSERT INTO tab VALUES (i);
END FOR;
END PROCEDURE;

```

Specifying a Labelled FOR Loop

To create a labeled FOR loop, declare a loop label before the initial FOR keyword, and repeat the label after the END FOR keywords, as in this example:

```

CREATE PROCEDURE ex_cont_ex()
  DEFINE i,s,j, INT;
  <<for_lab>>
  FOR j = 1 TO 20
    IF j > 10 THEN
      CONTINUE FOR;
    END IF
    LET i,s = j,0;
    WHILE i > 0
      LET i = i -1;
      IF i = 5 THEN
        EXIT for_lab;
      END IF
    END WHILE
  END FOR for_lab
END PROCEDURE;

```

Here the EXIT for_lab statement has the same effect that the EXIT or EXIT FOR keywords would have, terminating both the FOR loop and the routine. In this example, the statement that includes the EXIT for_lab statement has the same effect that EXIT for_lab WHEN i = 5 would have.

You can also label a LOOP statement that begins with a loop <<label>> specification that immediately precedes the initial FOR keyword. In this type of loop, the CONTINUE LOOP, EXIT LOOP, and END LOOP keywords replace the CONTINUE FOR, EXIT FOR, and END FOR keywords. Both the LOOP and FOR keywords are optional after the CONTINUE and EXIT keywords, but the END LOOP keywords are required in SPL loop statements that include the LOOP keyword.

You can use similar syntax to create an unlabeled loop that omits the <<label>> specification that immediately precedes the initial FOR keyword. In this case, you must also omit the undelimited loop label identifier that follows the END LOOP keywords. See the LOOP statement for a description and examples of these forms of labeled and unlabeled loop statements that enable you to combine FOR statement syntax, with a finite number of loop iterations, with the "loop forever" syntax of the LOOP statement.

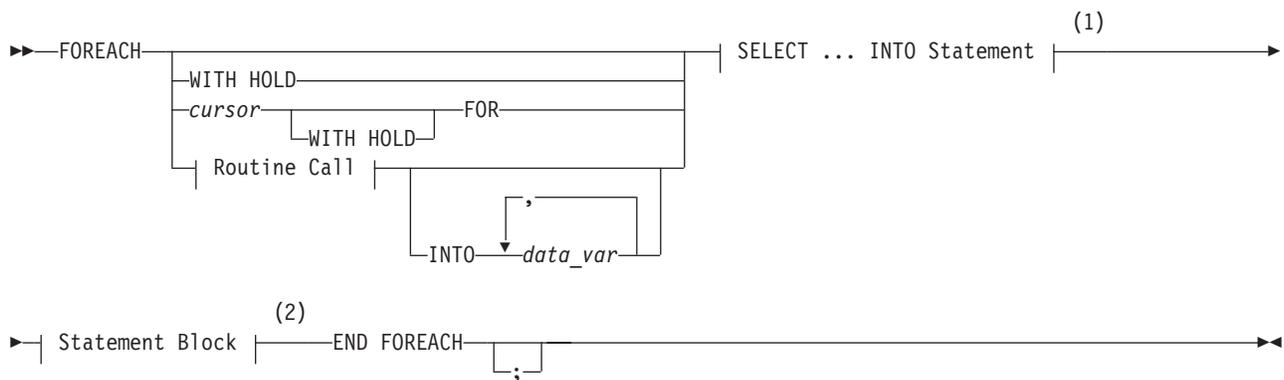
Related Statements

"<< Label >>" on page 3-1, "CONTINUE" on page 3-7, "EXIT " on page 3-17, "LOOP " on page 3-34, "FOREACH ," "WHILE " on page 3-49

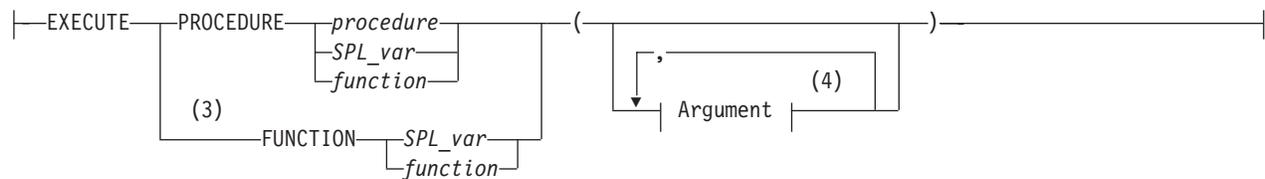
FOREACH

Use the FOREACH statement to select and manipulate more than one row or collection.

Syntax



Routine Call:



Notes:

- 1 See "Using a SELECT ... INTO Statement" on page 3-24
- 2 See "Statement Block " on page 5-74
- 3 Dynamic Server only
- 4 See "Arguments " on page 5-1

Element	Description	Restrictions	Syntax
<i>cursor</i>	Identifier that you supply as a name for this FOREACH loop	Each cursor name within a routine must be unique	"Identifier " on page 5-21
<i>data_var</i>	SPL variable in the calling routine that receives the returned values	Data type of <i>data_var</i> must be appropriate for returned value	"Identifier " on page 5-21
<i>function, procedure</i>	SPL function or procedure to execute	Function or procedure must exist	"Database Object Name " on page 5-16
<i>SPL_var</i>	SPL variable that contains the name of a routine to execute	Must be a CHAR, VARCHAR, NCHAR, or NVARCHAR type	"Identifier " on page 5-21

Usage

A FOREACH loop is the procedural equivalent of using a cursor. To execute a FOREACH statement, the database server takes these actions:

1. It declares and implicitly opens a cursor.
2. It obtains the first row from the query contained within the FOREACH loop, or else the first set of values from the called routine.
3. It assigns to each variable in the variable list the value of the corresponding value from the active set that the SELECT statement or the called routine creates.

4. It executes the statement block.
5. It fetches the next row from the SELECT statement or called routine on each iteration, and it repeats steps 3 and 4.
6. It terminates the loop when it finds no more rows that satisfy the SELECT statement or called routine. It closes the implicit cursor when the loop terminates.

Because the statement block can contain additional FOREACH statements, cursors can be nested. No limit exists on the number of nested cursors.

An SPL routine that returns more than one row, collection element, or set of values is called a *cursor function*. An SPL routine that returns only one row or value is a *noncursor function*.

This SPL procedure illustrates FOREACH statements with a SELECT ... INTO clause, with an explicitly named cursor, and with a procedure call:

```
CREATE PROCEDURE foreach_ex()
  DEFINE i, j INT;
  FOREACH SELECT c1 INTO i FROM tab ORDER BY 1
    INSERT INTO tab2 VALUES (i);
  END FOREACH
  FOREACH cur1 FOR SELECT c2, c3 INTO i, j FROM tab
    IF j > 100 THEN
      DELETE FROM tab WHERE CURRENT OF cur1;
      CONTINUE FOREACH;
    END IF
    UPDATE tab SET c2 = c2 + 10 WHERE CURRENT OF cur1;
  END FOREACH
  FOREACH EXECUTE PROCEDURE bar(10,20) INTO i
    INSERT INTO tab2 VALUES (i);
  END FOREACH
END PROCEDURE; -- foreach_ex
```

A Select cursor is closed when any of the following situations occur:

- The cursor returns no further rows.
- The cursor is a Select cursor without a HOLD specification, and a transaction completes using COMMIT or ROLLBACK statements.
- An EXIT statement executes, which transfers control out of the FOREACH statement.
- An exception occurs that is not trapped inside the body of the FOREACH statement. (See "ON EXCEPTION " on page 3-38.)
- A cursor in the calling routine that is executing this cursor routine (within a FOREACH loop) closes for any reason.

Using a SELECT ... INTO Statement

As indicated in the diagram for "FOREACH " on page 3-22, not all clauses and options of the SELECT statement are available for you to use in a FOREACH statement. The SELECT statement in the FOREACH statement must include the INTO clause. It can also include UNION and ORDER BY clauses, but it cannot use the INTO TEMP clause. For a complete description of SELECT syntax and usage, see "SELECT" on page 2-528. The data type and count of each variable in the variable list must match each value that the SELECT ... INTO statement returns.

Using the ORDER BY Clause of the SELECT Statement

The ORDER BY clause of the SELECT statement implies that the query returns more than one row. The database server issues an error if you specify the ORDER BY clause without a FOREACH loop to process the returned rows individually within an SPL routine.

Using Hold Cursors

The WITH HOLD keywords specify that the cursor should remain open when a transaction closes (by being committed or by being rolled back).

Updating or Deleting Rows Identified by Cursor Name

Specify a *cursor* name in the FOREACH statement if you intend to use the WHERE CURRENT OF *cursor* clause in UPDATE or DELETE statements that operate on the current row of *cursor* within the FOREACH loop. Although you cannot include the FOR UPDATE keywords in the SELECT ... INTO segment of the FOREACH statement, the cursor behaves like a FOR UPDATE cursor.

For a discussion of locking, see the section on “ Locking with an Update Cursor” on page 2-292. For a discussion of isolation levels, see the description of “SET ISOLATION” on page 2-631.

Using Collection Variables (IDS)

The FOREACH statement allows you to declare a cursor for an SPL collection variable. Such a cursor is called a *Collection cursor*. Use a collection variable to access the elements of a collection (SET, MULTISSET, LIST) column. Use a cursor when you want to access one or more elements in a collection variable.

Restrictions

When you use a Collection cursor to fetch individual elements from a collection variable, the FOREACH statement has the following restrictions:

- It cannot contain the WITH HOLD keywords.
- It must contain a restricted SELECT statement in the FOREACH loop.

In addition, the SELECT statement that you associate with the Collection cursor has the following restrictions:

- Its general structure is SELECT... INTO ... FROM TABLE. The statement selects one element at a time from a collection variable specified after the TABLE keyword into another variable called an *element variable*.
- It cannot contain an expression in the projection list.
- It cannot include the following clauses or options: WHERE, GROUP BY, ORDER BY, HAVING, INTO TEMP, and WITH REOPTIMIZATION.
- The data type of the element variable must be the same as the element type of the collection.
- The data type of the element variable can be any opaque, distinct, or collection data type, or any built-in data type except BIGSERIAL, BLOB, BYTE, CLOB, SERIAL, SERIAL8, or TEXT.
- If the collection contains opaque, distinct, built-in, or collection types, the projection list must be an asterisk (*) symbol.
- If the collection contains ROW types, the projection list can be a list of one or more field names.

Examples

The following excerpt from an SPL routine shows how to fill a collection variable and then how to use a cursor to access individual elements:

```
DEFINE a SMALLINT;
DEFINE b SET(SMALLINT NOT NULL);
SELECT numbers INTO b FROM table1 WHERE id = 207;
FOREACH cursor1 FOR
    SELECT * INTO a FROM TABLE(b);
...
END FOREACH;
```

In this example, the SELECT statement selects one element at a time from the collection variable **b** into the element variable **a**. The projection list is an asterisk, because the collection variable **b** contains a collection of built-in types. The variable **b** is used with the TABLE keyword as a Collection-Derived Table. For more information, see “Collection-Derived Table” on page 5-4.

The next example also shows how to fill a collection variable and then how to use a cursor to access individual elements. This example, however, uses a list of ROW-type fields in its projection list:

```
DEFINE employees employee_t;
DEFINE n VARCHAR(30);
DEFINE s INTEGER;

SELECT emp_list into employees FROM dept_table
    WHERE dept_no = 1057;
FOREACH cursor1 FOR
    SELECT name,salary
        INTO n,s FROM TABLE( employees ) AS e;
...
END FOREACH;
```

Here the collection variable **employees** contains a collection of ROW types. Each ROW type contains the fields **name** and **salary**. The collection query selects one name and salary combination at a time, placing **name** into **n** and **salary** into **s**. The AS keyword declares **e** as an alias for the collection-derived table **employees**. The alias exists as long as the SELECT statement executes.

Modifying Elements in a Collection Variable

To update an element of a collection within an SPL routine, you must first declare a cursor with the FOREACH statement.

Then, within the FOREACH loop, select elements one at a time from the collection variable, using the collection variable as a collection-derived table in a SELECT query.

When the cursor is positioned on the element to be updated, you can use the WHERE CURRENT OF clause, as follows:

- The UPDATE statement with the WHERE CURRENT OF clause updates the value in the current element of the collection variable.
- The DELETE statement with the WHERE CURRENT OF clause deletes the current element from the collection variable.

Using Select Cursors with FOREACH

When using the FOREACH statement, if the result set from a query is to be modified, do not use this result set as an exit criterion for the FOREACH loop. For example, if the FOREACH statement declares a Select cursor that is expected to

return 30 rows, but DELETE, INSERT, or UPDATE operations within the FOREACH loop modify the result set of the query, this might cause unexpected behavior. To ensure that a FOREACH loop works as intended, make sure that any Select cursor in the FOREACH statement completes its execution before you begin modifying its result set.

One way to avoid unexpected results from a FOREACH loop that performs DML operations on the rows returned by a query is to use an ORDER BY clause in the SELECT statement to materialize the result set.

Calling a UDR in the FOREACH Loop

In general, use these guidelines for calling another UDR from an SPL routine:

- To call a user-defined procedure, use EXECUTE PROCEDURE *procedure name*.
- To call a user-defined function, use EXECUTE FUNCTION *function name* (or EXECUTE PROCEDURE *function name* if the user-defined function was created with the CREATE PROCEDURE statement).

In Extended Parallel Server, you must use EXECUTE PROCEDURE. Extended Parallel Server does not support the EXECUTE FUNCTION statement.

In Dynamic Server, if you use EXECUTE PROCEDURE, the database server looks first for a user-defined procedure of the name you specify. If it finds the procedure, the database server executes it. If it does not find the procedure, it looks for a user-defined function of the same name to execute. If the database server finds neither a function nor a procedure, it issues an error message. If you use EXECUTE FUNCTION, the database server looks for a user-defined function of the name you specify. If it does not find a function of that name, the database server issues an error message.

An SPL function can return zero (0) or more values or rows.

The data type and count of each variable in the variable list must match each value that the function returns.

Related Statements

“CONTINUE” on page 3-7, “EXIT ” on page 3-17, “FOR ” on page 3-19, “LOOP ” on page 3-34, “WHILE ” on page 3-49

GOTO

Use the GOTO statement to transfer control of program execution to the statement that has a specified statement label.

Syntax

►► GOTO *label* ; ◀◀

Element	Description	Restrictions	Syntax
<i>label</i>	Name of the loop label for this loop	Must be unique among labels in this SPL routine	“Identifier ” on page 5-21

Usage

The GOTO statement branches to a statement label unconditionally. The statement label must be unique within its scope and must precede an executable statement. When successfully executed, the GOTO statement transfers control to the labeled statement or statement block.

In the following program fragment, the **jump_back** function transfers control to a LET statement that has the statement label **back** if the value of variable **j** is greater than 100.

```
CREATE FUNCTION jump_back()  
  RETURNING INT;  
  DEFINE i,j INT;  
  ...  
  <<back>>  
  LET j = j + i  
  FOR i IN (1 TO 52 STEP 5)  
    IF i < 11 THEN  
      LET j = j + 3  
      CONTINUE FOR;  
    END IF;  
    IF j > 100 THEN  
      GOTO back  
    END IF;  
    RETURN j WITH RESUME;  
  END FOR;  
END FUNCTION;
```

The GOTO statement is not valid in an ON EXCEPTION statement block.

The identifier of the statement label that the GOTO statement references must exist in the database, must be unique among statement labels and loop labels the SPL routine, and must be within a scope that the GOTO statement can reach.

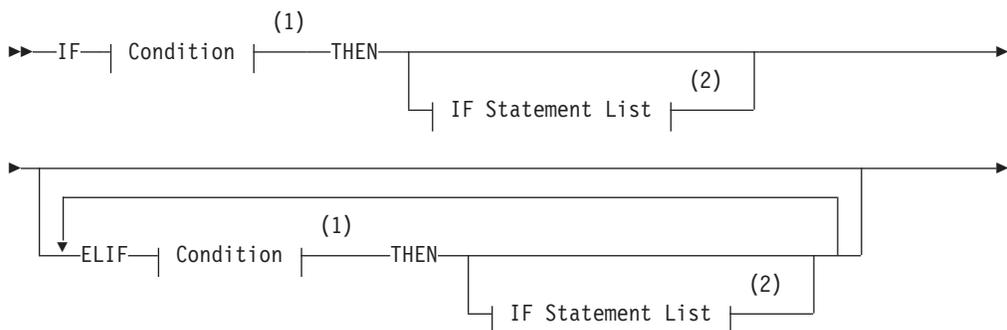
Related Statements

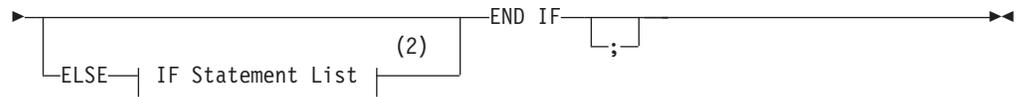
"<< Label >>" on page 3-1

IF

Use the IF statement to create a logical branch within an SPL routine.

Syntax





Notes:

- 1 See “Condition ” on page 4-5
- 2 See “IF Statement List” on page 3-30

Usage

The database server processes the IF statement by the following steps:

1. If the condition that follows the IF keyword is true, any statements that follow the first THEN keyword of the IF statement execute, and the IF statement terminates.
2. If the result of the initial IF condition is false, but an ELIF clause exists, the database server evaluates the condition that follows the ELIF keyword.
3. If the result of the ELIF condition is true, any statements that follow the THEN keyword of the ELIF clause execute, and the IF statement terminates.
4. If the result of the condition in the first ELIF clause is also false, but one or more additional ELIF clauses exist, the database server evaluates the condition in the next ELIF clause, and proceeds as in the previous step if it is true. If it is false, the database server evaluates the condition in successive ELIF clauses, until it finds a condition that is true, in which case it executes the statement list that follows the THEN keyword of that ELIF clause, and the IF statement terminates.
5. If no condition in the IF statement is true, but the ELSE clause exists, statements that follow the ELSE keyword execute, and the IF statement terminates.
6. If none of the conditions in the IF statement are true, and no ELSE clause exists, the IF statement terminates without executing any statement list.

ELIF Clause

Use the ELIF clause to specify one or more additional conditions to evaluate. If the IF condition is false, the ELIF condition is evaluated. If the ELIF condition is true, the statements that follow the THEN keyword in the ELIF clause execute.

If no statement follows the THEN keyword of the ELIF clause when the ELIF condition is true, program control passes from the IF statement to the next statement.

ELSE Clause

The ELSE clause executes if no true previous condition exists in the IF clause or any of the ELIF clauses.

In the following example, the SPL function uses an IF statement with both an ELIF clause and an ELSE clause. The IF statement compares two strings.

The function displays 1 to indicate that the first string comes before the second string alphabetically, or -1 if the first string comes after the second string alphabetically. If the strings are the same, a zero (0) is returned.

```

CREATE FUNCTION str_compare (str1 CHAR(20), str2 CHAR(20))
  RETURNING INT;
  DEFINE result INT;
  IF str1 > str2 THEN LET result =1;
  ELIF str2 > str1 THEN LET result = -1;
  ELSE LET result = 0;
  END IF
  RETURN result;
END FUNCTION -- str_compare

```

Conditions in an IF Statement

Just as in the WHILE statement, if any expression in the *condition* evaluates to NULL, then the condition cannot be true, unless you are explicitly testing for NULL using the IS NULL operator. The following rules summarize NULL values in conditions:

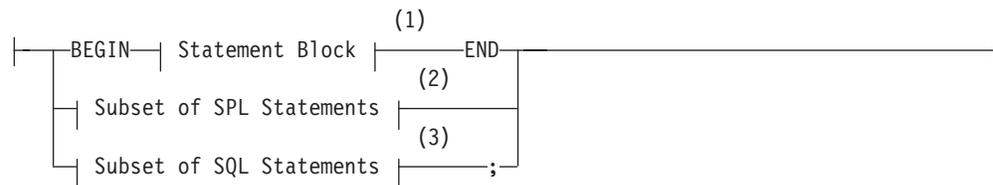
1. If the expression x evaluates to NULL, then x is not true by definition. Furthermore, NOT (x) is also not true .
2. IS NULL is the only operator that can return true for x. That is, x IS NULL is true, and x IS NOT NULL is not true.

If an expression in the condition has an UNKNOWN value from an uninitialized SPL variable, the statement terminates and raises an exception.

You can specify a trigger-type Boolean operator (DELETING, INSERTING, SELECTING, or UPDATING) as a condition in an IF statement only within a trigger routine.

IF Statement List

IF Statement List:



Notes:

- 1 See "Statement Block " on page 5-74
- 2 See "Subset of SPL Statements Allowed in the IF Statement List"
- 3 See "SQL Statements Not Valid in an IF Statement" on page 3-31

Subset of SPL Statements Allowed in the IF Statement List

You can use any of the following SPL statements in the IF statement list:

- CALL
- CONTINUE
- EXIT FOR
- FOREACH
- GOTO
- IF
- LET

- LOOP
- RAISE EXCEPTION
- RETURN
- SYSTEM
- TRACE
- WHILE

The “Subset of SPL Statements” syntax diagram for the “IF Statement List” on page 3-30 refers to the SPL statements that are listed in the preceding table.

SQL Statements Not Valid in an IF Statement

The “Subset of SQL Statements” element in the syntax diagram for the “IF Statement List” on page 3-30 refers to all SQL statements, except for the following SQL statements, which are not valid in the IF statement list.

- ALLOCATE DESCRIPTOR
- CLOSE DATABASE
- CONNECT
- CREATE DATABASE
- CREATE PROCEDURE
- DATABASE
- DEALLOCATE DESCRIPTOR
- DESCRIBE
- DISCONNECT
- EXECUTE
- FLUSH
- GET DESCRIPTOR
- GET DIAGNOSTICS
- INFO
- LOAD
- OUTPUT
- PUT
- SET AUTOFREE
- SET CONNECTION
- SET DESCRIPTOR
- UNLOAD
- WHENEVER

You can use a SELECT statement only if you use the INTO TEMP clause to store the result set of the SELECT statement in a temporary table.

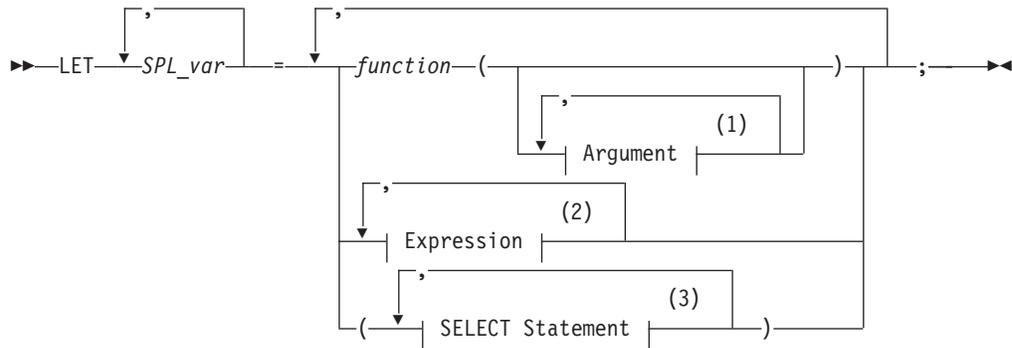
Related Statements

“WHILE ” on page 3-49

LET

Use the LET statement to assign values to variables or to call a user-defined SPL routine and assign the returned value or values to SPL variables.

Syntax



Notes:

- 1 See "Arguments " on page 5-1
- 2 See "Expression" on page 4-40
- 3 See "SELECT" on page 2-528

Element	Description	Restrictions	Syntax
<i>function</i>	SPL function to be invoked	Must exist in the database	"Identifier " on page 5-21
<i>SPL_var</i>	SPL variable to receive a value that the <i>function</i> , expression, or query returns	Must be defined and in scope within the statement block	"Identifier " on page 5-21;

Usage

The LET statement can assign a value returned by an expression, function, or query to an SPL variable. At runtime, the value to be assigned is calculated first. The resulting value is cast to the data type of *SPL_var*, if possible, and the assignment occurs. If conversion is not possible, an error occurs, and the value of the variable remains undefined. (A LET operation that assigns a single value to a single SPL variable is called a *simple assignment*.)

A *compound assignment* assigns multiple expressions to multiple SPL variables. The data types of expressions in the expression list do not need to match the data types of the corresponding variables in the variable list, because the database server automatically converts the data types. (For a detailed discussion of casting, see the *IBM Informix Guide to SQL: Reference*.)

In multiple-assignment operations, the number of variables to the left of the equal (=) sign must match the number of values returned by the functions, expressions, and queries listed on the right of the equal (=) sign. The following example shows several LET statements that assign values to SPL variables:

```

LET a = c + d ;
LET a,b = c,d ;
LET expire_dt = end_dt + 7 UNITS DAY;
LET name = 'Brunhilda';
LET sname = DBSERVERNAME;
LET this_day = TODAY;

```

You cannot use multiple values to the right of the equal (=) sign to operate on other values. For example, the following statement is not valid:

```
LET a,b = (c,d) + (10,15); -- INVALID EXPRESSION
```

Using a SELECT Statement in a LET Statement

The examples in this section use a SELECT statement in a LET statement. You can use a SELECT statement to assign values to one or more variables on the left side of the equals (=) operator, as the following example shows:

```

LET a,b = (SELECT c1,c2 FROM t WHERE id = 1);
LET a,b,c = (SELECT c1,c2 FROM t WHERE id = 1), 15;

```

You cannot use a SELECT statement to make multiple values operate on other values. The following example is invalid:

```
LET a,b = (SELECT c1,c2 FROM t) + (10,15); -- INVALID CODE
```

Because a LET statement is equivalent to a SELECT ... INTO statement, the two statements in the following example have the same results: a=c and b=d:

```

CREATE PROCEDURE proof()
  DEFINE a, b, c, d INT;
  LET a,b = (SELECT c1,c2 FROM t WHERE id = 1);
  SELECT c1, c2 INTO c, d FROM t WHERE id = 1
END PROCEDURE

```

If the SELECT statement returns more than one row, you must enclose the SELECT statement in a FOREACH loop.

For a description of SELECT syntax and usage, see “SELECT” on page 2-528.

Calling a Function in a LET Statement

You can call a user-defined function in a LET statement and assign the returned values to an SPL variable that receives the values that the function returns.

An SPL function can return multiple values (that is, values from multiple columns in the same row) into a list of variable names. In other words, the function can have multiple values in its RETURN statement and the LET statement can have multiple variables to receive the returned values.

When you call the function, you must specify all the necessary arguments to the function unless the arguments of the function have default values. If you specify the name of one of the parameters in the called function with syntax such as **name = 'smith'**, you must name all of the parameters.

An SPL function that selects and returns more than one row must be enclosed in a FOREACH loop.

The following two examples show valid LET statements:

```

LET a, b, c = func1(name = 'grok', age = 17);
LET a, b, c = 7, func2('orange', 'green');

```


Element	Description	Restrictions	Syntax
<i>increment_expr</i>	Positive or negative value by which <i>loop_var</i> is incremented	Must return an integer. Cannot return 0.	"Expression" on page 4-40
<i>label</i>	Name of the loop label for this loop	Must be unique among labels in this SPL routine	"Identifier " on page 5-21
<i>left_expression</i>	Starting expression of a range	Value must match SMALLINT or INT data type of <i>loop_var</i>	"Expression" on page 4-40
<i>loop_var</i>	Variable that determines how many times the loop executes	Must be defined and in scope within this statement block	"Identifier " on page 5-21
<i>right_expression</i>	Ending expression in the range	Same as for <i>left_expression</i>	"Expression" on page 4-40

Usage

The LOOP statement is an iterative statement that resembles the FOR and WHILE statements. Like FOR and WHILE, the LOOP statement can have an optional loop label, It can include the CONTINUE statement to specify another iteration, and the EXIT statement to terminate execution of the loop.

Besides resembling FOR and WHILE in its functionality, the LOOP statement can use the syntax of FOR or WHILE that precedes the statement block. Sections that follow describe several forms of the LOOP statement, including these:

- Simple LOOP statements that iterate a statement loop indefinitely
- FOR LOOP statements, that use FOR statement syntax specify a finite number of iterations
- WHILE LOOP statements, that iterate while a specified condition is true
- Labeled versions of each of these LOOP statements, which can terminate deeply nested loops.

Simple LOOP Statements

The following program fragment illustrates a simple form of the LOOP statement.

```

LOOP
LET i = i + 1;
  IF i = 5 THEN EXIT;
  ELSE
  CONTINUE;
  END IF
END LOOP;

```

In this example the IF statement limits the number of iterations. Here the CONTINUE and EXIT statements omit the optional LOOP keyword, but the END LOOP statement is required at the end of the statement loop. A similar FOR or WHILE keyword would have required the FOR or WHILE keywords, respectively, in the CONTINUE and EXIT statements.

The next example uses a conditional EXIT statement to terminate the loop:

```

LOOP
LET i = i + 1;
  EXIT WHEN i = 4;
END LOOP;

```

No keyword identifying the type of loop statement is required after the EXIT statement, as would be the case for an EXIT statement in a FOR, WHILE, or FOREACH statement. When the *i* = 4 condition becomes true, program control passes from the LOOP statement to whatever statement follows the END LOOP keywords.

FOR LOOP Statements

The FOR LOOP statement uses FOR statement syntax to specify a variable and a range of values that the variable can take. The loop iterates until the specified limit to these values is reached, or until control is transferred outside the loop, as by the unconditional EXIT statement in the following example:

```
FOR i IN (1 TO 5) LOOP
  IF i = 5 THEN EXIT;
  ELSE
    CONTINUE;
END LOOP;
```

In the FOR LOOP statement, the FOR keyword can follow the EXIT or CONTINUE keyword, but the FOR keyword is not required, as it is in an ordinary FOR statement.

The following example replaces the IF statement with a functionally equivalent conditional EXIT statement:

```
FOR i IN (1 TO 5) LOOP
  EXIT WHEN i = 5;
END LOOP;
```

WHILE LOOP Statements

To create a WHILE LOOP statement, loop, you can immediately follow a WHILE *condition* specification with a LOOP statement. The resulting loop terminates after the *condition* becomes false, or when some other statement transfers program control from the loop. In the following WHILE LOOP statement, the condition specifies that the loop terminates after the loop variable *i* has been incremented to the value of 6:

```
WHILE (i < 6) LOOP
  LET i = i + 1;
  IF i = 5 THEN EXIT;
  ELSE
    CONTINUE;
  END IF
END LOOP;
```

As in the FOR LOOP statement, the EXIT and CONTINUE keywords do not need to specify the type of loop statement, but the example would not be affected if EXIT WHILE and CONTINUE WHILE replaced the EXIT and CONTINUE keywords. The END LOOP keywords are required, however, because Dynamic Server treats the WHILE LOOP (and FOR LOOP) statements as LOOP statements, despite their initial FOR and WHILE specifications.

Labeled LOOP Statements

All forms of the LOOP statement, including the FOR LOOP, WHILE LOOP, and simple LOOP statements can have statement labels. You can create a labeled LOOP statement in the following steps:

1. Write a valid LOOP, FOR LOOP, or WHILE LOOP statement.
2. Create a statement label by enclosing an SQL identifier (that is not already the name of a label in the same SPL routine) between angle brackets (<<loop_label>>) immediately before the first line of the LOOP, FOR LOOP, or WHILE LOOP statement.
3. Enter the same SQL identifier, but without angle bracket delimiters, immediately after the END LOOP keywords that terminate the statement, which is now a labeled loop statement.

One advantage of labeled LOOP statements is that they can be referenced in EXIT statements. When the EXIT *label* statement executes, program control passes from the EXIT statement to the statement that follows the specified loop label.

In the following example, a labeled WHILE LOOP loop, whose loop label identifier is **endo**, is part of the statement block of a labeled LOOP statement whose loop label identifier is **voort**. If the conditional EXIT statement EXIT endo WHEN x = 7: detects that its condition is true, program control passes to the LET x = x + 1 statement that follows the END LOOP endostatement. If the conditional EXIT statement EXIT voort WHEN x > 9: detects that its condition is true, program control passes to the LET x = x + 1 statement that follows the END LOOP voortstatement, and the value of x is not incremented by the LET statement,

```
<<voort>>
LOOP
  LET x = x+1;
  <<endo>>
  WHILE ( i < 10 ) LOOP
    LET x = x+1;
    EXIT endo WHEN x = 7;
    EXIT voort WHEN x > 9;
    END LOOP endo;
    LET x = x+1;
  END LOOP voort;
```

uses FOR statement syntax to specify a variable and a range of values that the variable can take. The loop iterates until the specified limit to these values is reached, or until control is transferred outside the loop, as by the unconditional EXIT statement in the following example:

```
FOR i IN (1 TO 5) LOOP
  IF i = 5 THEN EXIT;
  ELSE
    CONTINUE;
  END LOOP;
```

In the FOR LOOP statement, the FOR keyword can follow the EXIT or CONTINUE keyword, but the FOR keyword is not required, as it is in an ordinary FOR statement.

The following example replaces the IF statement with a functionally equivalent conditional EXIT statement:

```
FOR i IN (1 TO 5) LOOP
  EXIT WHEN i = 5;
  END LOOP;
```

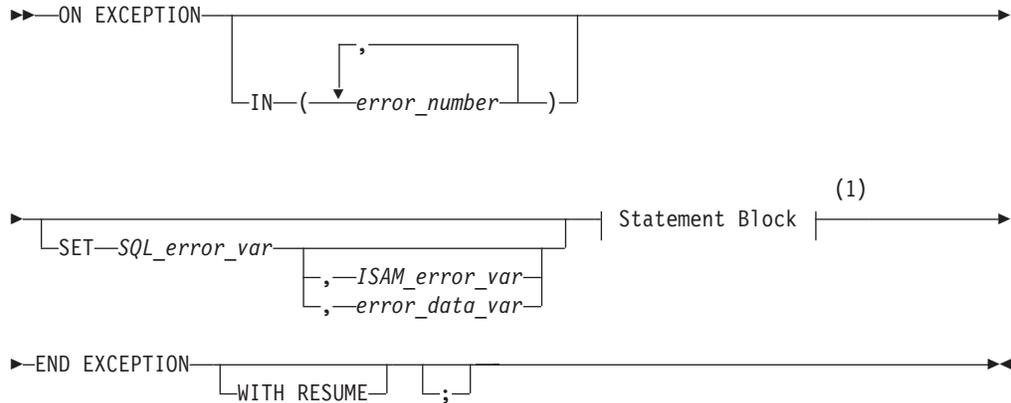
Related Statements

“<< Label >>” on page 3-1, “FOR ” on page 3-19, “WHILE ” on page 3-49

ON EXCEPTION

Use the ON EXCEPTION statement to specify actions to be taken for any error, or for a list of one or more specified errors, during execution of a statement block.

Syntax



Notes:

- 1 See "Statement Block " on page 5-74

Element	Description	Restrictions	Syntax
<i>error_data_var</i>	SPL variable to receive a string returned by an SQL error or by a user-defined exception	Must be a character type to receive the error information. Must be valid in current statement block.	"Identifier " on page 5-21
<i>error_number</i>	SQL error number or a number defined by a RAISE EXCEPTION statement that is to be trapped	Must be of integer type. Must be valid in current statement block.	"Literal Number " on page 4-185
<i>ISAM_error_var</i>	SPL variable that receives the ISAM error number of the exception raised	Same as for <i>error_number</i>	"Identifier " on page 5-21
<i>SQL_error_var</i>	SPL variable that receives the SQL error number of the exception raised	Same as for <i>ISAM_error_var</i>	"Identifier " on page 5-21

Usage

The ON EXCEPTION statement, together with the RAISE EXCEPTION statement, provides an error-trapping and error-recovery mechanism for SPL. ON EXCEPTION can specify the errors that you want to trap as the SPL routine executes, and specifies the action to take if the error occurs within the statement block. ON EXCEPTION can specify an error number list in the IN clause, or can include no IN clause. If the IN clause is omitted, then all errors are trapped.

A statement block can include more than one ON EXCEPTION statement. The exceptions that are trapped can be either system-defined or user-defined.

The scope of an ON EXCEPTION statement is the statement block that follows the ON EXCEPTION statement, and all the statement blocks that are nested within that statement block.

When an exception is trapped, the error status is cleared.

If you specify a variable to receive an ISAM error, but no accompanying ISAM error exists, a zero (0) is assigned to the variable. If you specify a variable to receive the error text, but none exists, the variable stores an empty string.

Placement of the ON EXCEPTION Statement

The ON EXCEPTION statement is a declarative statement, not an executable statement. For this reason, ON EXCEPTION must precede any executable statement and must follow any DEFINE statement within an SPL routine.

The following example shows the correct placement of an ON EXCEPTION statement. Function `add_salesperson()` inserts a set of values into a table. If the table does not exist, it is created, and the values are inserted. The function also returns the total number of rows in the table after the insert occurs:

```
CREATE FUNCTION add_salesperson(last CHAR(15), first CHAR(15))
  RETURNING INT;
  DEFINE x INT;
  ON EXCEPTION IN (-206) -- If no table was found, create one
    CREATE TABLE emp_list
      (lname CHAR(15), fname CHAR(15), tele CHAR(12));
    INSERT INTO emp_list VALUES -- and insert values
      (last, first, '800-555-1234');
  END EXCEPTION WITH RESUME;
  INSERT INTO emp_list VALUES (last, first, '800-555-1234');
  SELECT count(*) INTO x FROM emp_list;
  RETURN x;
END FUNCTION;
```

When an error occurs, the database server searches for the last ON EXCEPTION statement that traps the error code. If the database server finds no pertinent ON EXCEPTION statement, the error code passes back to the calling context (the SPL routine, application, or interactive user), and execution terminates.

In the previous example, the minus sign (-) is required in the IN clause that specifies error -206; most error codes are negative integers.

The next example uses two ON EXCEPTION statements with the same error number so that error code 691 can be trapped in two levels of nesting. All of the DELETE statements except the one that is marked { 6 } are within the scope of the first ON EXCEPTION statement. The DELETE statements that are marked { 1 } and { 2 } are within the scope of the inner ON EXCEPTION statement:

```
CREATE PROCEDURE delete_cust (cnum INT)
  ON EXCEPTION IN (-691) -- children exist
  BEGIN -- Begin-end so no other DELETES get caught in here.
    ON EXCEPTION IN (-691)
      DELETE FROM another_child WHERE num = cnum; { 1 }
      DELETE FROM orders WHERE customer_num = cnum; { 2 }
    END EXCEPTION -- for error -691
      DELETE FROM orders WHERE customer_num = cnum; { 3 }
    END
      DELETE FROM cust_calls WHERE customer_num = cnum; { 4 }
      DELETE FROM customer WHERE customer_num = cnum; { 5 }
    END EXCEPTION
      DELETE FROM customer WHERE customer_num = cnum; { 6 }
  END PROCEDURE
```

Using the IN Clause to Trap Specific Exceptions

An error is trapped if the SQL error code or the ISAM error code matches an exception code in the list of error numbers. The search through the list of errors begins from the left and stops with the first match. You can use a combination of an ON EXCEPTION statement without an IN clause and one or more ON EXCEPTION statements with an IN clause. When an error occurs, the database server searches for the last declaration of the ON EXCEPTION statement that traps the particular error code.

```
CREATE PROCEDURE ex_test()
  DEFINE error_num INT;
  ...
  ON EXCEPTION SET error_num
  -- action C
  END EXCEPTION
  ON EXCEPTION IN (-300)
  -- action B
  END EXCEPTION
  ON EXCEPTION IN (-210, -211, -212) SET error_num
  -- action A
  END EXCEPTION
```

A summary of the sequence of statements in the previous example would be:

1. Test for an error.
2. If error -210, -211, or -212 occurs, take action A.
3. If error -300 occurs, take action B.
4. If any other error occurs, take action C.

Receiving Error Information in the SET Clause

If you use the SET clause, when an exception occurs, the SQL error code and (optionally) the ISAM error code are inserted into the variables that are specified in the SET clause. If you provide an *error_data_var*, any error text that the database server returns is put into the *error_data_var*. Error text includes information such as the offending table or column name.

Forcing Continuation of the Routine

The first example in “Placement of the ON EXCEPTION Statement” on page 3-39 uses the WITH RESUME keyword to indicate that after the statement block in the ON EXCEPTION statement executes, execution is to continue at the LET x = SELECT COUNT(*) FROM emp_list statement, which is the line following the line that raised the error. For this function, the result is that the count of salespeople names occurs even if the error occurred.

Continuing Execution After an Exception Occurs

If you omit the WITH RESUME keywords, the next statement that executes after an exception occurs depends on the placement of the ON EXCEPTION statement, as the following scenarios describe:

- If the ON EXCEPTION statement is inside a statement block with a BEGIN and an END keyword, execution resumes with the first statement (if any) after that BEGIN ... END block. That is, it resumes after the scope of the ON EXCEPTION statement.
- If the ON EXCEPTION statement is inside a loop (FOR, WHILE, FOREACH), the rest of the loop is skipped, and execution resumes with the next iteration of the loop.

- If no statement or block, but only the SPL routine, contains the ON EXCEPTION statement, the routine executes a RETURN statement with no arguments, returning a successful status and no values.

To prevent an infinite loop, if an error occurs during execution of the statement block, then the search for another ON EXCEPTION statement to trap the error does not include the current ON EXCEPTION statement.

Related Statements

“RAISE EXCEPTION ”

RAISE EXCEPTION

Use the RAISE EXCEPTION statement to simulate the generation of an error.

Syntax

```

▶▶ RAISE EXCEPTION SQL_error_var [ , ISAM_error ] [ , error_text ] ;

```

Element	Description	Restrictions	Syntax
<i>error_text</i>	SPL variable or expression that contains error message text for error -746	Must be a character data type and be valid in the statement block	“Identifier ” on page 5-21; “Expression” on page 4-40
<i>ISAM_error</i>	SPL variable or other expression that represents an ISAM error number. The default is 0.	Must return a value in SMALLINT range. You can specify a unary minus sign before error number.	“Expression” on page 4-40
<i>SQL_error</i>	SPL variable or other expression that represents an SQL error number	Same as for <i>ISAM_error</i>	“Expression” on page 4-40

Usage

Use the RAISE EXCEPTION statement to simulate an error or to generate an error with a custom message. An ON EXCEPTION statement can trap the generated error.

If you omit *ISAM_error*, the database server sets the ISAM error code to zero (0) when the exception is raised. If you want to specify *error_text* but not specify a value for *ISAM_error*, specify zero (0) as the value of *ISAM_error*.

The RAISE EXCEPTION statement can raise either system-generated exceptions or user-generated exceptions. For example, the following statement raises the error number -208:

```
RAISE EXCEPTION -208, 0;
```

Here the minus (-) symbol is required after the EXCEPTION keyword for error -208; most error codes are negative integers.

Special Error Number -746

The special error number -746 allows you to produce a customized message. For example, the following statement raises the error number -746 and returns the quoted text:

```
RAISE EXCEPTION -746, 0, 'You broke the rules';
```

In the following example, a negative value for **alpha** raises exception -746 and provides a specific message that describes the problem. The code should contain an ON EXCEPTION statement that traps for an exception of -746.

```
FOREACH SELECT c1 INTO alpha FROM sometable
IF alpha < 0 THEN
RAISE EXCEPTION -746, 0, 'a < 0 found' -- emergency exit
END IF
END FOREACH
```

When the SPL routine executes and the IF condition is met, the database server returns the following error:

```
-746: a < 0 found.
```

For more information about the scope and compatibility of exceptions, see “ON EXCEPTION ” on page 3-38.

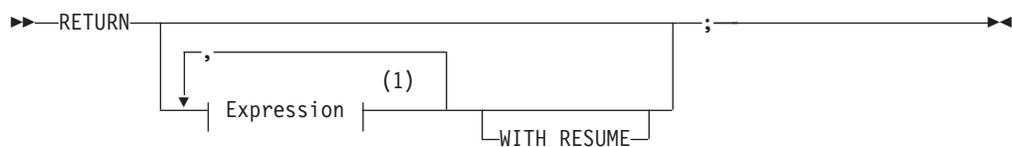
Related Statements

“ON EXCEPTION ” on page 3-38

RETURN

Use the RETURN statement to specify what values (if any) the SPL function returns to the calling context.

Syntax



Notes:

- 1 See “Expression” on page 4-40

Usage

In Dynamic Server, for backward compatibility, you can use the RETURN statement inside a CREATE PROCEDURE statement to create an SPL function. By only using RETURN in CREATE FUNCTION statements, however, you can maintain the convention of using CREATE FUNCTION to define routines that return a value, and CREATE PROCEDURE for other routines.

All RETURN statements in the SPL function must be consistent with the RETURNING clause of the CREATE FUNCTION (or CREATE PROCEDURE) statement that defines the function. Any RETURN list of expressions must match

in cardinality (and be of data types compatible with) the ordered list of data types in the RETURNING clause of the function definition.

Alternatively, however, the RETURN statement can specify no expressions, even if the RETURNING clause lists one or more data types. In this case, a RETURN statement that specifies no expression is equivalent to returning the expected number of NULL values to the calling context. A RETURN statement without any expressions exits only if the SPL function is declared as not returning any values. Otherwise it returns NULL values.

The following SPL function has two valid RETURN statements:

```
CREATE FUNCTION two_returns (stockno INT) RETURNING CHAR (15);
  DEFINE des CHAR(15);
  ON EXCEPTION (-272)      -- if user does not have select privilege
    RETURN;                -- return no values.
  END EXCEPTION;
  SELECT DISTINCT descript INTO des FROM stock
  WHERE stocknum = stockno;
  RETURN des;
END FUNCTION;
```

A program that calls the function in the previous example should test whether no values are returned and act accordingly.

WITH RESUME Keyword

If you use the WITH RESUME keywords, then after the RETURN statement completes execution, the next invocation of the SPL function (upon the next FETCH or FOREACH statement) starts from the statement that follows the RETURN statement. Any function that executes a RETURN WITH RESUME statement must be invoked within a FOREACH loop, or else in the FROM clause of a SELECT statement. If an SPL routine executes a RETURN WITH RESUME statement, a FETCH statement in an Informix ESQL/C application can call the SPL routine.

The following example shows a cursor function that another UDR can call. After the RETURN WITH RESUME statement returns each value to the calling UDR or program, the next line of **series** executes the next time **series** is called. If the variable **backwards** equals zero (0), no value is returned to the calling UDR or program, and execution of **series** stops:

```
CREATE FUNCTION series (limit INT, backwards INT) RETURNING INT;
  DEFINE i INT;
  FOR i IN (1 TO limit)
    RETURN i WITH RESUME;
  END FOR;
  IF backwards = 0 THEN
    RETURN;
  END IF;
  FOR i IN (limit TO 1 STEP -1)
    RETURN i WITH RESUME;
  END FOR;
END FUNCTION; -- series
```

Returning Values from Another Database (IDS)

If an SPL function uses the Return clause to return values from another database of the local Dynamic Server instance, the following data types are supported as the returned data type:

- Built-in data types that are not opaque

- Most of the *built-in opaque data types*, as listed in “Data Types in Cross-Database Transactions” on page 2-536
- DISTINCT of the built-in types that are referenced in the two lines above
- DISTINCT of any DISTINCT data type in this list
- Any opaque user-defined type (UDT) that is cast explicitly to one of the built-in data types in this list.

The definitions of the UDF and of the type hierarchies, casts, DISTINCT types, and UDTs must be exactly the same in each of the participating databases. The same data-type restrictions apply to a value that an external function returns from another database of the local Dynamic Server instance. For more information about data types that are supported in distributed operations across two or more databases of the same database server, see “Data Types in Cross-Database Transactions” on page 2-536.

UDRs can return only the following data types from tables in databases of other database servers:

- Any non-opaque built-in data type
- BOOLEAN
- LVARCHAR
- DISTINCT of non-opaque built-in types
- DISTINCT of BOOLEAN
- DISTINCT of LVARCHAR
- DISTINCT of any DISTINCT type that appears in this list.

UDRs can return these DISTINCT types from databases of other Dynamic Server instances only if the DISTINCT types are cast explicitly to built-in types. The definitions of the DISTINCT data types, their type hierarchies, and their casts must be exactly the same in all databases that participate in the distributed operations. For queries or other DML operations in cross-server UDRs that use the data types in the preceding list as parameters or as returned data types, the UDR must be defined in each participating database, and the participating Dynamic Server instances must support the data type as a returned value in cross-server operations.

For additional information about the data types that Dynamic Server can access in distributed operations, see “Data Types in Distributed Queries” on page 2-535.

External Functions and Iterator Functions (IDS): In an SPL program, you can use a C or Java language external function as an expression in a RETURN statement, provided that the external function is not an iterator function. An *iterator function* is an external function that returns one or more rows of data (and therefore requires a cursor to execute).

SPL iterator functions must include the RETURN WITH RESUME statement. For information about using an iterator function with a virtual table interface in the FROM clause of a query, see “Iterator Functions (IDS)” on page 2-551.

SYSTEM

Use the SYSTEM statement to issue an operating-system command from within an SPL routine.

Syntax

►—SYSTEM— $\left\{ \begin{array}{l} \text{expression} \\ \text{SPL_var} \end{array} \right.$ —;—►

Element	Description	Restrictions	Syntax
<i>expression</i>	Evaluates to a user-executable operating-system command	You cannot specify that the command run in the background	Operating-system dependent
<i>SPL_var</i>	SPL variable containing a command	Must be of a character data type	"Identifier " on page 5-21;

Usage

If the specified *expression* is not a character expression, it is converted to a character expression and passed to the operating system for execution.

The command that SYSTEM specifies cannot run in the background. The database server waits for the operating system to complete execution of the command before it continues to the next statement in the SPL routine. The SPL routine cannot use any returned values from the command.

If the operating-system command fails (that is, returns a nonzero status for the command), an exception is raised that contains the returned operating-system status as the ISAM error code and an appropriate SQL error code.

A rollback does not terminate a system call, so a suspended transaction can wait indefinitely for the call to return. For instructions on recovery from a deadlock during a long transaction rollback, see the *IBM Informix Administrator's Guide*.

The dynamic log feature of Dynamic Server automatically adds log files until the long transaction completes or rolls back successfully.

In DBA- and owner-privileged SPL routines that contain SYSTEM statements, the command runs with the access privileges of the user who executes the routine.

Executing the SYSTEM Statement on UNIX

The SYSTEM statement in the following example of an SPL routine causes the UNIX operating system to send a mail message to the system administrator:

```
CREATE PROCEDURE sensitive_update()
...
LET mailcall = 'mail headhoncho < alert';
-- code to execute if user tries to execute a specified
-- command, then sends email to system administrator
SYSTEM mailcall;
...
END PROCEDURE; -- sensitive_update
```

You can use a double-pipe symbol (||) to concatenate expressions with a SYSTEM statement, as the following example shows:

```
CREATE PROCEDURE sensitive_update2()
DEFINE user1 char(15);
DEFINE user2 char(15);
LET user1 = 'joe';
LET user2 = 'mary';
...
```

```

-- code to execute if user tries to execute a specified
-- command, then sends email to system administrator
SYSTEM 'mail -s violation' ||user1 || ' ' || user2
      || '< violation_file';
...
END PROCEDURE; --sensitive_update2

```

Executing the SYSTEM Statement on Windows

On Windows systems, any SYSTEM statements in an SPL routine are executed only if the current user who is executing the SPL routine has logged on with a password. The database server must have the password and login name of the user in order to execute a command on behalf of that user.

The first SYSTEM statement in the following example of an SPL routine causes Windows to send an error message to a temporary file and to put the message in a system log that is sorted alphabetically. The second SYSTEM statement causes the operating system to delete the temporary file:

```

CREATE PROCEDURE test_proc()
...
SYSTEM 'type errormess101 > %tmp%tmpfile.txt |
      sort >>'
SYSTEM 'del %tmp%tmpfile.txt'
...
END PROCEDURE; --test_proc

```

The expressions that follow the SYSTEM statements in this example contain variables `%tmp%` and `%SystemRoot%` that are defined by Windows.

Setting Environment Variables in SYSTEM Commands

When the operating-system command that SYSTEM specifies is executed, no guarantee exists that any environment variables that the user application sets are passed to the operating system. If you set an environment variable in a SYSTEM command, the setting is only valid during that SYSTEM command.

To avoid this potential problem, the following method is recommended to ensure that any environment variables that the user application requires are carried forward to the operating system.

To Change Environment Settings for an Operating System Command

1. Create a shell script (on UNIX systems) or a batch file (on Windows platforms) that sets up the desired environment and then executes the operating system command.
2. Use the SYSTEM command to execute the shell script or batch file.

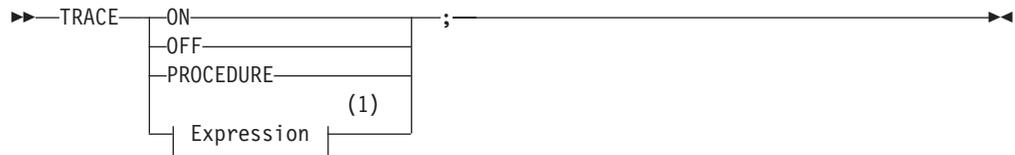
This solution has an additional advantage: if you subsequently need to change the environment, you can modify the shell script or the batch file without needing to recompile the SPL routine.

For information about operating system commands that set environment variables, see the *IBM Informix Guide to SQL: Reference*.

TRACE

Use the TRACE statement to control the generation of debugging output.

Syntax



Notes:

- 1 See “Expression” on page 4-40

Usage

The TRACE statement generates output that is sent to the file that the SET DEBUG FILE TO statement specifies. Tracing writes to the debug file the current values of the following program objects:

- SPL variables
- Routine arguments
- Return values
- SQL error codes
- ISAM error codes

The output of each executed TRACE statement appears on a separate line.

If you use the TRACE statement before you specify a DEBUG file to contain the output, an error is generated.

Any routine that the SPL routine calls inherits the trace state. That is, a called routine (on the same database server) assumes the same trace state (ON, OFF, or PROCEDURE) as the calling routine. The called routine can set its own trace state, but that state is not passed back to the calling routine.

A routine that is executed on a remote database server does not inherit the trace state.

TRACE ON

If you specify the keyword ON, all statements are traced. The values of variables (in expressions or otherwise) are printed before they are used. To turn tracing ON implies tracing both routine calls and statements in the body of the routine.

TRACE OFF

If you specify the keyword OFF, all tracing is turned off.

TRACE PROCEDURE

If you specify the keyword PROCEDURE, only the routine calls and return values, but not the body of the routine, are traced.

The TRACE statement supports no ROUTINE or FUNCTION keywords. Use the TRACE PROCEDURE keywords when the SPL routine that you trace is a function.

Displaying Expressions

You can use the TRACE statement with a quoted string or an expression to display values or comments in the output file. If the expression is not a literal expression, the expression is evaluated before it is written to the output file.

You can use the TRACE statement with an expression even if you used a TRACE OFF statement earlier in a routine. You must first, however, use the SET DEBUG statement to establish a trace output file.

The next example uses a TRACE statement with an expression after using a TRACE OFF statement. The example uses UNIX file naming conventions:

```
CREATE PROCEDURE tracing ()
  DEFINE i INT;
BEGIN
  ON EXCEPTION IN (1)
  END EXCEPTION; -- do nothing
  SET DEBUG FILE TO '/tmp/foo.trace';
  TRACE OFF;
  TRACE 'Forloop starts';
  FOR i IN (1 TO 1000)
  BEGIN
    TRACE 'FOREACH starts';
    FOREACH SELECT...INTO a FROM t
      IF <some condition> THEN
        RAISE EXCEPTION 1 -- emergency exit
      END IF
    END FOREACH -- return some value
  END
END FOR -- do something
END;
END PROCEDURE
```

Example Showing Different Forms of TRACE

The following example shows several different forms of the TRACE statement. The example uses Windows file naming conventions:

```
CREATE PROCEDURE testproc()
  DEFINE i INT;
  SET DEBUG FILE TO 'C:\tmp\test.trace';
  TRACE OFF;
  TRACE 'Entering foo';
  TRACE PROCEDURE;
  LET i = test2();

  TRACE ON;
  LET i = i + 1;

  TRACE OFF;
  TRACE 'i+1 = ' || i+1;
  TRACE 'Exiting testproc';

  SET DEBUG FILE TO 'C:\tmp\test2.trace';

END PROCEDURE
```

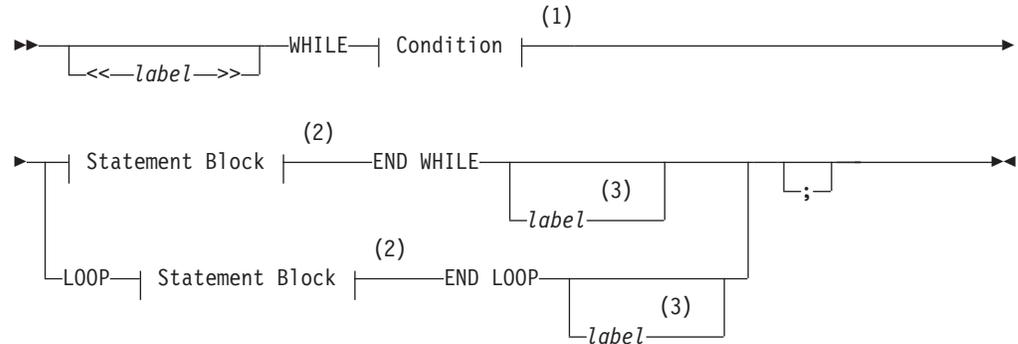
Looking at the Traced Output

To see the traced output, use a text editor or similar utility to display or read the contents of the file.

WHILE

Use the WHILE statement to establish a loop with variable end conditions.

Syntax



Notes:

- 1 See "Condition " on page 4-5
- 2 See "Statement Block " on page 5-74
- 3 Valid only if <<label>> precedes the first WHILE keyword

Element	Description	Restrictions	Syntax
<i>label</i>	Name of the loop label for this loop	Must be unique among labels in this SPL routine	"Identifier " on page 5-21

Usage

The *condition* is evaluated before the *statement block* first runs and before each subsequent iteration. Iterations continue as long as the *condition* remains true. The loop terminates when the *condition* evaluates to not true.

If any expression within the *condition* evaluates to NULL, the *condition* becomes not true unless you are explicitly testing for NULL with the IS NULL operator.

If an expression within the *condition* has an UNKNOWN value because it references uninitialized SPL variables, an immediate error results. In this case, the loop terminates, raising an exception.

Example of WHILE Loops in an SPL Routine

The following example illustrates the use of WHILE loops in an SPL routine. In the SPL procedure, **simp_while**, the first WHILE loop executes a DELETE statement. The second WHILE loop executes an INSERT statement and increments the value of an SPL variable.

```
CREATE PROCEDURE simp_while()  
  DEFINE i INT;  
  WHILE EXISTS (SELECT fname FROM customer  
    WHERE customer_num > 400)  
    DELETE FROM customer WHERE id_2 = 2;  
  END WHILE;  
  LET i = 1;
```

```

WHILE i < 10
  INSERT INTO tab_2 VALUES (i);
  LET i = i + 1;
END WHILE;
END PROCEDURE;

```

Labeled WHILE Loops

To create a labeled WHILE loop, you can declare a loop label before the initial WHILE keyword, and repeat the label after the END WHILE keywords, as in the two WHILE loops of the following example:

```

CREATE PROCEDURE ex_cont_ex()
  DEFINE i,s,j, INT;
  <<while_jlab>>
  WHILE j < 20
    IF j > 10 THEN
      CONTINUE WHILE;
    END IF
    LET i,s = j,0;
    <<while_slab>>
    WHILE i > 0
      LET i = i -1;
      IF i = 5 THEN
        EXIT while_jlab;
      END IF
    END WHILE while_slab
  END WHILE while_jlab
END PROCEDURE;

```

Here the EXIT while_jlab statement has the same effect that the EXIT or EXIT FOR keywords would have, terminating both the outer WHILE loop and the routine. In this example, the statement that includes the EXIT while_jlab statement has the same effect that EXIT while_jlab WHEN i = 5 would have.

You can also label a LOOP statement that begins with a loop <<label>> specification that immediately precedes the initial WHILE keyword and *condition*. In this type of loop, the CONTINUE LOOP, EXIT LOOP, and END LOOP keywords replace the CONTINUE WHILE, EXIT WHILE, and END WHILE keywords. Both the LOOP and WHILE keywords are optional after the CONTINUE and EXIT keywords, but the END LOOP keywords are required in SPL loop statements that include the LOOP keyword.

You can use similar syntax to create an unlabeled loop that omits the <<label>> declaration that immediately precedes the WHILE *condition* specification. In this case, you must also omit the undelimited loop label identifier that follows the END LOOP keywords. See the LOOP statement for a description and examples of these forms of labeled and unlabeled loop statements that enable you to combine WHILE statement syntax, with its condition-based number of loop iterations, with the "loop forever" syntax of the LOOP statement.

Related Statements

"<< Label >>" on page 3-1, "CONTINUE" on page 3-7, "EXIT " on page 3-17, "LOOP " on page 3-34

Chapter 4. Data Types and Expressions

In This Chapter

This chapter describes the data types and expressions that Dynamic Server and Extended Parallel Server support. These fundamental syntax segments can appear in data definition language (DDL) and data manipulation language (DML) statements, and in other types of SQL statements. Some SPL statements can also specify data types or expressions. You can use these features of a relational or object-relational database in various contexts, such as to define the schema of a table, to specify the signature and arguments of a routine, or to represent or calculate specific data values.

Scope of Segment Descriptions

The description of each segment includes the following information:

- A brief introduction that explains the effect of the segment
- A syntax diagram that shows how to enter the segment correctly
- A table that explains the terms in the syntax diagram for which you must substitute names, values, or other specific information
- Rules of usage, typically including examples that illustrate these rules

If a segment consists of multiple parts, the segment description provides similar information about each part. Some descriptions conclude with references to related information in this document and in other documents.

Use of Segment Descriptions

The syntax diagram within each segment description is not a stand-alone diagram. Rather, it is a subdiagram of the syntax of the SQL statements (in Chapter 2, “SQL Statements,” on page 2-1) or of SPL statements (in Chapter 3, “SPL Statements,” on page 3-1) that can include the segment.

SQL or SPL syntax descriptions can refer to segment descriptions in two ways:

- A *subdiagram reference* in a syntax diagram can list a segment name and the page in this document where the segment description begins.
- The **Syntax** column of the table that immediately follows a syntax diagram can list a segment name and the page where the segment description begins.

If the syntax diagram for a statement includes a reference to a segment, turn to that segment description to see the complete syntax for the segment.

For example, if you want to write a CREATE VIEW statement that includes a *database* and *database server* qualifiers of the *view* name, first look up the syntax diagram for the “CREATE VIEW ” on page 2-274. The table beneath that diagram refers to the Database Object Name segment for the syntax of *view*. Then use the Database Object Name segment syntax to enter a valid CREATE VIEW statement that also specifies the *database* and *database server* name for the view (and for Extended Parallel Server, the *coserver* identifier). In the following example, the CREATE VIEW statement defines a view called **name_only** in the **sales** database on the **boston** database server:

```
CREATE VIEW sales@boston:name_only AS
  SELECT customer_num, fname, lname FROM customer;
```

Besides the Data Types and Expressions syntax segments that this chapter documents, Chapter 5, “Other Syntax Segments,” on page 5-1 provides additional syntax segments that are referenced in the syntax diagrams of this document.

Segments in This Chapter

This chapter describes the following Data Type and Expression segments:

- Data Type
- DATETIME Field Qualifier
- INTERVAL Field Qualifier
- Expression
- Aggregate Expression
- AVG, COUNT, MAX, MIN, SUM, RANGE, STDEV, VARIANCE, *and* User-Defined Aggregates
- Arithmetic Expressions
- Binary (+, -, *, /) Operators, Operator Functions, *and* Unary (+, -) Operators
- Cast Expressions
- CAST function, Cast (::) Operator
- Collection Subquery
- Column Expressions
- Column Name, ROWID, *and* Substring ([...]) Operator
- Concatenation (||) Operator
- Conditional Expressions
- Condition
- Comparison Condition: AND, OR, NOT, BETWEEN, IS NULL, LIKE, MATCHES, *and* Relational Operators
- Condition with Subquery: IN, EXISTS, ALL, ANY, *and* SOME Operators
- Boolean UDF
- CASE Expressions
- NVL Function
- DECODE Function
- Constant Expressions: CURRENT, TODAY, DBSERVERNAME, SITENAME, UNITS, *and* USER Operators
- Literal Value
- Literal Collection
- Literal DATETIME
- Literal INTERVAL
- Literal Number
- Literal Row
- Quoted String
- Constructor Expressions
- Collection Constructor
- ROW Constructor Function Expressions

- Algebraic Functions: ABS, MOD, POW, POWER ROOT, ROUND, SQRT, and TRUNC Functions
- CARDINALITY Function
- DBINFO Function
- Encryption and Decryption Functions: DECRYPT_BINARY, DECRYPT_CHAR, ENCRYPT_AES, ENCRYPT_TDES, and GETHINT Functions
- Exponential and Logarithmic Functions: EXP, LOGN, and LOG10 Functions
- HEX Function
- IFX_ALLOW_NEWLINE Function
- Length Functions: CHARACTER_LENGTH, CHAR_LENGTH, LENGTH, and OCTET_LENGTH Functions
- Sequence Operators
- CURRVAL, NEXTVAL
- Smart Large Object Functions: FILETOBLOB, FILETOCLOB, LOCOPY, and LOTOFILE Functions
- String-Manipulation Functions: LPAD, RPAD, TRIM, REPLACE, SUBSTR, SUBSTRING, INITCAP, LOWER, and UPPER Functions
- Time Functions: DATE, DAY, EXTEND, MDY, MONTH, TO_CHAR, TO_DATE, WEEKDAY, and YEAR Functions
- Trigger-Type Boolean Operators: DELETING, INSERTING, SELECTING, and UPDATING
- Trigonometric Functions: ACOS, ASIN, ATAN, ATAN2, COS, SIN, and TAN Functions
- User-Defined Functions Statement-Local
- Variable Expressions

You can also use host variables or SPL variables as expressions. For an alphabetic list of expressions with page references, see “List of Expressions” on page 4-42.

Collection Subquery

You can use a Collection Subquery to create a MULTISET collection from the results of a subquery. Only Dynamic Server supports this syntax, which is an extension to the ANSI/ISO standard for SQL.

Syntax

Collection Subquery:

```
(1)
|-----MULTISET-----(|-----subquery-----|)-----|
|-----SELECT ITEM-----singleton_select-----|
```

Notes:

- 1 Informix extension

Element	Description	Restrictions	Syntax
<i>singleton_select</i>	Subquery returning exactly one row	Subquery cannot repeat the SELECT keyword, nor include the ORDER BY clause	“SELECT” on page 2-528

Element	Description	Restrictions	Syntax
<i>subquery</i>	Embedded query	Cannot contain the ORDER BY clause	"SELECT" on page 2-528

Usage

The MULTISSET and SELECT ITEM keywords have the following significance:

- MULTISSET specifies a collection of elements that can contain duplicate values, but that has no specific order of elements.
- SELECT ITEM supports only one expression in the projection list. You cannot repeat the SELECT keyword in the singleton subquery.

You can use a collection subquery in the following contexts:

- The Projection clause and WHERE clause of the SELECT statement
- The VALUES clause of the INSERT statement
- The SET clause of the UPDATE statement
- Wherever you can use a collection expression (that is, any expression that evaluates to a single collection)
- As an argument passed to a user-defined routine

The following restrictions apply to a collection subquery:

- The Projection clause cannot contain duplicate column (field) names.
- It cannot contain aliases for table names. (But it can use aliases for column (field) names, as in some of the examples that follow.)
- It is read-only.
- It cannot be opened twice.
- It cannot contain NULL values.
- It cannot contain syntax that attempts to seek within the subquery.

A collection subquery returns a multiset of unnamed ROW data types. The fields of this ROW type are elements in the projection list of the subquery. Examples that follow access the tables and the ROW types that these statements define:

```
CREATE ROW TYPE rt1 (a INT);
CREATE ROW TYPE rt2 (x int, y rt1);
CREATE TABLE tab1 (col1 rt1, col2 rt2);
CREATE TABLE tab2 OF TYPE rt1;
CREATE TABLE tab3 (a ROW(x INT));
```

The following examples of collection subqueries return the MULTISSET collections that are listed to the right of the subquery.

Collection Subquery	Resulting Collections
MULTISSET (SELECT * FROM tab1)...	MULTISSET(ROW(col1 rt1, col2 rt2))
MULTISSET (SELECT col2.y FROM tab1)...	MULTISSET(ROW(y rt1))
MULTISSET (SELECT * FROM tab2)...	MULTISSET(ROW(a int))
MULTISSET(SELECT p FROM tab2 p)...	MULTISSET(ROW(p rt1))
MULTISSET (SELECT * FROM tab3)...	MULTISSET(ROW(a ROW(x int)))

The following is another collection subquery:

```
SELECT f(MULTISET(SELECT * FROM tab1 WHERE tab1.x = t.y))
FROM t WHERE t.name = 'john doe';
```

The following collection subquery includes the UNION operator:

```
SELECT f(MULTISET(SELECT id FROM tab1
UNION
SELECT id FROM tab2 WHERE tab2.id2 = tab3.id3)) FROM tab3;
```

Table Expressions in the FROM Clause

Dynamic Server supports ANSI/ISO standard syntax for table expressions in the FROM clause of SELECT queries and subqueries as a substitute for the Informix-extension collection subquery syntax. The keywords TABLE and MULTISET are required in version 10.00 and in earlier releases. These extensions to the ANSI/ISO standard for SQL are supported but are no longer required for collection subqueries in the FROM clause of SELECT statements.

The following two queries return the same result set, but only the second query complies with the ANSI/ISO standard:

```
SELECT * FROM TABLE(MULTISET(SELECT col1 FROM tab1 WHERE col1 = 100))
AS vtab(c1),
(SELECT col1 FROM tab1 WHERE col1 = 10) AS vtab1(vc1) ORDER BY c1;

SELECT * FROM (SELECT col1 FROM tab1 WHERE col1 = 100) AS vtab(c1),
(SELECT col1 FROM tab1 WHERE col1 = 10) AS vtab1(vc1)
ORDER BY c1;
```

The same SELECT statement can combine instances of both the Informix-extension and ANSI/ISO syntax for collection subqueries:

```
SELECT * FROM (select col1 FROM tab1 WHERE col1 = 100) AS vtab(c1),
TABLE(MULTISET(SELECT col1 FROM tab1 WHERE col1 = 10)) AS vtab1(vc1)
ORDER BY c1;
```

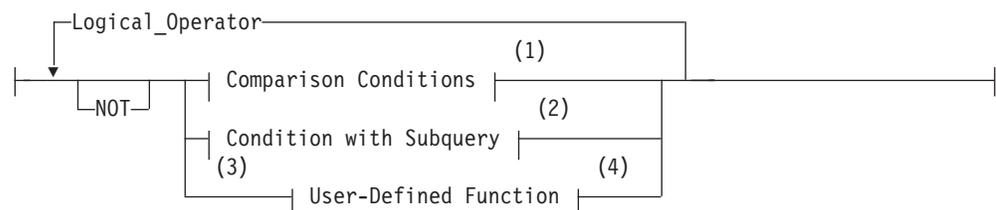
The collection subquery must be delimited by parentheses in both formats, but the outer set of parentheses (()) that immediately follows the TABLE keyword and encloses the MULTISET collection subquery specification is an extension to the ANSI/ISO syntax. This ANSI/ISO syntax is valid only in the FROM clause of the SELECT statement. You cannot omit these keywords and parentheses from a collection subquery specification in any other context.

Condition

Use a condition to test whether data meets certain qualifications. Use this segment wherever you see a reference to a condition in a syntax diagram.

Syntax

Condition:



Notes:

- 1 See "Comparison Conditions (Boolean Expressions)"
- 2 See "Condition with Subquery" on page 4-15
- 3 Dynamic Server only
- 4 See "User-Defined Functions" on page 4-161

Element	Description	Restrictions	Syntax
<i>Logical _Operator</i>	Combines two conditions	Valid options are OR (= <i>logical union</i>) or AND (= <i>logical intersection</i>)	"Conditions with AND or OR" on page 4-20

Usage

A *condition* is a search criterion, optionally connected by the logical operators AND or OR. Conditions can be classified into the following categories:

- Comparison conditions (also called filters or Boolean expressions)
- Conditions with a subquery
- User-defined functions (Dynamic Server only)

A condition can contain an aggregate function only if it is used in the HAVING clause of a SELECT statement or in the HAVING clause of a subquery.

No aggregate function can appear in a condition in the WHERE clause of a DELETE, SELECT, or UPDATE statement unless both of the following are TRUE:

- Aggregate is on a correlated column originating from a parent query.
- The WHERE clause appears in a subquery within a HAVING clause.

In Dynamic Server, user-defined functions are not valid as conditions in the following contexts:

- In the HAVING clause of a SELECT statement
- In the definition of a check constraint

SPL routines are not valid as conditions in the following contexts:

- In the definition of a check constraint
- In the ON clause of a SELECT statement
- In the WHERE clause of a DELETE, SELECT, or UPDATE statement

External routines are not valid as conditions in the following contexts:

- In the definition of a check constraint
- In the ON clause of a SELECT statement
- In the WHERE clause of a DELETE, SELECT, or UPDATE statement
- In the WHEN clause of CREATE TRIGGER
- In the IF, CASE, or WHILE statements of SPL

Comparison Conditions (Boolean Expressions)

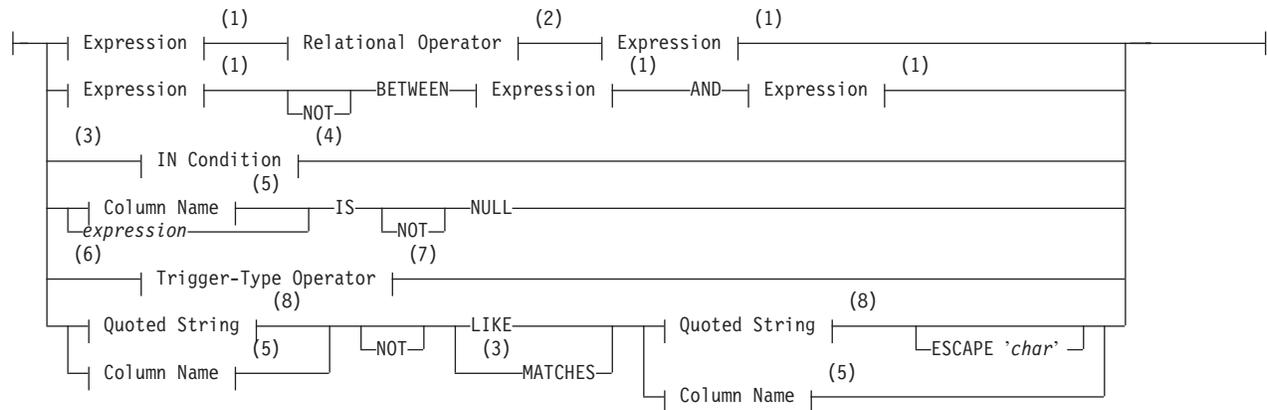
Six kinds of comparison conditions exist:

- Relational operator
- [NOT] BETWEEN ... AND
- IN

- IS [NOT] NULL
- Trigger-type operators
- LIKE and MATCHES

Comparison conditions are often called *Boolean expressions* because they return a TRUE or FALSE result. Their syntax is summarized in this diagram and explained in the sections that follow.

Comparison Conditions:



Notes:

- 1 See "Expression" on page 4-40
- 2 See "Relational Operator" on page 4-193
- 3 Informix extension
- 4 See "IN Condition" on page 4-10
- 5 See "Column Name" on page 4-8
- 6 SPL trigger routines only
- 7 See "Trigger-Type Boolean Operator" on page 4-12
- 8 See "Quoted String" on page 4-189

Element	Description	Restrictions	Syntax
<i>char</i>	An ASCII character to be the nondefault escape character in the quoted string. Single (') and double (") quotation marks are not valid as <i>char</i> .	See "ESCAPE with LIKE" on page 4-15 and "ESCAPE with MATCHES" on page 4-15	"Quoted String" on page 4-189
<i>expression</i>	An SQL expression that returns a single value	None	"Expression" on page 4-40

The following sections describe the different types of comparison conditions:

- "Relational-Operator Condition" on page 4-9
- "BETWEEN Condition" on page 4-9
- "IN Condition" on page 4-10
- "IS NULL and IS NOT NULL Conditions" on page 4-12
- "LIKE and MATCHES Condition" on page 4-13.

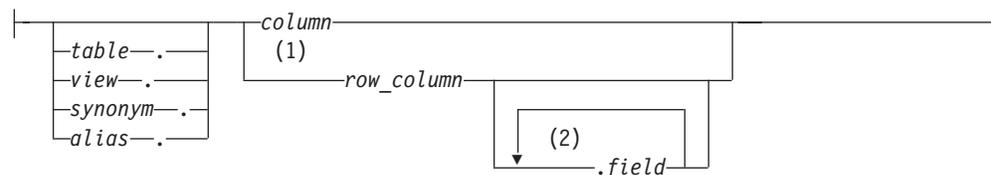
For a discussion of comparison conditions in the context of the SELECT statement, see “Using a Condition in the WHERE Clause” on page 2-561.

Warning: A literal DATE or DATETIME value in a comparison condition should specify 4 digits for the year. When you specify a 4-digit year, the **DBCENTURY** environment variable has no effect on the result. When you specify a 2-digit year, **DBCENTURY** can affect how the database server interprets the comparison condition, which might not work as you intended. For more information about **DBCENTURY**, see the *IBM Informix Guide to SQL: Reference*.

Column Name

The Column Name segment can be an element in comparison conditions. The name of a column (or of one or more fields within a column of a ROW data type) is not the subject of the comparison, but the database server uses this SQL identifier to access the data value in the specified column or field of a row in a database table or view.

Column Name:



Notes:

- 1 Dynamic Server only
- 2 Repeat no more than three times

Element	Description	Restrictions	Syntax
<i>alias</i>	Temporary alternative name for table or view	Must be defined in the FROM clause of the SELECT statement	“Identifier ” on page 5-21
<i>column</i>	Name of a column	Must exist in the specified table	“Identifier ” on page 5-21
<i>field</i>	A field to compare in a ROW type column	Must be a component of <i>row-column name</i> or <i>field name</i> (for nested rows)	“Identifier ” on page 5-21
<i>row_column</i>	A column of type ROW	Must be an existing named ROW type or unnamed ROW type	“Identifier ” on page 5-21
<i>synonym, table, view</i>	Name of a synonym, table, or view	The <i>synonym</i> and the table or view to which it points must exist in the database	“Identifier ” on page 5-21

For more information on the meaning of the *column* name in these conditions, see the “IS NULL and IS NOT NULL Conditions” on page 4-12 and the “LIKE and MATCHES Condition” on page 4-13.

Quotation Marks in Conditions

When you compare a column expression with a constant expression in any comparison condition, observe the following rules:

- If the column has a numeric data type, do not enclose the constant expression between quotation marks.

- If the column has a character data type, enclose the constant expression between quotation marks.
- If the column has a time data type, enclose the constant expression between quotation marks.

Otherwise, you might get unexpected results.

The following example shows the correct use of quotation marks in comparison conditions. Here the **ship_instruct** column has a character data type, the **order_date** column has a date data type, and the **ship_weight** column has a numeric data type.

```
SELECT * FROM orders
  WHERE ship_instruct = 'express'
     AND order_date > '05/01/98'
     AND ship_weight < 30;
```

Relational-Operator Condition

The following examples show some relational-operator conditions:

```
city[1,3] = 'San'
```

```
o.order_date > '6/12/98'
```

```
WEEKDAY(paid_date) = WEEKDAY(CURRENT- (31 UNITS DAY))
```

```
YEAR(ship_date) < YEAR (TODAY)
```

```
quantity <= 3
```

```
customer_num <> 105
```

```
customer_num != 105
```

If an expression within the *condition* has an UNKNOWN value because it references an uninitialized variable, the database server raises an exception.

If any expression within the *condition* evaluates to NULL, the *condition* cannot be true, unless you are explicitly testing for NULL by using the IS NULL operator. For example, if the **paid_date** column has a NULL value, then neither of the following statements can retrieve that row:

```
SELECT customer_num, order_date FROM orders
  WHERE paid_date = '';
```

```
SELECT customer_num, order_date FROM orders
  WHERE NOT (paid_date != '');
```

The IS NULL operator tests for a NULL value, as the next example shows. The IS NULL operator is described in “IS NULL and IS NOT NULL Conditions” on page 4-12.

```
SELECT customer_num, order_date FROM orders
  WHERE paid_date IS NULL;
```

BETWEEN Condition

For a BETWEEN test to be TRUE, the value of the expression on the left of the BETWEEN keyword must be in the inclusive range of the values of the two expressions on the right of the BETWEEN keyword.

NULL values do not satisfy the condition, and you cannot use NULL for either expression that defines the range.

The following examples show some BETWEEN conditions:

```
order_date BETWEEN '6/1/97' and '9/7/97'
```

```
zipcode NOT BETWEEN '94100' and '94199'
```

```
EXTEND(call_dtime, DAY TO DAY) BETWEEN  
(CURRENT - INTERVAL(7) DAY TO DAY) AND CURRENT
```

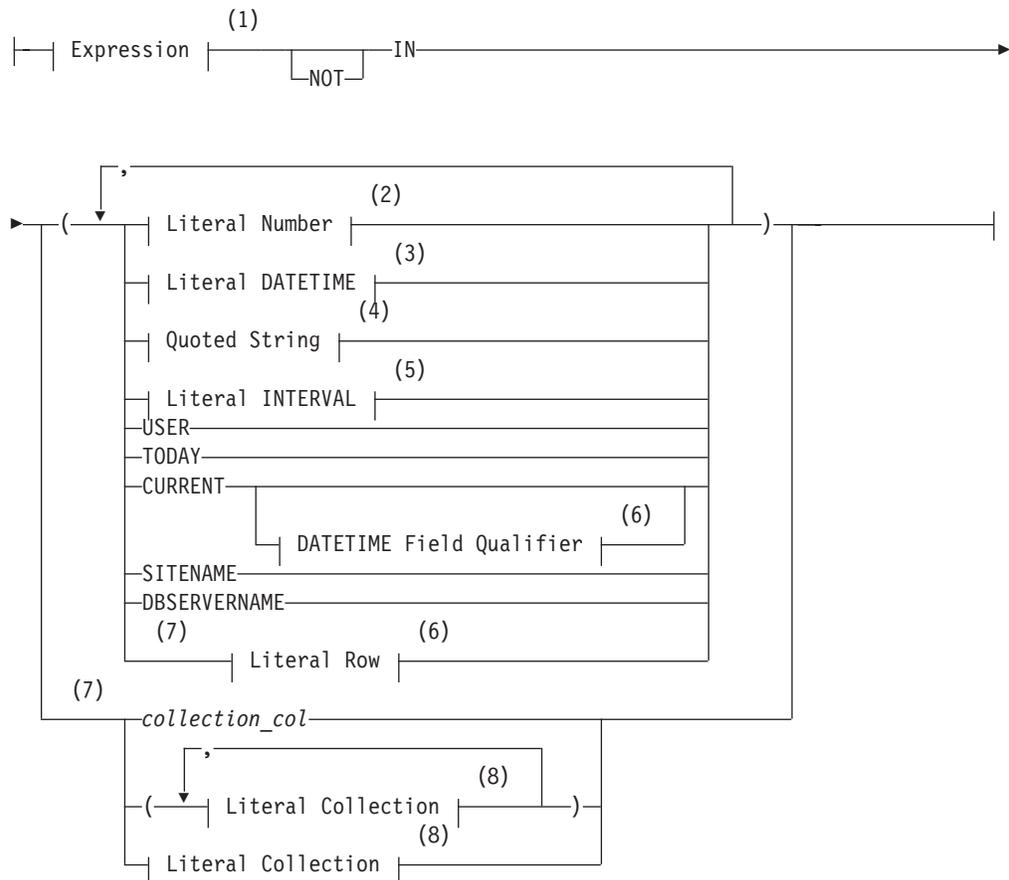
```
lead_time BETWEEN INTERVAL (1) DAY TO DAY  
AND INTERVAL (4) DAY TO DAY
```

```
unit_price BETWEEN loprice AND hiprice
```

IN Condition

The IN condition is satisfied when the expression to the left of the keyword IN is included in the list of items.

IN Condition:



Notes:

- 1 See "Expression" on page 4-40
- 2 See "Literal Number " on page 4-185
- 3 See "Literal DATETIME " on page 4-181

- 4 See "Quoted String" on page 4-189
- 5 See "Literal INTERVAL" on page 4-183
- 6 See " DATETIME Field Qualifier" on page 4-38
- 7 Dynamic Server only
- 8 See "Literal Row" on page 4-186

Element	Description	Restrictions	Syntax
<i>collection_col</i>	Name of a collection column that is used in an IN condition	The column must exist in the specified table	"Identifier " on page 5-21

If you specify the NOT operator, the IN condition is TRUE when the expression is not in the list of items. NULL values do not satisfy the IN condition.

The following examples show some IN conditions:

```
WHERE state IN ('CA', 'WA', 'OR')
WHERE manu_code IN ('HRO', 'HSK')
WHERE user_id NOT IN (USER)
WHERE order_date NOT IN (TODAY)
```

In Informix ESQL/C, the built-in TODAY function is evaluated at execution time. The built-in CURRENT function is evaluated when a cursor opens or when the query executes, if it is a singleton SELECT statement.

The built-in USER function is case sensitive; for example, it interprets **minnie** and **Minnie** as different values.

Using the IN Operator with Collection Data Types (IDS)

You can use the IN operator to determine if an element is contained in a collection. The collection can be a simple or nested collection. (In a *nested* collection type, the element type of the collection is also a collection type.) When you use IN to search for an element of a collection, the expression to the left or right of the IN keyword cannot contain a BYTE or TEXT data type.

Suppose you create the following table that contains two collection columns:

```
CREATE TABLE tab_coll
(
  set_num SET(INT NOT NULL),
  list_name LIST(SET(CHAR(10) NOT NULL) NOT NULL)
);
```

The following statement fragments show how you might use the IN operator for search conditions on the collection columns of the **tab_coll** table:

```
WHERE 5 IN set_num
WHERE 5.0::INT IN set_num
WHERE "5" NOT IN set_num
WHERE set_num IN ("SET{1,2,3}", "SET{7,8,9}")
WHERE "SET{'john', 'sally', 'bill'}" IN list_name
WHERE list_name IN ("LIST{""SET{'bill','usha'}""",
  ""SET{'ann' 'moshi'}""",
  "LIST{""SET{'bob','ramesh'}""",
  ""SET{'bomani' 'ann'}""")
```

In general, when you use the IN operator on a collection data type, the database server checks whether the value on the left of the IN operator is an element in the set of values on the right of the IN operator.

IS NULL and IS NOT NULL Conditions

The IS NULL condition is satisfied if the term that immediately precedes the IS keyword specifies one of the following undefined values:

- The name of a *column* that contains a null value.
- An *expression* that evaluates to null.

Conversely, if you use the IS NOT NULL operator, the condition is satisfied when the *column* contains a value that is not null, or when the *expression* that immediately precedes the IS NOT NULL keywords does not evaluate to null.

Suppose that you wish to perform an arithmetic computation on a column that can contain NULL values. You can create a table, insert values into the table, and then perform a query that uses a generic CASE expression that converts null values to 0 for the purpose of arithmetic calculations:

```
CREATE TABLE employee (emp_id INT, savings_in_401k INT, total_salary INT);

INSERT INTO employee VALUES(1, 5000, 40000);
INSERT INTO employee VALUES(2, 0, 40000);
INSERT INTO employee VALUES(3, NULL, 100000);

SELECT emp_id, savings_in_401k AS employer_match FROM employee WHERE
    CASE WHEN(savings_in_401k IS NULL) THEN 0
         ELSE savings_in_401k END * 0.06 > 0;
```

This example shows that by using IS NULL in the CASE expression, you can provide a value for the entries that otherwise are not computable because null is not a valid numeric value.

The IS NULL condition is satisfied if the column contains a null value or if the expression cannot be evaluated because it contains one or more null values. If you use the IS NOT NULL operator, the condition is satisfied when the operand is column value that is not null, or an expression that does not evaluate to null.

Trigger-Type Boolean Operator

The trigger-type Boolean operators of Dynamic Server can test at runtime whether the currently executing triggered action was triggered by the specified type of DML event. These operators take no operands.

Trigger-Type Boolean Operator:



These operators return TRUE ('t') if the triggering event of the currently executing trigger is the DML operation corresponding to the name of the operator, and they return FALSE ('f') otherwise. These operators are valid in IF statements, in CASE expressions, and in other contexts within an SPL trigger routines where a Boolean condition is valid.

For example, in the following statement fragment, the LET statement in the first THEN clause is executed only if the currently executing trigger was activated by

an INSERT event, and the LET statement in the second THEN clause is executed only if the trigger was activated by a DELETE event:

```
IF (INSERTING = 't') THEN
    LET square = NEW.X * NEW.X
ELIF (DELETING = 't') THEN
    LET square = 0
```

The SELECTING, DELETING, INSERTING, and UPDATING operators are valid only in trigger UDRs that are invoked in the FOR EACH ROW triggered action of a trigger on a table, or (for the DELETING, INSERTING, and UPDATING operators) of an INSTEAD OF trigger on a view. An error is issued if you attempt to use a trigger-type Boolean operator in any other context.

LIKE and MATCHES Condition

A LIKE or MATCHES condition tests for matching character strings. The condition is TRUE, or satisfied, when either of the following tests is TRUE:

- The value of the column on the left matches the pattern that the quoted string specifies. You can use wildcard characters in the string. NULL values do not satisfy the condition.
- The value of the column on the left matches the pattern that the column on the right specifies. The value of the column on the right serves as the matching pattern in the condition.

You can use the single quotation mark (') only with the quoted string to match a literal single quotation mark; you cannot use the ESCAPE clause. You can use the single quotation mark character as the escape character in matching any other pattern if you write it as '' ''.

Important: You cannot specify a ROW-type column in a LIKE or MATCHES condition. A ROW-type column is a column that is declared as a named or unnamed ROW type.

NOT Operator

The NOT operator makes the search condition successful when the column on the left has a value that is not NULL and that does not match the pattern that the quoted string specifies.

For example, the following conditions exclude all rows that begin with the characters Baxter in the **Iname** column:

```
WHERE Iname NOT LIKE 'Baxter%'
WHERE Iname NOT MATCHES 'Baxter*'
```

LIKE Operator

The LIKE operator supports these wildcard characters in the quoted string.

Wildcard

Effect

- | | |
|---|---|
| % | Matches zero or more characters |
| _ | Matches any single character |
| \ | Removes the special significance of the next character (to match a literal % or _ or \ by specifying \% or _ or \\) |

Using the backslash (\) symbol as the default escape character is an Informix extension to the ANSI/ISO-standard for SQL.

In an ANSI-compliant database, you can only use an escape character to escape a percent sign (%), an underscore (_), or the escape character itself.

The following condition tests for the string tennis, alone or in a longer string, such as tennis ball or table tennis paddle:

```
WHERE description LIKE '%tennis%'
```

The next example tests for **description** rows containing an underscore. Here backslash (\) is necessary because underscore (_) is a wildcard character.

```
WHERE description LIKE '%\_%'
```

The LIKE operator has an associated operator function called **like()**. You can define a **like()** function to handle your own user-defined data types. See also *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

MATCHES Operator

The MATCHES operator supports wildcard characters in the quoted string.

Wildcard

Effect

- * Matches any string of zero or more characters
- ? Matches any single character
- [. . .] Matches any of the enclosed characters, including ranges, as in [a-z]. Characters within the brackets cannot be escaped.
- ^ As first character within the brackets, matches any character that is not listed. Thus, [^abc] matches any character except a, b, or c.
- \ Removes the special significance of the next character (to match a literal \ or any other wildcard by specifying \\ or * or \? and so forth)

The following condition tests for the string tennis, alone or within a longer string, such as tennis ball or table tennis paddle:

```
WHERE description MATCHES '*tennis*'
```

The following condition is TRUE for the names Frank and frank:

```
WHERE fname MATCHES '[Ff]rank'
```

The following condition is TRUE for any name that begins with either F or f:

```
WHERE fname MATCHES '[Ff]*'
```

The next condition is TRUE for any name that ends with the letters a, b, c, or d:

```
WHERE fname MATCHES '*[a-d]'
```

MATCHES has an associated **matches()** operator function. You can define a **matches()** function for your own user-defined data types. For more information, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

If **DB_LOCALE** or **SET COLLATION** specifies a nondefault locale supporting a localized collation, and you specify a range for the MATCHES operator using bracket ([. . .]) symbols, the database server uses the localized collating order, instead of code-set order, to interpret the range and to compare values that have CHAR, CHARACTER VARYING, LVARCHAR, NCHAR, NVARCHAR, and VARCHAR data types.

This behavior is an exception to the usual rule that only NCHAR and NVARCHAR data types can be compared in a localized collating order. For more information on the GLS aspects of conditions that include the MATCHES or LIKE operators, see the *IBM Informix GLS User's Guide*.

ESCAPE with LIKE

The ESCAPE clause can specify a nondefault escape character. For example, if you specify z in the ESCAPE clause, then a quoted string operand that included z_ is interpreted as including a literal underscore (_) character, rather than _ as a wildcard. Similarly, z% is interpreted as a literal percent (%) sign, rather than % as a wildcard. Finally, the characters zz in a string would be interpreted as single literal z. The following statement retrieves rows from the **customer** table in which the **company** column includes a literal underscore character:

```
SELECT * FROM customer WHERE company LIKE '%z_%' ESCAPE 'z';
```

You can also use a host variable that contains a single character. The next statement uses a host variable to specify an escape character:

```
EXEC SQL BEGIN DECLARE SECTION;
    char escp='z';
    char fname[20];
EXEC SQL END DECLARE SECTION;
EXEC SQL select fname from customer
    into :fname where company like '%z_%' escape :escp;
```

ESCAPE with MATCHES

The ESCAPE clause can specify a nondefault escape character. Use this as you would the backslash to include a question mark (?), an asterisk (*), a caret (^), or a left ([) or right (]) bracket as a literal character within the quoted string, to prevent them from being interpreted as special characters. If you choose to use z as the escape character, the characters z? in a string stand for a literal question mark (?). Similarly, the characters z* stand for a literal asterisk (*). Finally, the characters zz in the string stand for the single character z.

The following example retrieves rows from the **customer** table in which the value of the **company** column includes the question mark (?):

```
SELECT * FROM customer WHERE company MATCHES '*z?*' ESCAPE 'z';
```

Stand-Alone Condition

A stand-alone condition can be any expression that is not explicitly listed in the syntax for the comparison condition. Such an expression is valid as a condition only if it returns a BOOLEAN value. For example, the following example returns a value of the BOOLEAN data type:

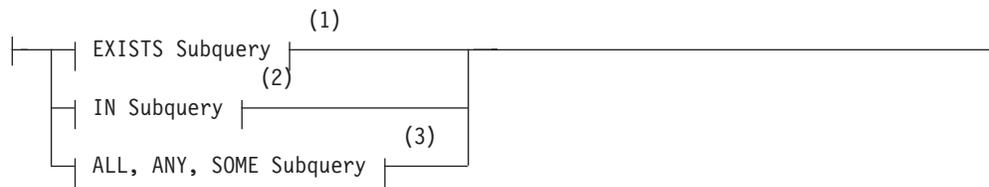
```
funcname(x)
```

Condition with Subquery

Include a SELECT statement within a condition specifies a condition with subquery. You can use a subquery in a SELECT, INSERT, DELETE, or UPDATE statement to perform tasks like the following:

- Compare an expression to the result of the query.
- Determine if an expression is included in the results of the query.
- Ask whether the query selects any rows.

Condition with Subquery:



Notes:

- 1 See “ EXISTS Subquery” on page 4-18
- 2 See “ IN Subquery” on page 4-17
- 3 See “ ALL, ANY, and SOME Subqueries” on page 4-18

The subquery can depend on the current row that the outer SELECT statement is evaluating; in this case, the subquery is called a *correlated subquery*. (For a discussion of correlated subqueries and their impact on performance, see the *IBM Informix Guide to SQL: Tutorial*.)

The following sections describe subquery conditions and their syntax.

- For a discussion of types of subquery conditions in the context of the SELECT statement, see “Using a Condition in the WHERE Clause” on page 2-561.
- For a discussion of types of subquery conditions in the context of the INSERT statement, see “Subset of SELECT Statement” on page 2-441..
- For a discussion of types of subquery conditions in the context of the DELETE statement, see “Subqueries in the WHERE Clause of DELETE” on page 2-306..
- For a discussion of types of subquery conditions in the context of the UPDATE statement, see “Subqueries in the WHERE Clause of UPDATE” on page 2-700.

A subquery can return a single value, no value, or a set of values, depending on its context. If a subquery returns a value, it must select only a single column. If the subquery simply checks whether a row (or rows) exists, it can select any number of rows and columns.

A subquery cannot reference BYTE or TEXT columns, nor can it contain an ORDER BY clause. A subquery that specifies a table expression in the FROM clause, however, can include the ORDER BY clause.

A subquery and its outer DML statement operate on the same table object if the FROM clause of the subquery specifies the same table or view that the outer statement references in one of these clauses:

- in the FROM clause of the DELETE or SELECT statement
- in the INTO clause of the INSERT statement
- in the Table Options or Collection Derived Table specification of the UPDATE statement.

Subqueries that return more than one row and that operate on the same table or view as the enclosing DML statement are valid only in the WHERE clause of the DELETE or UPDATE statement. Even in this context, such subqueries return error -360 unless all of the following conditions are satisfied:

- The subquery does not reference any column name in its FROM list that is in a table not specified in the projection list

- The subquery is specified using the Condition with Subquery syntax.
- Any SPL routines within the subquery cannot reference the table that is being modified.

The following program fragment includes examples of conditions with subqueries in UPDATE and DELETE statements:

```
CREATE TABLE t1 ( a INT, a1 INT)
CREATE TABLE t2 ( b INT, b1 INT) ;
.
.
UPDATE t1 SET a = a + 10 WHERE EXISTS
  (SELECT a FROM t1 WHERE a > 1);
UPDATE t1 SET a = a + 10 WHERE a IN
  (SELECT a FROM t1, t2 WHERE a > b
   AND a IN
     (SELECT a FROM t1 WHERE a > 50 ) );
DELETE FROM t1 WHERE EXISTS
  (SELECT a FROM t1);
```

For more information about subqueries in the DELETE statement, see “Subqueries in the WHERE Clause of DELETE” on page 2-306.

For more information about subqueries in the UPDATE statement, see “Subqueries in the WHERE Clause of UPDATE” on page 2-700.

IN Subquery

An IN subquery condition is TRUE if the value of the expression matches one or more of the values from the subquery. (The subquery must return only one row, but it can return more than one column.) The keyword IN is equivalent to the =ANY specification. The keywords NOT IN are equivalent to the !=ALL specification. See the “ ALL, ANY, and SOME Subqueries” on page 4-18.

IN Subquery:

```

(1)
|-----| Expression |-----| IN (---subquery---)
|-----|
|-----| NOT |-----|

```

Notes:

- 1 See “Expression” on page 4-40

Element	Description	Restrictions	Syntax
<i>subquery</i>	Embedded query	Cannot contain the FIRST nor the ORDER BY clause	“SELECT” on page 2-528

The following example of an IN subquery finds the order numbers for orders that do not include baseball gloves (stock_num = 1):

```
WHERE order_num NOT IN
  (SELECT order_num FROM items WHERE stock_num = 1)
```

Because the IN subquery tests for the presence of rows, duplicate rows in the subquery results do not affect the results of the main query. Therefore, the UNIQUE or DISTINCT keyword in the subquery has no effect on the query results, although not testing duplicates can improve query performance.

EXISTS Subquery

An EXISTS subquery condition evaluates to TRUE if the subquery returns a row. With an EXISTS subquery, one or more columns can be returned. The subquery always contains a reference to a column of the table in the main query. If you use an aggregate function in an EXISTS subquery that includes no HAVING clause, at least one row is always returned.

EXISTS Subquery:



Element	Description	Restrictions	Syntax
<i>subquery</i>	Embedded query	Cannot contain the FIRST nor the ORDER BY clause	"SELECT" on page 2-528

The following example of a SELECT statement with an EXISTS subquery returns the stock number and manufacturer code for every item that has never been ordered (and is therefore not listed in the **items** table). You can appropriately use an EXISTS subquery in this SELECT statement because you use the subquery to test both **stock_num** and **manu_code** in **items**.

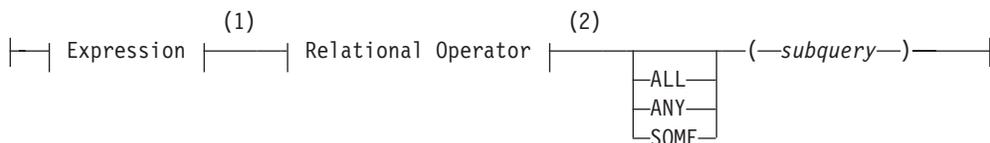
```
SELECT stock_num, manu_code FROM stock
WHERE NOT EXISTS (SELECT stock_num, manu_code FROM items
WHERE stock.stock_num = items.stock_num AND
stock.manu_code = items.manu_code);
```

The preceding example works equally well if you use SELECT * in the subquery in place of the column names, because the existence of the entire row is tested; specific column values are not tested.

ALL, ANY, and SOME Subqueries

Use the ALL, ANY, and SOME keywords to specify what makes the condition TRUE or FALSE. A search condition that is TRUE when the ANY keyword is used might not be TRUE when the ALL keyword is used, and vice versa.

ALL, ANY, SOME Subquery:



Notes:

- 1 See "Expression" on page 4-40
- 2 See "Relational Operator" on page 4-193

Element	Description	Restrictions	Syntax
<i>subquery</i>	Embedded query	Cannot contain the FIRST or the ORDER BY clause	"SELECT" on page 2-528

Using the ALL Keyword: The ALL keyword specifies that the search condition is TRUE if the comparison is TRUE for every value that the subquery returns. If the subquery returns no value, the condition is TRUE.

In the following example, the first condition tests whether each **total_price** is greater than the total price of every item in order number 1023. The second condition uses the **MAX** aggregate function to produce the same results.

```
total_price > ALL (SELECT total_price FROM items
                  WHERE order_num = 1023)
```

```
total_price > (SELECT MAX(total_price) FROM items
              WHERE order_num = 1023)
```

Using the NOT keyword with an ALL subquery tests whether an expression is not TRUE for at least one element that the subquery returns. For example, the following condition is TRUE when the expression **total_price** is not greater than all the selected values. That is, it is TRUE when **total_price** is not greater than the highest total price in order number 1023.

```
NOT total_price > ALL (SELECT total_price FROM items
                     WHERE order_num = 1023)
```

Using the ANY or SOME Keywords: The ANY keyword denotes that the search condition is TRUE if the comparison is TRUE for at least one of the values that is returned. If the subquery returns no value, the search condition is FALSE. The SOME keyword is a synonym for ANY.

The following conditions are TRUE when the total price is greater than the total price of at least one of the items in order number 1023. The first condition uses the ANY keyword; the second uses the MIN aggregate function:

```
total_price > ANY (SELECT total_price FROM items
                  WHERE order_num = 1023)
```

```
total_price > (SELECT MIN(total_price) FROM items
              WHERE order_num = 1023)
```

Using the NOT keyword with an ANY subquery tests whether an expression is not TRUE for all elements that the subquery returns. For example, the following condition is TRUE when the expression **total_price** is not greater than any selected value. That is, it is TRUE when **total_price** is greater than none of the total prices in order number 1023.

```
NOT total_price > ANY (SELECT total_price FROM items
                     WHERE order_num = 1023)
```

Omitting the ANY, ALL, or SOME Keywords: You can omit the keywords ANY, ALL, or SOME in a subquery if you know that the subquery will return exactly one value. If you omit the ANY, ALL, or SOME keywords, and the subquery returns more than one value, you receive an error. The subquery in the following example returns only one row because it uses an aggregate function:

```
SELECT order_num FROM items
   WHERE stock_num = 9 AND quantity =
         (SELECT MAX(quantity) FROM items WHERE stock_num = 9);
```

NOT Operator

If you preface a condition with the keyword NOT, the test is TRUE only if the condition that NOT qualifies is FALSE. If the condition that NOT qualifies has a NULL or an UNKNOWN value, the NOT operator has no effect.

The following truth table shows the effect of NOT with 3-valued Boolean operands. Here T represents a TRUE condition, F represents a FALSE condition, and a question mark (?) represents an UNKNOWN condition. (An UNKNOWN value can occur when an operand is NULL).

NOT	
T	F
?	?
F	T

The left column shows the value of the operand of the NOT operator, and the right column shows the returned value after NOT is applied to the operand.

Conditions with AND or OR

You can combine simple conditions with the logical operators AND or OR to form complex conditions. The following SELECT statements contain examples of complex conditions in their WHERE clauses:

```
SELECT customer_num, order_date FROM orders
  WHERE paid_date > '1/1/97' OR paid_date IS NULL;
SELECT order_num, total_price FROM items
  WHERE total_price > 200.00 AND manu_code LIKE 'H';
SELECT lname, customer_num FROM customer
  WHERE zipcode BETWEEN '93500' AND '95700'
  OR state NOT IN ('CA', 'WA', 'OR');
```

The following truth tables show the effect of the AND and OR operators. The letter T represents a TRUE condition, F represents a FALSE condition, and the question mark (?) represents an UNKNOWN value. An UNKNOWN value can occur when part of an expression that uses a logical operator is NULL.

OR	T	?	F	AND	T	?	F
T	T	T	T	T	T	?	F
?	T	?	?	?	?	?	F
F	T	?	F	F	F	F	F

The marginal values at the left represent the first operand, and values in the top row represent the second operand. Values within each 3x3 matrix show the returned value after the operator is applied to operands of those values.

If the Boolean expression evaluates to UNKNOWN, the condition is not satisfied.

Consider the following example within a WHERE clause:

```
WHERE ship_charge/ship_weight < 5
  AND order_num = 1023
```

The row where **order_num** = 1023 is a row where **ship_weight** is NULL. Because **ship_weight** is NULL, **ship_charge/ship_weight** is also NULL; therefore, the truth value of **ship_charge/ship_weight** < 5 is UNKNOWN. Because **order_num** = 1023 is TRUE, the AND table states that the truth value of the entire condition is UNKNOWN. Consequently, that row is not chosen. If the condition used an OR in place of the AND, the condition would be TRUE.

Related Information

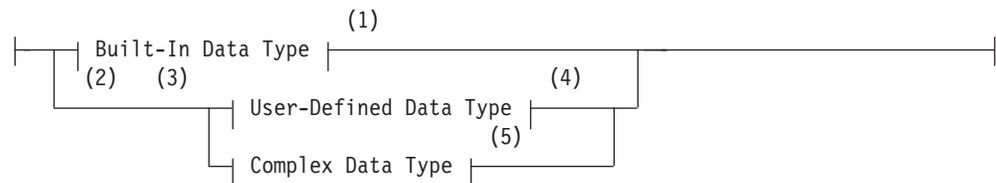
For discussions of comparison conditions in the SELECT statement and of conditions with a subquery, see the *IBM Informix Guide to SQL: Tutorial*. For information on the GLS aspects of conditions, see the *IBM Informix GLS User's Guide*.

Data Type

The Data Type segment specifies the data type of a column, of a component of a collection, of a field within a ROW type, of a routine parameter, or of a value returned by an expression or by a cast function. Use this segment whenever you see a reference to a data type in a syntax diagram.

Syntax

Data Type:



Notes:

- 1 See "Built-In Data Types"
- 2 Informix extension
- 3 Dynamic Server only
- 4 See "User-Defined Data Type (IDS)" on page 4-32
- 5 See "Complex Data Type (IDS)" on page 4-34

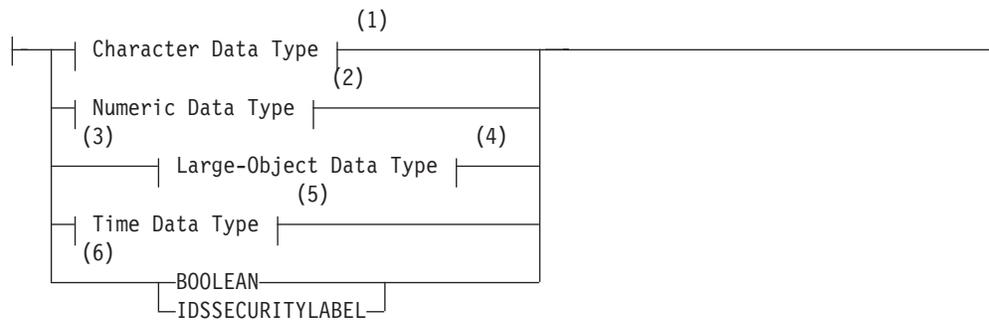
Usage

Sections that follow summarize these data types. For more information, see the chapter about data types in the *IBM Informix Guide to SQL: Reference*.

Built-In Data Types

Built-in data types are data types that are defined by the database server.

Built-In Data Type:



Notes:

- 1 See “Character Data Types” on page 4-23
- 2 See “Numeric Data Types” on page 4-26
- 3 Informix extension
- 4 See “ Large-Object Data Types” on page 4-29
- 5 See “Time Data Types” on page 4-31
- 6 Dynamic Server only

These are “built into the database server” in the sense that the information and support functions required to interpret and transfer these data types is part of the database server software, which supports *character*, *numeric*, *large-object*, and *time* categories of built-in data types. These are described in sections that follow.

BOOLEAN and Other Built-In Opaque Data Types (IDS)

Dynamic Server also supports the BOOLEAN data type, which is a *built-in opaque data type* that can store true, false, or NULL values. The symbol t represents a literal BOOLEAN true value, and f represents a literal BOOLEAN false value.

BOOLEAN and LVARCHAR are the only built-in opaque data types that can be returned by cross-server distributed queries or by other cross-server distributed DML operations. Column values of other built-in opaque data types cannot be retrieved by a distributed query (nor modified by INSERT, DELETE, or UPDATE operations on a remote database) unless all of the tables that the DML operation accesses are in databases of the local Dynamic Server instance.

Similarly, in UDRs that perform distributed operations on databases of other Dynamic Server instances, BOOLEAN and LVARCHAR are the only built-in opaque types that are valid as a parameter or as the returned data type of the UDR, which must be defined in all participating databases.

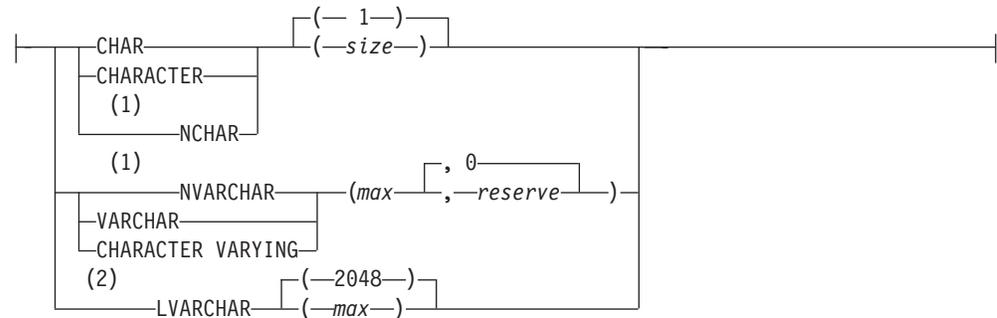
Besides the BOOLEAN type, other built-in opaque data types of Dynamic Server include BLOB, CLOB, LVARCHAR, IFX_LO_SPEC, IFX_LO_STAT, INDEXKEYARRAY, POINTER, RTNPARAMTYPES, SELFUNCARGS, STAT, CLIENTBINVAL, and XID data types. These twelve built-in opaque types are supported in the local database and in distributed operations across databases of the same server instance. The first three of these types are discussed in subsequent sections of this chapter.

Dynamic Server also supports the built-in opaque data types LOLIST, IMPEX, IMPEXBIN, and SENDRECV. These types cannot, however, be accessed in a remote database by DML operations, nor returned from a remote database by a UDR, because these data types do not have the required support functions. For more information about the data types that Dynamic Server supports in distributed transactions, see “Data Types in Distributed Queries” on page 2-535.

Character Data Types

The character data types enable the database server to store text strings.

Character Data Type:



Notes:

- 1 Localized Collation
- 2 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>max</i>	Maximum size in bytes. For VARCHAR, this is required. LVARCHAR default is 2048	VARCHAR: Integer; $1 \leq max \leq 255$ (or $1 \leq max \leq 254$, if indexed) LVARCHAR: $1 \leq max \leq 32,739$	“Literal Number ” on page 4-185
<i>reserve</i>	Bytes reserved. Default is 0.	Integer; $0 \leq reserve \leq max$	“Literal Number ” on page 4-185
<i>size</i>	Size in bytes. Default is 1.	Integer; $1 \leq size \leq 32,767$	“Literal Number ” on page 4-185

The database server issues an error if the data type declaration includes empty parentheses, such as LVARCHAR(). To declare a CHAR or LVARCHAR data type of the default length, simply omit any (*size*) or (*max*) specification. The CREATE TABLE statement of Dynamic Server accepts VARCHAR and NVARCHAR column declarations that have no (*max*) nor (*max, reserve*) specifications, using (1, 0) as the (*max, reserve*) defaults for the column.

The following table summarizes the built-in character data types.

Data Type	Description
CHAR	Stores single-byte or multibyte text strings of fixed length (up to 32,767 bytes); supports code-set order in collation of text data. Default size is 1 byte.
CHARACTER	Synonym for CHAR
CHARACTER VARYING	ANSI-compliant synonym for VARCHAR

Data Type	Description
LVARCHAR (IDS)	Stores single-byte or multibyte text strings of varying length (up to 32,739 bytes). The size of other columns in the same table can further reduce this upper limit. Default size is 2,048 bytes.
NCHAR	Stores single-byte or multibyte text strings of fixed length (up to 32,767 bytes); supports localized collation of text data.
NVARCHAR	Stores single-byte or multibyte text strings of varying length (up to 255 bytes); supports localized collation of text data.
VARCHAR	Stores single-byte or multibyte text strings of varying length (up to 255 bytes); supports code-set order collation of text data.

Single-Byte and Multi-Byte Characters and Locales

All built-in character data types can support single- and multibyte characters in the code set that the `DB_LOCALE` setting specifies. Locales for most European and Middle Eastern languages support only single-byte code sets, but some East Asian locales (and the `UTF-8` Unicode locale) support multibyte characters.

The `TEXT` and `CLOB` data types also support single-byte or multibyte character data, but most built-in functions for manipulating character strings do not support `TEXT` nor `CLOB` data. For more information, see “Large-Object Data Types” on page 4-29.

Fixed- and Varying-Length Character Data Types

The database server supports storage of fixed-length and varying-length character data. A *fixed-length* column requires the defined number of bytes regardless of the actual size of the data. The `CHAR` data type is of fixed-length. For example, a `CHAR(25)` column requires 25 bytes of storage for all values, so the string "This is a text string" uses 25 bytes of storage.

A *varying-length* column size can be the number of bytes occupied by its data. `NVARCHAR`, `VARCHAR`, and (for Dynamic Server only) the `LVARCHAR` data types are varying-length character data types. For example, a `VARCHAR(25)` column reserves up to 25 bytes of storage for the column value, but the character string "This is a text string" uses only 21 bytes of the reserved 25 bytes. The `VARCHAR` data type can store up to 255 bytes of data. For information about the `IFX_PAD_VARCHAR` environment variable, whose setting controls how the database server sends and receives `VARCHAR` and `NVARCHAR` data values, see *IBM Informix Guide to SQL: Reference*.

Because of the maximum row size limit of 32,767 bytes, a single table cannot be created with more than approximately 195 varying-length or `ROW` type columns.

LVARCHAR Data Type

The `LVARCHAR` type of Dynamic Server can store up to 32,739 bytes of text, but if you specify no *size* in an `LVARCHAR` data type declaration, the default length is 2,048 bytes. `LVARCHAR` is a built-in opaque data type. Unlike most of the built-in opaque types, `LVARCHAR` column values can be accessed in a database of a non-local Dynamic Server instance in a distributed query or other DML operations, and `LVARCHAR` can be the data type of a parameter or of a returned value of a UDR that accesses data outside the local database.

Dynamic Server uses the LVARCHAR data type in cross-server I/O operations on opaque data types. In this context, the maximum size of the LVARCHAR data value is limited only by the operating system.

Light scans are not supported on tables that include LVARCHAR columns.

NCHAR and NVARCHAR Data Types

The character data types CHAR, LVARCHAR, and VARCHAR support code-set order collation of data. The database server collates text data in columns of these types by the order that their characters are defined in the code set of the locale, which the **DB_LOCALE** environment variable specifies. For information on how the settings (or the default values) of the **DB_LOCALE**, **CLIENT_LOCALE**, and **SERVER_LOCALE** environment variables determine which locale is used for collation, see the *IBM Informix GLS User's Guide*.

Some locales specify an order of collation that is not identical to the code-set order. To support any locale-specific order of collation, you can use the NCHAR and NVARCHAR data types. The NCHAR data type is a fixed-length character data type that supports localized collation. The NVARCHAR data type is a varying-length character data type that can store up to 255 bytes of text data and supports localized collation.

For NCHAR or NVARCHAR values, the SET COLLATION statement of SQL can override the localized collation order of the current session, by specifying another locale. Indexes on NCHAR or NVARCHAR columns sort values according to the localized collation order that was in effect when the index was created, if that is different from the current collation order.

Dynamic Server does not support light scans on tables that have NVARCHAR columns.

In Dynamic Server, if you specify no parameters in CREATE TABLE or ALTER TABLE statements that declare VARCHAR or NVARCHAR columns, then the new columns default to a *max* size of 1 byte and a *reserve* size of zero.

IDSSECURITYLABEL Data Type

The IDSSECURITYLABEL type of Dynamic Server stores a security label in a table that is protected by a security policy. Only a user who holds the DBSECADM role can create, alter, or drop a column of this data type. This is a built-in DISTINCT OF VARCHAR(128) data type, but it is not classified as a character data type because its use is restricted to label-based access control. A table that has a security policy can have no more than one IDSSECURITYLABEL column, and a table associated with no security policy can have none.

The DBSECADM can use the GRANT statement to associate a specific security label with a user, and the REVOKE statement can cancel a security label that a user holds. For a given security policy, a user can have no more than one label that supports both read and write access, or no more than one label for write access and no more than one label for read access. For data protected by a security policy, but for which the user has been granted discretionary access privileges, the database server determines whether a specific user can access the data by comparing the security label of the data with the security label of the user, while also taking into consideration any exemptions to the security policy rules that the user holds.

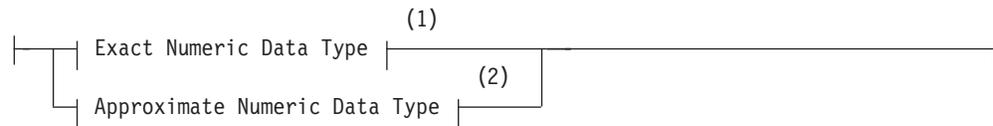
For information on how to specify an IDSSECURITYLABEL value, see “ Security Label Support Functions (IDS)” on page 4-119.

For a discussion of security policies, security components, security labels, and other concepts of label-based access control (LBAC), see the Informix Dynamic Server Security Guide.

Numeric Data Types

Numeric data types enable the database server to store numbers such as integers and real numbers in a column.

Numeric Data Type:



Notes:

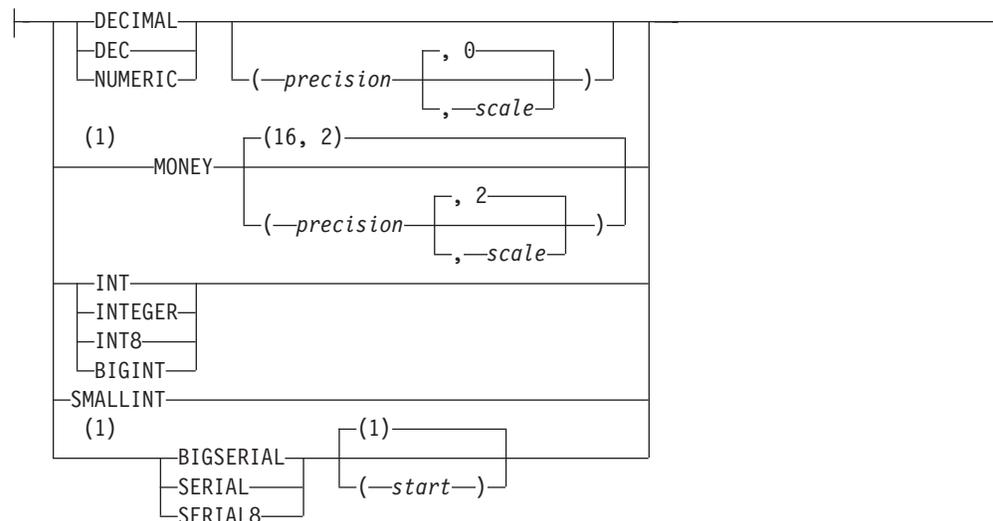
- 1 See “Exact Numeric Data Types”
- 2 See “Approximate Numeric Data Types” on page 4-28

The values of numbers are stored either as *exact numeric* data types or as *approximate numeric* data types.

Exact Numeric Data Types

An *exact numeric* data type stores numbers of a specified precision and scale.

Exact Numeric Data Type:



Notes:

- 1 Informix extension

Element	Description	Restrictions	Syntax
<i>precision</i>	Significant digits	Must be an integer; $1 \leq \textit{precision} \leq 32$	"Literal Number " on page 4-185
<i>scale</i>	Digits in fractional part	Must be an integer; $1 \leq \textit{scale} \leq \textit{precision}$	"Literal Number " on page 4-185
<i>start</i>	Integer starting value	For SERIAL: $1 \leq \textit{start} \leq 2,147,483,647$; For BIGSERIAL and SERIAL8: $1 \leq \textit{start} \leq 9,223,372,036,854,775,807$	"Literal Number " on page 4-185

The *precision* of a data type is the number of digits that the data type stores. The *scale* is the number of digits to the right of the decimal separator.

The following table summarizes the exact numeric data types available.

Data Type	Description
DEC(<i>p,s</i>)	Synonym for DECIMAL(<i>p,s</i>)
DECIMAL(<i>p,s</i>)	Stores fixed-point decimal values of real numbers, with up to 30 significant digits in the fractional part, or up to 32 significant digits to the left of the decimal point.
INT	Synonym for INTEGER
INTEGER	Stores a 4-byte integer value. These values can be in the range from $-(2^{31}-1)$ to $2^{31}-1$ (from -2,147,483,647 to 2,147,483,647).
BIGINT and INT8	Stores an 8-byte integer value. These values can be in the range from $-(2^{63}-1)$ to $2^{63}-1$ (the range -9,223,372,036,854,775,807 to 9,223,372,036,854,775,807). BIGINT has storage and processing advantages over INT8.
MONEY(<i>p,s</i>)	Stores fixed-point currency values. These values have same internal data format as a fixed-point DECIMAL(<i>p,s</i>) value.
NUMERIC(<i>p,s</i>)	ANSI-compliant synonym for DECIMAL(<i>p,s</i>)
SERIAL	Stores a 4-byte positive integer that the database server generates. Values can range from 1 to $2^{31}-1$ (that is, from 1 to 2,147,483,647).
BIGSERIAL and SERIAL8	Stores an 8-byte positive integer value that the database server generates. Values can range from 1 to $2^{63}-1$ (that is, from 1 to 9,223,372,036,854,775,807). BIGSERIAL has storage and processing advantages over SERIAL8.
SMALLINT	Stores a 2-byte integer value. These values can be in the range from $-(2^{15}-1)$ to $2^{15}-1$ (that is, from -32,767 to 32,767).

DECIMAL(*p,s*) Data Types:

The *p* parameter specifies the *precision* (the total number of digits) and the second parameter, (*s*), specifies the *scale* (the number of digits in the fractional part). If you provide only one parameter, an ANSI-compliant database interprets it as the precision of a fixed-point number and the default scale is 0. If you specify no parameters, and the database is ANSI-compliant, then by default the precision is 16 and the scale is 0.

If the database is not ANSI-compliant, and you specify fewer than 2 parameters, you declare a floating-point DECIMAL, which is not an exact number data type. (See instead the section "Approximate Numeric Data Types" on page 4-28.)

DECIMAL(*p*, *s*) values are stored internally with the first byte representing a sign bit and a 7-bit exponent in excess-65 format. The other bytes express the mantissa as base-100 digits. This implies that DECIMAL(32, *s*) data types store only *s*-1 decimal digits to the right of the decimal point, if *s* is an odd number.

Serial Data Types:

You can declare columns of SERIAL, BIGSERIAL, or SERIAL8 data types. If user-defined routines require whole-number values for variables, arguments, or returned data types, specify INT, BIGINT, or INT8 as the data types, rather than SERIAL, BIGSERIAL, or SERIAL8. These data types are integer data types that differ primarily in their names, their range, and their storage requirements. Columns of serial data types cannot store values less than 1. A table can have no more than one SERIAL column and no more than one BIGSERIAL or SERIAL8 column. Because the serial values are assigned by the database server, you cannot use the UPDATE statement to change an existing serial value in the database.

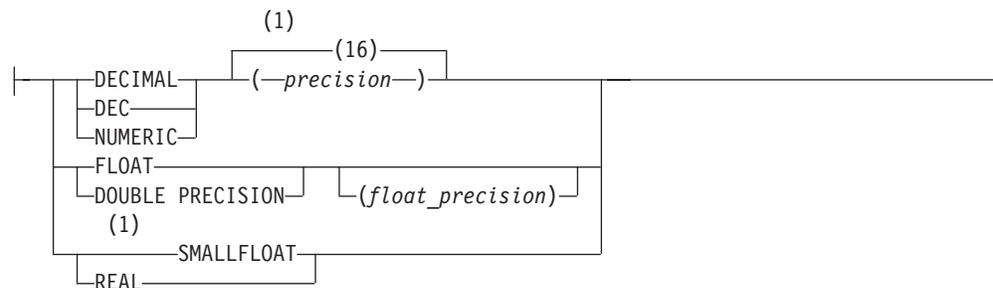
To insert an explicit value into a SERIAL, BIGSERIAL, or SERIAL8 column, specify any integer greater than zero. For details of an alternative way to generate integer values, see "CREATE SEQUENCE" on page 2-184.

A SERIAL, BIGSERIAL, or SERIAL8 column is unique only if you set a unique index on the column. (The index can also be in the form of a primary key or a unique constraint.) With a unique index, values in serial data type columns are guaranteed to be unique, but successive values are not necessarily contiguous.

Approximate Numeric Data Types

An *approximate numeric* data type represents numeric values approximately.

Approximate Numeric Data Type:



Notes:

- 1 Informix extension

Element	Description	Restrictions	Syntax
<i>float_precision</i>	The <i>float_precision</i> is ignored, but is ANSI/ISO compliant.	Must be a positive integer. Specified value has no effect.	"Literal Number " on page 4-185
<i>precision</i>	Significant digits. Default is 16.	An integer; $1 \leq precision \leq 32$	"Literal Number " on page 4-185

Use approximate numeric data types for very large and very small numbers that can tolerate some degree of rounding during arithmetic operations.

The following table summarizes the built-in approximate numeric data types.

Data Type	Description
DEC(<i>p</i>)	Synonym for DECIMAL(<i>p</i>)
DECIMAL(<i>p</i>)	Stores floating-point decimal values in the approximate range from 1.0E-130 to 9.99E+126 The <i>p</i> parameter specifies the precision. If no precision is specified, the default is 16. This floating-point data type is available as an approximate numeric type only in a database that is not ANSI-compliant. In an ANSI-compliant database, DECIMAL(<i>p</i>) is implemented as a fixed-point DECIMAL; see “Exact Numeric Data Types” on page 4-26.
DOUBLE PRECISION	ANSI-compliant synonym for FLOAT. The <i>float_precision</i> term is not valid when you use this synonym in data type declarations.
FLOAT	Stores double-precision floating-point numbers with up to 16 significant digits. The <i>float_precision</i> parameter is accepted in data-type declarations for compliance with the ANSI/ISO standard for SQL, but this parameter has no effect on the actual precision of values that the database server stores.
NUMERIC(<i>p</i>)	ANSI-compliant synonym for DECIMAL(<i>p</i>) In an ANSI-compliant database, this is implemented as an exact numeric type, with the specified precision and a scale of zero, rather than an approximate numeric (floating-point) data type.
REAL	ANSI-compliant synonym for SMALLFLOAT
SMALLFLOAT	Stores single-precision floating-point numbers with approximately 8 significant digits

The built-in number data types of Informix database servers support real numbers. They cannot directly store imaginary or complex numbers.

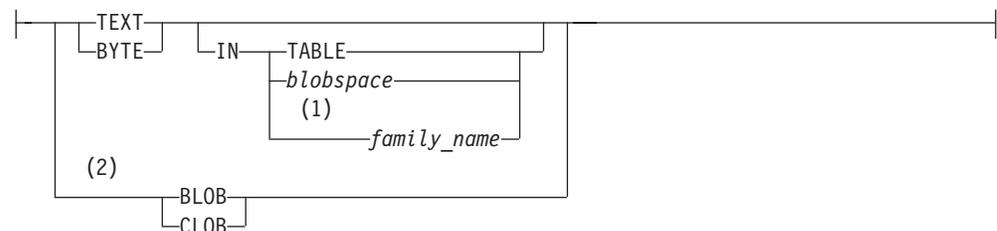
In Dynamic Server, you must create a user-defined data type for applications that support values that can have an imaginary part.

No more than nine arguments to an external UDR can be DECIMAL data types of SQL that the UDR declares as BigDecimal data types of the Java language.

Large-Object Data Types

Large-object data types can store extremely large column values, such as images and documents, independently of the column.

Large-Object Data Type:



Notes:

- 1 Optical Subsystem only

2 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>blobspace</i>	Name of an existing blobspace	Must exist	"Identifier " on page 5-21
<i>family_name</i>	Family name or variable in the optical family	Must exist	"Quoted String" on page 4-189.

The large object data types can be classified in two categories:

- Simple large objects: TEXT and BYTE
- Smart large objects: CLOB and BLOB

Simple-Large-Object Data Types:

These are the simple-large-object data types:

Data Type

Description

TEXT Stores text data of up to 2^{31} bytes

BYTE Stores any digitized data of up to 2^{31} bytes

Do not supply a BYTE value where TEXT is expected. No built-in cast supports BYTE to TEXT data-type conversion.

Because of the maximum row size limit of 32,767 bytes, you cannot create a table with more than approximately 195 BYTE or TEXT columns. (This restriction also applies to all varying-length and ROW data types.)

Storing BYTE and TEXT Data:

A simple-large-object data type can store text or binary data in blobspaces or in tables. The database server can access a BYTE or TEXT value in one piece. When you specify a BYTE or TEXT data type, you can specify the location in which it is stored. You can store data with the table or in a separate blobspace.

In Dynamic Server, if you are creating a named ROW data type that has a BYTE or TEXT field, you cannot use the IN clause to specify a separate storage space.

The following example shows how blobspaces and dbspaces are specified. The user creates the **resume** table. The data values are stored in the **employ** dbspace. The data in the **vita** column is stored with the table, but the data associated with the **photo** column is stored in a blobspace named **photo_space**.

```
CREATE TABLE resume
(
  fname      CHAR(15),
  lname      CHAR(15),
  phone      CHAR(18),
  recd_date  DATETIME YEAR TO HOUR,
  contact_date DATETIME YEAR TO HOUR,
  comments   VARCHAR(250, 100),
  vita       TEXT IN TABLE,
  photo      BYTE IN photo_space
)
IN employ;
```

Smart-Large-Object Data Types (IDS):

A smart-large-object data type stores text or binary data in sbspaces. The database server can provide random access to a smart-large-object value. That is, it can access any portion of the smart-large-object value. These data types are recoverable. The following table summarizes the smart-large-object data types that Dynamic Server supports.

Data Type

Description

BLOB Stores binary data of up to 4 terabytes (4*2⁴⁰ bytes)

CLOB Stores text data of up to 4 terabytes (4*2⁴⁰ bytes)

Both of these are built-in opaque data types. Like most opaque types, they cannot be accessed in a database of a non-local database server by a distributed query or by other DML operations, nor can they be returned from a database of another database server by a UDR. For information on accessing BLOB or CLOB values in other databases of the local server, however, see "BOOLEAN and Other Built-In Opaque Data Types (IDS)" on page 4-22.

Smart large object data types are not parallelizable. The PDQ feature of Dynamic Serve has no effect on operations that load or unload BLOB or CLOB values, or that process them in queries or in other DML operations.

For more information about these "smart blob" data types, see the *IBM Informix Guide to SQL: Reference*.

For information on how to create blobspaces, see your *IBM Informix Administrator's Guide*.

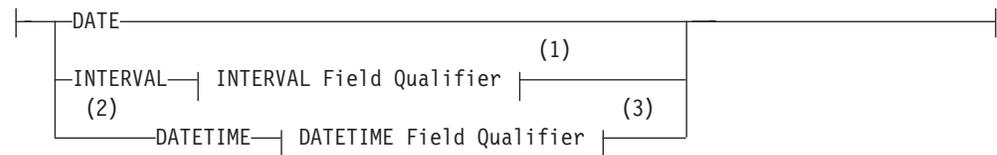
For information about optical families, see the *IBM Informix Optical Subsystem Guide*.

For information about the built-in functions that you can use to import, export, and copy smart large objects, see "Smart-Large-Object Functions (IDS)" on page 4-122 and the *IBM Informix Guide to SQL: Tutorial*.

Time Data Types

The time data types allow the database server to store calendar dates, points in time, and intervals of time.

Time Data Types:



Notes:

- 1 See "INTERVAL Field Qualifier " on page 4-176
- 2 Informix extension
- 3 See " DATETIME Field Qualifier" on page 4-38

The following table summarizes the built-in time data types.

Data Type
Description

DATE Stores a date value as a Julian date in the range from January 1 of the year 1 up to December 31, 9999.

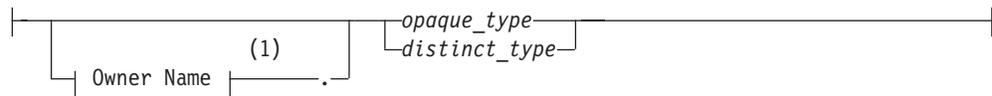
DATETIME
 Stores a point-in-time date (*year, month, day*) and time-of-day (*hour, minute, second, and fraction of second*), in the range of years 1 to 9999. Also supports contiguous subsets of these time units.

INTERVAL
 Stores spans of time, in *years* and/or *months*, or in smaller time units (*days, hours, minutes, seconds, and/or fractions of second*), with up to 9 digits of precision in the largest time unit, if this is not FRACTION. Also supports contiguous subsets of these time units.

User-Defined Data Type (IDS)

A user-defined data type is one that a user defines for the database server. Dynamic Server supports two categories of user-defined data types, namely *distinct data types* and *opaque data types*. This is the declaration syntax for user-defined data types:

User-Defined Data Type:



Notes:

- 1 See " Owner Name " on page 5-45

Element	Description	Restrictions	Syntax
<i>distinct_type</i>	Distinct data type with same structure as an existing data type	Must be unique among data type names in the database	"Identifier " on page 5-21
<i>opaque_type</i>	Name of the opaque data type	Must be unique among data type names in the database	"Identifier " on page 5-21

In this document, *user-defined data type* is usually abbreviated as UDT.

Distinct Data Types

A DISTINCT data type is a user-defined data type that is based on one of the following data types:

- a built-in type (including built-in opaque types)
- a user-defined opaque type
- a named ROW type
- an existing DISTINCT type.

The base type of a DISTINCT type cannot be any of the following data types:

- an unnamed ROW type
- a LIST, MULTISSET, SET, or generic COLLECTION type.

The DISTINCT type inherits the length and the alignment of its base type in storage. Dynamic Server automatically creates explicit casts between the DISTINCT type and its base type. To create a DISTINCT type, you must use the CREATE DISTINCT TYPE statement. (For more information, see “CREATE DISTINCT TYPE ” on page 2-100.)

DISTINCT Types in Distributed Operations:

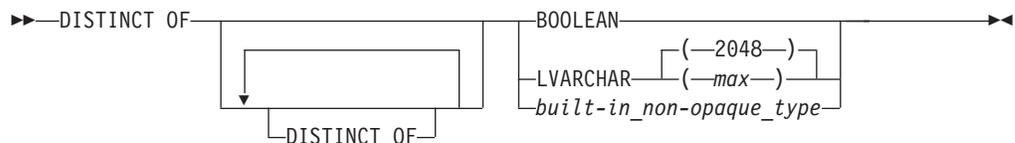
DISTINCT column values cannot be retrieved from another database of the same Dynamic Server instance by a distributed query (nor modified by INSERT, DELETE, or UPDATE cross-database distributed operations) unless all of the following conditions are true:

- the DISTINCT type is defined on one of the following base types:
 - a non-opaque built-in data type
 - a BOOLEAN or LVARCHAR data type
 - a DISTINCT type defined on BOOLEAN, on LVARCHAR, or on a non-opaque built-in data type.

(This condition also applies recursively to DISTINCT types of DISTINCT types, where the ultimate base type is BOOLEAN, or LVARCHAR, or a non-opaque built-in data type.)

- the DISTINCT type is explicitly cast to BOOLEAN, to LVARCHAR, or to a non-opaque built-in type
- the DISTINCT type, its type hierarchy, and its explicit cast to a built-in type are defined exactly the same way in all participating databases.

For DISTINCT data types in distributed operations, the data type hierarchy must have one of these forms, which cannot vary across the participating databases:



The IDSSecurityLabel data type, which stores the security label in rows of protected tables, is a built-in DISTINCT type that satisfies this requirement, because its base type is the built-in VARCHAR(128) data type.

A user-defined routine can return to the local database a DISTINCT data type from another database of the same Dynamic Server instance only if all of the conditions listed above are true, and the UDR is defined in all of the participating databases.

The same rules that apply to DISTINCT data types in distributed operations across databases of the same Dynamic Server instance also apply to DISTINCT data types in cross-server distributed operations on databases of different Dynamic Server instances.

For additional information about the data types that Dynamic Server supports in distributed operations, see “Data Types in Distributed Queries” on page 2-535.

Opaque Data Types

An opaque data type is a user-defined data type that can be used in the same way as a built-in data type. To create an opaque type, you must use the CREATE

OPAQUE TYPE statement. Because an opaque type is encapsulated, you create support functions to access the individual components of an opaque type. The internal storage details of the type are hidden or opaque.

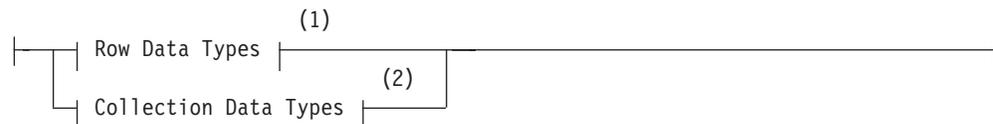
For more information about how to create an opaque data type and its support functions, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

Because of the maximum row size limit of 32,767 bytes, when you create a new table, no more than approximately 195 columns can be varying-length opaque or distinct user-defined data types. (The same restriction applies to BYTE, TEXT, VARCHAR, LVARCHAR, NVARCHAR, and ROW type columns. See "ROW Data Types" for additional information about ROW data types.)

Complex Data Type (IDS)

Complex data types are ROW types or COLLECTION types that you create from built-in types, opaque types, distinct types, or other complex types.

Complex Data Type:



Notes:

- 1 See "CREATE ROW TYPE " on page 2-168
- 2 See " Collection Data Types" on page 4-35

A single complex data type can include multiple components. When you create a complex type, you define the components of the complex type. Unlike an opaque type, however, a complex type is not encapsulated. You can use SQL to access the individual components of a complex data type. The individual components of a complex data type are called *elements*.

Dynamic Server supports the following categories of complex data types:

- ROW data types: Named ROW types and unnamed ROW types
- COLLECTION data types: SET, MULTISET, and LIST

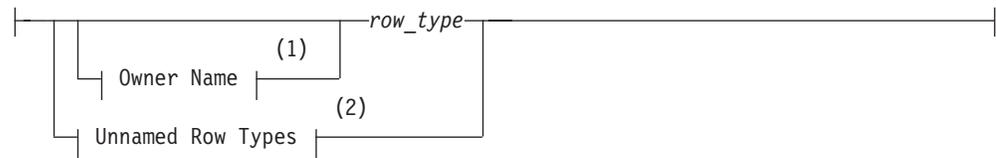
The elements of a COLLECTION data type must all be of the same data type. You can use the keyword COLLECTION in SPL data type declarations to specify an untyped collection variable. NULL values are not supported in elements of COLLECTION data types.

The elements of a ROW data type can be of different data types, but the pattern of data types from the first to the last element cannot vary for a given ROW data type. NULL values are supported in elements of ROW data types, unless you specify otherwise in the data type declaration or in a constraint.

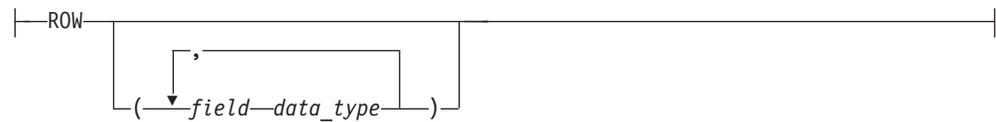
ROW Data Types

This is the syntax to define a column as a named or unnamed ROW type.

Row Data Types:



Unnamed Row Types:



Notes:

- 1 See “ Owner Name ” on page 5-45
- 2 See “CREATE ROW TYPE ” on page 2-168

Element	Description	Restrictions	Syntax
<i>data_type</i>	Data type of <i>field</i>	Any data type except BYTE or TEXT	“Data Type ” on page 4-21
<i>field</i>	Name of a field within <i>row_type</i>	Must be unique among fields of the same ROW type	“Identifier ” on page 5-21
<i>row_type</i>	Some ROW data type defined by CREATE ROW TYPE statement	ROW type must exist in the database	“Identifier ” on page 5-21; “Data Type ” on page 4-21

You can assign a named ROW type to a table, to a column, or to an SPL variable. A named ROW type that you use to create a typed table or to define a column must already exist. For information on how to create a named ROW data type, see “CREATE ROW TYPE ” on page 2-168.

To specify a named ROW data type in an ANSI-compliant database, you must qualify the *row_type* with its *owner* name, if you are not the owner of *row_type*.

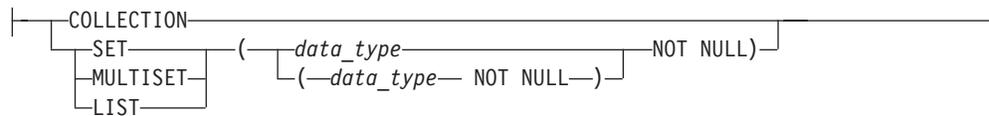
An unnamed ROW data type is identified by its structure, which specifies fields that you create with its ROW constructor. You can define a column or an SPL variable as an unnamed ROW data type. For the syntax to specify values for an unnamed ROW type, see “ROW Constructors” on page 4-82.

Because of the maximum row size limit of 32,767 bytes, a single table cannot be created with more than approximately 195 ROW type columns.

Collection Data Types

This diagram shows the syntax to define a column or an SPL variable as a collection data type. (A table can include no more than 97 columns of collection data types.) For the syntax to specify values of collection elements, see “ Collection Constructors” on page 4-83.

Collection Data Type:



Element	Description	Restrictions	Syntax
<i>data_type</i>	Data type of each of the elements of the collection	Can be any data type except BIGSERIAL, BYTE, SERIAL, or SERIAL8, or TEXT	"Data Type " on page 4-21

A SET is an unordered collection of elements, each of which has a unique value. Define a column as a SET data type when you want to store collections whose elements contain no duplicate values and have no associated order.

A MULTISET is an unordered collection of elements that can have duplicate values. You can define a column as a MULTISET collection type when you want to store collections whose elements might not be unique and have no specific order associated with them.

A LIST is an ordered collection of elements that can include duplicate elements. A LIST differs from a MULTISET in that each element in a LIST collection has an ordinal position in the collection. You can define a column as a LIST collection type when you want to store collections whose elements might not be unique but have a specific order associated with them.

The keyword COLLECTION can be used in SPL data type declarations to specify an untyped collection variable.

If you attempt to insert a collection that includes one or more duplicate values into a SET column, Dynamic Server issues no error, but the duplicate values are ignored, and only the unique values are inserted.

Duplicate Elements in DML Operations on SET Columns: The SET data type does not allow duplicate element values in the same collection. If you attempt to insert duplicate elements into a SET data type, or to update a SET column or variable to a value that includes duplicate elements, the database server issues no error or warning when the INSERT or UPDATE statement executes, but only one of the duplicate elements is stored in the SET column or variable.

For example, suppose you create table **t3** with column **a** of the SET data type, and then you insert four rows, some of which include elements that have identical values:

```

> CREATE TABLE t3(a SET(INT NOT NULL));

Table created.

> INSERT INTO t3 VALUES( SET{10, 20, 30} );

1 row(s) inserted.

> INSERT INTO t3 VALUES( SET{10, 20, 10});

1 row(s) inserted.

> INSERT INTO t3 VALUES( SET{10, 10, 10});
  
```

```

1 row(s) inserted.
> INSERT INTO t3 VALUES( SET{10,10,10});
1 row(s) inserted.

```

When you look at the data values that were inserted into column **t3.a**, the four inserted rows include no duplicate element values:

```

> SELECT * FROM t3;
a SET{10      ,20      ,30      }
a SET{10      ,20      }
a SET{10      }
a SET{10      }

4 row(s) retrieved.

```

In this example, Dynamic Server silently discarded all but one instance of the duplicated elements from what the VALUES clause of the INSERT statement specified for each SET value.

Similar behavior occurs if the SET clause of the UPDATE statement includes duplicate elements within the same SET value. Declare collection columns of the MULTISSET data type, rather than of the SET data type, if you want the database to store unordered sets that can include duplicate elements within the same collection

Defining the Element Type: The element type can be any data type except TEXT, BYTE, SERIAL, or SERIAL8. You can nest collection types, using elements of a collection type.

Every element must be of the same type. For example, if the element type of a collection data type is INTEGER, every element must be of type INTEGER.

An exception to this restriction occurs if the database server determines that some elements of a collection of character strings are VARCHAR data types (whose length is limited to 255 or fewer bytes) but other elements are longer than 255 bytes. In this case, the collection constructor can assign a CHAR(*n*) data type to all elements, for *n* the length in bytes of the longest element. If this is undesirable, you can cast the collection to LVARCHAR, to prevent padding extra length in elements of the collection, as in this example:

```

LIST {'first character string longer than 255 bytes . . . ',
      'second character string longer than 255 bytes . . . ',
      'another character string'} ::LIST (LVARCHAR NOT NULL)

```

See “ Collection Constructors” on page 4-83 for additional information.

If the element type of a collection is an unnamed ROW type, the unnamed ROW type cannot contain fields that hold unnamed ROW types. That is, a collection cannot contain nested unnamed ROW data types.

The elements of a collection cannot be NULL. When you define a column as a collection data type, you must use the NOT NULL keywords to specify that the elements of the collection cannot be NULL.

Privileges on a collection data type are those of the database column. You cannot specify privileges on individual elements of a collection.

Related Information

For more information about choosing a data type for your database, see the *IBM Informix Database Design and Implementation Guide*.

For more information about the specific qualities of individual data types, see the chapter on data types in the *IBM Informix Guide to SQL: Reference*.

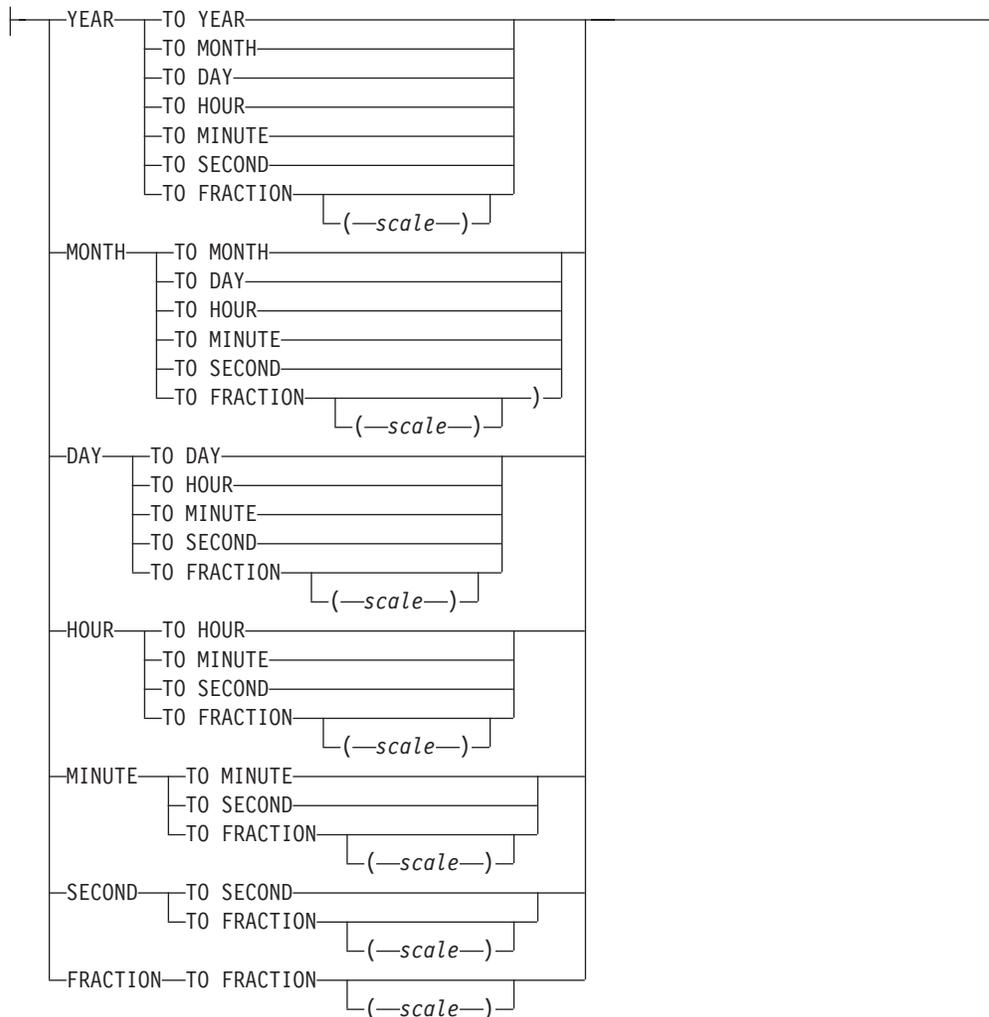
For more information about data types for storing character data in multibyte locales, see the discussion of the NCHAR and NVARCHAR data types and the GLS aspects of other character data types in the *IBM Informix GLS User's Guide*.

DATETIME Field Qualifier

Use a DATETIME Field Qualifier to specify the largest and smallest unit of time in a DATETIME column or value. Use this segment whenever you see a reference to a DATETIME Field Qualifier in a syntax diagram.

Syntax

DATETIME Field Qualifier:



Element	Description	Restrictions	Syntax
<i>scale</i>	Fraction of a second. Default is 3.	Integer ($1 \leq scale \leq 5$)	"Literal Number " on page 4-185

Usage

This segment specifies the precision and scale of a DATETIME data type.

Specify, as the first keyword, the largest time unit that the DATETIME column will store. After the keyword TO, specify the smallest unit as the last keyword. These can be the same keyword. If they are different, the qualifier implies that any intermediate time units between the first and last are also recorded by the DATETIME data type.

The keywords can specify the following time units for the DATETIME column.

Unit of Time

Description

YEAR Specifies a year, in the range from A.D. 1 to 9999

MONTH

Specifies a month, in the range from 1 (January) to 12 (December)

DAY Specifies a day, in the range from 1 to 28, 29, 30, or 31 (depending on the specific month)

HOUR

Specifies an hour, in the range from 0 (midnight) to 23

MINUTE

Specifies a minute, in the range from 0 to 59

SECOND

Specifies a second, in the range from 0 to 59

FRACTION

Specifies a fraction of a second, with up to five decimal places

The default scale is three digits (thousandth of a second).

Unlike INTERVAL qualifiers, DATETIME qualifiers cannot specify nondefault precision (except for FRACTION, when FRACTION is the smallest unit in the qualifier). Some examples of DATETIME qualifiers follow:

```
YEAR TO MINUTE           MONTH TO MONTH
DAY TO FRACTION(4)       MONTH TO DAY
```

On some platforms, the system clock cannot support precision greater than FRACTION(3).

An error results if the first keyword represents a smaller time unit than the last, or if you use the plural form of a keyword (such as MINUTES).

Operations on DATETIME values that do not include YEAR in their qualifier use values from the system clock-calendar to supply any additional precision. If the first term in the qualifier is DAY, and the current month has fewer than 31 days, unexpected results can occur.

Related Information

For an explanation of the DATETIME Field Qualifier, see the discussion of the DATETIME data type in the *IBM Informix Guide to SQL: Reference*.

For important differences between the syntax of DATETIME and INTERVAL field qualifiers, see “INTERVAL Field Qualifier ” on page 4-176.

Expression

Data values in SQL statements must be represented as expressions. An *expression* is a specification, which can include operators, operands, and parentheses, that the database server can evaluate to one or more values, or to a reference to some database object.

Expressions can refer to values already in a table of the database, or to values derived from such data, but some expressions (such as TODAY, USER, or literal values) can return values that are independent of the database. You can use expressions to specify values in data-manipulation statements, to define fragmentation strategies, and in other contexts. Use the Expression segment whenever you see a reference to an expression in a syntax diagram.

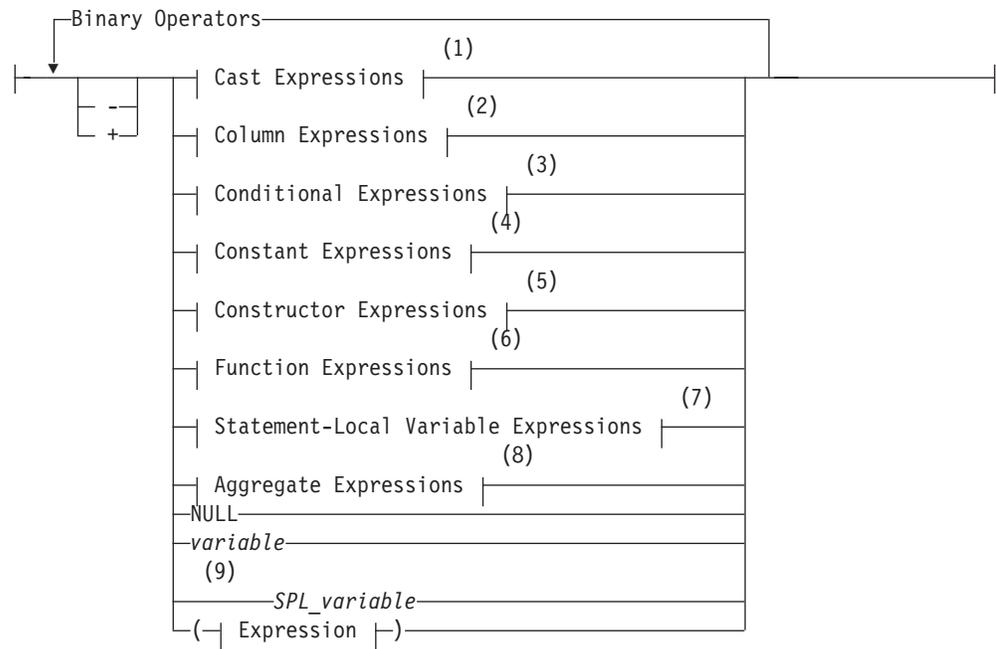
In most contexts, however, you are restricted to expressions whose returned value is of some specific data type, or of a data type that can be converted by the database server to some required data type.

For an alphabetical listing of the built-in operators and functions that are described in this segment, see “List of Expressions” on page 4-42.

Syntax of SQL Expressions

The sections that follow describe SQL expressions, which are specifications that return one or more values or references to database objects. IBM Informix database servers support the following categories of expressions:

SQL Expressions:



Binary Operators:



Notes:

- 1 See "Cast Expressions (IDS)" on page 4-58
- 2 See "Column Expressions" on page 4-59
- 3 See "Conditional Expressions" on page 4-65
- 4 See "Constant Expressions" on page 4-70
- 5 See "Constructor Expressions (IDS)" on page 4-82
- 6 See "Function Expressions" on page 4-87
- 7 See "Statement-Local Variable Expressions (IDS)" on page 4-164
- 8 See "Aggregate Expressions" on page 4-165
- 9 Stored Procedure Language only

Element	Description	Restrictions	Syntax
<i>SPL_variable</i>	In an SPL routine, a variable that contains some expression type that the syntax diagram shows	Must conform to the rules for expressions of that type	"Identifier " on page 5-21
<i>variable</i>	Host or program variable that contains some expression type that the syntax diagram shows	Must conform to the rules for expressions of that type	Language-specific rules for names

Usage

The following table lists the types of SQL expressions, as identified in the diagram for “Expression” on page 4-40, and describes what each type returns.

Expression Type	Description
Aggregate functions	Returns values from built-in or from user-defined aggregates
Arithmetic operators	Supports arithmetic operations on one (unary operators) or two (binary operators) numeric operands
Concatenation operator	Concatenates two string values
Cast operators	Explicit casts from one data type to another
Column expressions	Full or partial column values
Conditional expressions	Returns values that depend on conditional tests
Constant expressions	Literal values in data manipulation (DML) statements
Constructor expressions	Dynamically creates values for complex data types
Function expressions	Returns values from built-in or user-defined functions
Statement-Local Variable expressions	References a statement-local variable (SLV) in the same SQL statement where it was declared

You can also use host variables or SPL variables as expressions. For a complete list with page references to this chapter, see the following “List of Expressions.”

List of Expressions

Each category of SQL expression includes many individual expressions. The following table lists all the SQL expressions (and some operators) in alphabetical order. The columns in this table have the following meanings:

- **Name** gives the name of each expression.
- **Description** gives a short description of each expression.
- **Syntax** lists the page that shows the syntax of the expression.
- **Usage** shows the page that describes the usage of the expression.

Each expression listed in the following table is supported on all database servers unless otherwise noted. When an expression is not supported on all database servers, the **Name** column notes in parentheses the database server or servers that do support the expression.

Name	Description	Syntax	Usage
ABS function	Returns absolute value of a numeric argument	“Algebraic Functions” on page 4-89	“ABS Function” on page 4-90
ACOS function	Returns the arc cosine of a numeric argument	“Trigonometric Functions” on page 4-139	“ACOS Function” on page 4-140
ADD_MONTHS function	Adds a specified number of months	“Time Functions” on page 4-128	“ADD_MONTHS Function” on page 4-129
Addition (+) operator	Returns the sum of two numeric operands	“Expression” on page 4-40	“Arithmetic Operators” on page 4-51
ASCII function	Returns the ASCII codepoint of the first character in its string argument	“String-Manipulation Functions” on page 4-141	“ASCII Function” on page 4-146

Name	Description	Syntax	Usage
ASIN function	Returns the arc sine of a numeric argument	"Trigonometric Functions" on page 4-139	" ASIN Function" on page 4-141
ATAN function	Returns the arc tangent of numeric argument	"Trigonometric Functions" on page 4-139	" ATAN Function" on page 4-141
ATAN2 function	Computes the angular component of the polar coordinates (<i>r</i> , <i>q</i>) associated with (<i>x</i> , <i>y</i>)	"Trigonometric Functions" on page 4-139	" ATAN2 Function" on page 4-141
AVG function	Returns the mean of a set of numeric values	"Aggregate Expressions" on page 4-165	" AVG Function" on page 4-168
BITAND	Returns the bitwise AND of two arguments	"Bitwise Logical Functions" on page 4-52	"BITAND Function" on page 4-53
BITANDNOT	Returns the bitwise ANDNOT of two arguments	"Bitwise Logical Functions" on page 4-52	"BITANDNOT Function" on page 4-55
BITNOT	Returns the bitwise NOT of two arguments	"Bitwise Logical Functions" on page 4-52	"BITNOT Function" on page 4-55
BITOR	Returns the bitwise OR of two arguments	"Bitwise Logical Functions" on page 4-52	"BITOR Function" on page 4-54
BITXOR	Returns the bitwise XOR of two arguments	"Bitwise Logical Functions" on page 4-52	"BITXOR Function" on page 4-54
CARDINALITY function (IDS)	Returns the number of elements in a collection data type (SET, MULTISSET, or LIST)	"CARDINALITY Function (IDS)" on page 4-98	"CARDINALITY Function (IDS)" on page 4-98
CASE expression	Returns a value that depends on which of several conditional tests evaluates to true	"CASE Expressions" on page 4-66	"CASE Expressions" on page 4-66
CAST expression (IDS)	Converts an expression to a specified data type	"Cast Expressions (IDS)" on page 4-58	"Cast Expressions (IDS)" on page 4-58
Cast (::) operator	See "Double-colon (::) cast operator"	"Cast Expressions (IDS)" on page 4-58	"Cast Expressions (IDS)" on page 4-58
CEIL function	Returns the smallest integer that is greater than or equal to its single argument	"Algebraic Functions" on page 4-89	" CEIL Function" on page 4-90
CHARACTER_LENGTH function	See CHAR_LENGTH function. (In multibyte locales, this replaces the LENGTH function.)	" Length Functions" on page 4-118	" CHAR_LENGTH Function" on page 4-119
CHAR_LENGTH function	Returns count of logical characters in a string	" Length Functions" on page 4-118	" CHAR_LENGTH Function" on page 4-119
Column expression	Complete or partial column value from a table	"Column Expressions" on page 4-59	"Column Expressions" on page 4-59
CONCAT() operator function	Concatenates the results of two expressions	"String-Manipulation Functions" on page 4-141	"CONCAT Function" on page 4-142
Concatenation () operator	Concatenates the results of two expressions	"Expression" on page 4-40	"Concatenation Operator" on page 4-56
Constant expression	Expression with a literal, fixed, or variant value	"Constant Expressions" on page 4-70	"Constant Expressions" on page 4-70
COS function	Returns the cosine of a radian expression	"Trigonometric Functions" on page 4-139	" COS Function" on page 4-140

Name	Description	Syntax	Usage
COUNT (as a set of functions)	Functions that return frequency counts Each form of the COUNT function is listed below.	"Aggregate Expressions" on page 4-165	"Overview of COUNT Functions" on page 4-168
COUNT (ALL column) function	See COUNT (column) function.	"Aggregate Expressions" on page 4-165	"COUNT column Function" on page 4-169
COUNT (column) function	Returns the number of non-NULL values in a specified column	"Aggregate Expressions" on page 4-165	"COUNT column Function" on page 4-169
COUNT DISTINCT function	Returns the number of unique non-NULL values in a specified column	"Aggregate Expressions" on page 4-165	"COUNT DISTINCT and COUNT UNIQUE Functions" on page 4-169
COUNT UNIQUE function	See COUNT DISTINCT function.	"Aggregate Expressions" on page 4-165	"COUNT DISTINCT and COUNT UNIQUE Functions" on page 4-169
COUNT (*) function	Returns the cardinality of the set of rows that satisfy a query	"Aggregate Expressions" on page 4-165	"COUNT(*) Function" on page 4-169
CURRENT operator	Returns the current time as a DATETIME value that consists of the date and the time of day	"Constant Expressions" on page 4-70	"CURRENT Operator" on page 4-75
CURRENT_ROLE operator	Returns the currently enabled role of the user	"Constant Expressions" on page 4-70	"CURRENT_ROLE Operator" on page 4-73
sequence.CURRVAL (IDS)	Returns the current value of specified <i>sequence</i>	"Constant Expressions" on page 4-70	"Using CURRVAL" on page 4-80
DATE function	Converts a nondate argument to a DATE value	"Time Functions" on page 4-128	"DATE Function" on page 4-130
DAY function	Returns the day of the month as an integer	"Time Functions" on page 4-128	"DAY Function" on page 4-130
DBINFO (<i>option</i>)	Functions for retrieving database and session information. Each <i>option</i> is listed below.	"DBINFO Function" on page 4-99	"DBINFO Options" on page 4-100
DBINFO ('bigserial')	Returns most recently inserted SERIAL8 value	"DBINFO Function" on page 4-99	"Using the 'serial8' Option and 'bigserial' Option" on page 4-105
DBINFO ('coserverid', <i>table.column</i> , 'currentrow') (XPS)	Returns the coserver ID of the coserver where each row of a specified table is located	"DBINFO Function" on page 4-99	"Using the 'coserverid' Option Followed by Table and Column Names (XPS)" on page 4-107
DBINFO ('coserverid') (XPS)	Returns the coserver ID of the coserver to which the user who entered the query is connected	"DBINFO Function" on page 4-99	"Using the 'coserverid' Option with No Other Arguments (XPS)" on page 4-106
DBINFO ('dbhostname')	Returns the host name of the database server to which a client application is connected	"DBINFO Function" on page 4-99	"Using the 'dbhostname' Option" on page 4-104
DBINFO ('dbspace', <i>table.column</i> , 'currentrow') (XPS)	Returns the name of the dbspace where each row of a specified table is located	"DBINFO Function" on page 4-99	"Using the 'dbspace' Option Followed by Table and Column Names (XPS)" on page 4-107

Name	Description	Syntax	Usage
DBINFO ('dbspace', <i>tblspace_number</i>)	Returns the name of a dbspace corresponding to a tblspace number	" DBINFO Function" on page 4-99	"Using the ('dbspace', <i>tblspace_num</i>) Option" on page 4-101
DBINFO ('get_tz')	Returns the time zone of the current session	" DBINFO Function" on page 4-99	" Using the 'get_tz' Option" on page 4-105
DBINFO ('serial8')	Returns most recently inserted SERIAL8 value	" DBINFO Function" on page 4-99	" Using the 'serial8' Option and 'bigserial' Option" on page 4-105
DBINFO ('sessionid')	Returns the session ID of the current session	" DBINFO Function" on page 4-99	"Using the 'sessionid' Option" on page 4-103
DBINFO ('sqlca.sqlerrd1')	Returns the last serial value inserted in a table	" DBINFO Function" on page 4-99	"Using the 'sqlca.sqlerrd1' Option" on page 4-102
DBINFO ('sqlca.sqlerrd2')	Returns the number of rows processed by DML statements, and by EXECUTE PROCEDURE and EXECUTE FUNCTION statements	" DBINFO Function" on page 4-99	"Using the 'sqlca.sqlerrd2' Option" on page 4-102
DBINFO ('utc_current')	Returns the current Coordinated Universal Time (UTC) value.	" DBINFO Function" on page 4-99	" Using the 'utc_current' Option" on page 4-106
DBINFO ('utc_to_datetime', <i>expression</i>)	Returns the DATETIME value of an integer or column <i>expression</i> that specifies a UTC value.	" DBINFO Function" on page 4-99	" Using the 'utc_to_datetime' Option" on page 4-106
DBINFO ('version', <i>parameter</i>)	Returns all or part, as specified by the <i>parameter</i> , of the exact version of the database server to which the client application is connected.	" DBINFO Function" on page 4-99	"Using the 'version' Option" on page 4-104
DBSERVERNAME function	Returns the name of the database server	"Constant Expressions" on page 4-70	"DBSERVERNAME and SITENAME Operators" on page 4-74
DECODE function	Evaluates one or more expression pairs and compares the <i>when</i> expression in each pair with a specified value expression	" DECODE Function" on page 4-69	" DECODE Function" on page 4-69
DECRYPT_CHAR function (IDS)	Returns a plain-text string or CLOB after processing an encrypted argument	" Encryption and Decryption Functions" on page 4-107	"DECRYPT_CHAR Function" on page 4-113
DECRYPT_BINARY function (IDS)	Returns a plain-text BLOB data value after processing an encrypted BLOB argument	" Encryption and Decryption Functions" on page 4-107	"DECRYPT_BINARY Function" on page 4-114
DEFAULT_ROLE operator	Returns the default role of the current user	"Constant Expressions" on page 4-70	"DEFAULT_ROLE Operator" on page 4-73
DELETING Boolean operator	Returns 't' if triggering event is a DELETE	"Trigger-Type Boolean Operator" on page 4-12	"Trigger-Type Boolean Operator" on page 4-12
Division (/) operator	Returns the quotient of two numeric operands	"Expression" on page 4-40	"Arithmetic Operators" on page 4-51
Double-colon (::) cast operator (IDS)	Converts the value of an expression to a specified data type	"Cast Expressions (IDS)" on page 4-58	"Cast Expressions (IDS)" on page 4-58

Name	Description	Syntax	Usage
Double-pipe () concatenation operator	Returns a string that joins one string operand to another string operand	"Expression" on page 4-40	"Concatenation Operator" on page 4-56
ENCRYPT_AES function (IDS)	Returns an encrypted string or BLOB after processing a plain-text string, BLOB, or CLOB	"Encryption and Decryption Functions" on page 4-107	"ENCRYPT_AES Function" on page 4-114
ENCRYPT_TDES function (IDS)	Returns an encrypted string or BLOB after processing a plain-text string, BLOB, or CLOB	"Encryption and Decryption Functions" on page 4-107	"ENCRYPT_TDES Function" on page 4-115
EXP function	Returns the exponent of a numeric expression	"Exponential and Logarithmic Functions" on page 4-116	"EXP Function" on page 4-117
EXTEND function	Resets precision of DATETIME or DATE value	"Time Functions" on page 4-128	"EXTEND Function" on page 4-134
FILETOBLOB function (IDS)	Creates a BLOB value from data stored in a specified operating-system file	"Smart-Large-Object Functions (IDS)" on page 4-122	"FILETOBLOB and FILETOCLOB Functions" on page 4-123
FILETOCLOB function (IDS)	Creates a CLOB value from data stored in a specified operating-system file	"Smart-Large-Object Functions (IDS)" on page 4-122	"FILETOBLOB and FILETOCLOB Functions" on page 4-123
FLOOR function	Returns the largest integer that is smaller than or equal to its single argument	"Algebraic Functions" on page 4-89	"FLOOR Function" on page 4-90
FORMAT_UNITS function	Returns a character string that specifies a number and the names of abbreviated units of memory or of storage	"FORMAT_UNITS Function" on page 4-158	"FORMAT_UNITS Function" on page 4-158
GETHINT function (IDS)	Returns a plain-text hint string after processing an encrypted data-string argument	"Encryption and Decryption Functions" on page 4-107	"GETHINT Function" on page 4-116
HEX function	Returns the hexadecimal encoding of a base-10 integer argument	"HEX Function" on page 4-117	"HEX Function" on page 4-117
Host variable	See Variable.	"Syntax of SQL Expressions" on page 4-40	"Syntax of SQL Expressions" on page 4-40
INSERTING Boolean operator	Returns 't' if triggering event is an INSERT	"Trigger-Type Boolean Operator" on page 4-12	"Trigger-Type Boolean Operator" on page 4-12
IFX_ALLOW_NEWLINE function	Sets a newline session mode that allows or disallows newline characters in quoted strings	"IFX_ALLOW_NEWLINE Function" on page 4-160	"IFX_ALLOW_NEWLINE Function" on page 4-160
INITCAP function	Converts a string argument to a string in which only the initial letter of each word is uppercase	"Case-Conversion Functions" on page 4-156	"INITCAP Function" on page 4-157
LAST_DAY function	Returns the date of the last day of the month that its argument specifies	"Time Functions" on page 4-128	"LAST_DAY Function" on page 4-132

Name	Description	Syntax	Usage
LENGTH function	Returns the number of bytes in a character column, not including any trailing blank spaces	"Length Functions" on page 4-118	"LENGTH Function" on page 4-119
LIST collection constructor (IDS)	Constructor for collections whose elements are ordered and can contain duplicate values	"Collection Constructors" on page 4-83	"Collection Constructors" on page 4-83
Literal BOOLEAN	Literal representation of a BOOLEAN value	"Constant Expressions" on page 4-70	"Constant Expressions" on page 4-70
Literal collection (IDS)	Represents elements in a collection data type	"Constant Expressions" on page 4-70	"Literal Collection (IDS)" on page 4-82
Literal DATETIME	Represents a DATETIME value	"Constant Expressions" on page 4-70	"Literal DATETIME" on page 4-78
Literal INTERVAL	Represents an INTERVAL value	"Constant Expressions" on page 4-70	"Literal INTERVAL" on page 4-78
Literal number	Represents a numeric value	"Constant Expressions" on page 4-70	"Literal Number" on page 4-72
Literal opaque type (IDS)	Represents an opaque data type	"Constant Expressions" on page 4-70	"Constant Expressions" on page 4-70
Literal row (IDS)	Represents the elements in a ROW data type	"Constant Expressions" on page 4-70	"Literal Row (IDS)" on page 4-81
LOCOPY function (IDS)	Creates a copy of a smart large object	"Smart-Large-Object Functions (IDS)" on page 4-122	"LOCOPY Function" on page 4-126
LOGN function	Returns the natural log of a numeric argument	"Exponential and Logarithmic Functions" on page 4-116	"LOGN Function" on page 4-117
LOG10 function	Returns the base-10 logarithm of an argument	"Exponential and Logarithmic Functions" on page 4-116	"LOG10 Function" on page 4-117
LOTOFILE function (IDS)	Copies a BLOB or CLOB value to a file	"Smart-Large-Object Functions (IDS)" on page 4-122	"LOTOFILE Function" on page 4-125
LOWER function	Converts uppercase letters to lowercase	"Case-Conversion Functions" on page 4-156	"LOWER Function" on page 4-157
LPAD function	Returns a string that is left-padded by a specified number of pad characters	"String-Manipulation Functions" on page 4-141	"LPAD Function" on page 4-154
LTRIM function	Removes specified leading pad characters from a string.	"String-Manipulation Functions" on page 4-141	"LTRIM Function" on page 4-149
MAX function	Returns the largest in a specified set of values	"Aggregate Expressions" on page 4-165	"MAX Function" on page 4-172
MDY function	Returns a DATE value from integer arguments	"Time Functions" on page 4-128	"MDY Function" on page 4-135
MIN function	Returns the smallest in a specified set of values	"Aggregate Expressions" on page 4-165	"MIN Function" on page 4-172

Name	Description	Syntax	Usage
MOD function	Returns the modulus (the integer-division remainder value) from two numeric arguments	"Algebraic Functions" on page 4-89	"MOD Function" on page 4-91
MONTH function	Returns the month value from a DATE or DATETIME argument	"Time Functions" on page 4-128	"MONTH Function" on page 4-131
MONTHS_ BETWEEN function	Returns the difference in months between two time values	"Time Functions" on page 4-128	"MONTHS_BETWEEN Function" on page 4-131
Multiplication (*) operator	Returns the product of two numeric operands	"Expression" on page 4-40	"Arithmetic Operators" on page 4-51
MULTISET collection constructor (IDS)	Constructor for a non-ordered collection of elements that can contain duplicate value	"Collection Constructors" on page 4-83	"Collection Constructors" on page 4-83
NEXT_DAY function	Returns the earliest calendar date that satisfies each of two conditions	"Time Functions" on page 4-128	"NEXT_DAY Function" on page 4-133
<i>sequence</i> .NEXTVAL (IDS)	Increments value of the specified <i>sequence</i>	"Constant Expressions" on page 4-70	"Using NEXTVAL" on page 4-79
NULL keyword	Unknown, missing, or logically undefined value	"NULL Keyword" on page 4-85	"NULL Keyword" on page 4-85
NULLIF function	Returns NULL if both arguments are equal	"NULLIF Function" on page 4-68	"NULLIF Function" on page 4-68
NVL function	Returns the value of a not-NULL argument, or a specified value if the argument is NULL	"NVL Function" on page 4-68	"NVL Function" on page 4-68
OCTET_LENGTH function	Returns the number of bytes in a character column, including any trailing blank spaces	"Length Functions" on page 4-118	"OCTET_LENGTH Function" on page 4-119
POW function	Raises a base value to a specified power	"Algebraic Functions" on page 4-89	"POW Function" on page 4-91
POWER function	Synonym for POW function	"Algebraic Functions" on page 4-89	"POW Function" on page 4-91
Procedure-call expression	See user-defined function.	"User-Defined Functions" on page 4-161	"User-Defined Functions" on page 4-161
Program variable	See variable.	"Syntax of SQL Expressions" on page 4-40	"Syntax of SQL Expressions" on page 4-40
Quoted string	Literal character string	"Constant Expressions" on page 4-70	"Quoted String" on page 4-72
RANGE function	Returns the range of a specified set of values	"Aggregate Expressions" on page 4-165	"RANGE Function" on page 4-172
REPLACE function	Replaces specified characters in a source string	"String-Manipulation Functions" on page 4-141	"REPLACE Function" on page 4-153
ROOT function	Returns the root value of a numeric argument	"Algebraic Functions" on page 4-89	"ROOT Function" on page 4-91
ROUND function	Returns the rounded value of an argument	"Algebraic Functions" on page 4-89	"ROUND Function" on page 4-91

Name	Description	Syntax	Usage
ROW constructor (IDS)	Constructor for a named ROW data type	"Constructor Expressions (IDS)" on page 4-82	"ROW Constructors" on page 4-82
RPAD function	Returns a string that is right-padded by a specified number of pad characters	"String-Manipulation Functions" on page 4-141	"RPAD Function" on page 4-155
RTRIM function	Removes trailing pad characters from a string	"String-Manipulation Functions" on page 4-141	"RTRIM Function" on page 4-149
SECLABEL_BY_COMP function	Returns the individual components of a row security label for INSERT and UPDATE operations	"Security Label Support Functions (IDS)" on page 4-119	"SECLABEL_BY_NAME Function" on page 4-120
SECLABEL_BY_NAME function	Returns the identifier of a row security label for INSERT and UPDATE operations	"Security Label Support Functions (IDS)" on page 4-119	"SECLABEL_BY_COMP Function" on page 4-120
SECLABEL_TO_CHAR function	Returns the details of the individual components of a row security label for SELECT operations	"Security Label Support Functions (IDS)" on page 4-119	"SECLABEL_TO_CHAR Function" on page 4-121
SELECTING Boolean operator	Returns 't' if triggering event is a SELECT	"Trigger-Type Boolean Operator" on page 4-12	"Trigger-Type Boolean Operator" on page 4-12
SET collection constructor (IDS)	Constructor for an unordered collection of elements in which each value is unique	"Collection Constructors" on page 4-83	"Collection Constructors" on page 4-83
SIN function	Returns the sine of a radian expression	"Trigonometric Functions" on page 4-139	"SIN Function" on page 4-140
SITENAME function	See DBSERVERNAME function.	"Constant Expressions" on page 4-70	"DBSERVERNAME and SITENAME Operators" on page 4-74
SPL routine expression	See "User-defined functions"	"User-Defined Functions" on page 4-161	"User-Defined Functions" on page 4-161
SPL variable	SPL variable that stores an expression	"Syntax of SQL Expressions" on page 4-40	"Syntax of SQL Expressions" on page 4-40
SQLCODE function	Returns <code>sqlca.sqlcode</code> value to an SPL UDR	"SQLCODE Function (SPL)" on page 4-98	"SQLCODE Function (SPL)" on page 4-98
SQRT function	Returns the square root of a numeric argument	"Algebraic Functions" on page 4-89	"SQRT Function" on page 4-91
STDEV function	Returns the standard deviation of a data set	"Aggregate Expressions" on page 4-165	"STDEV Function" on page 4-173
SUBSTR function	Returns a substring of a string argument	"SUBSTR Function" on page 4-152	"SUBSTR Function" on page 4-152
SUBSTRING function	Returns a substring of a source string	"SUBSTRING Function" on page 4-150	"SUBSTRING Function" on page 4-150
Substring ([x, y]) operator	Returns a substring of a string operand	"Column Expressions" on page 4-59	"Using the Substring Operator" on page 4-63
Subtraction (-) operator	Returns the difference between two numbers	"Expression" on page 4-40	"Arithmetic Operators" on page 4-51
SUM function	Returns the sum of a specified set of values	"Aggregate Expressions" on page 4-165	"SUM Function" on page 4-172

Name	Description	Syntax	Usage
SLV expression	A statement-local variable (SLV) whose scope is the SQL statement that declares it	"Statement-Local Variable Declaration (IDS)" on page 4-162	"Statement-Local Variable Expressions (IDS)" on page 4-164
SYSDATE operator	Returns the current DATETIME value from the system clock.	"Constant Expressions" on page 4-70	"SYSDATE Operator" on page 4-76
TAN function	Returns the tangent of a radian expression	"Trigonometric Functions" on page 4-139	"TAN Function" on page 4-140
TO_CHAR function	Converts a time or number to a string	"Time Functions" on page 4-128	"TO_CHAR Function (IDS)" on page 4-135
TO_DATE function	Converts a string to a DATETIME value	"Time Functions" on page 4-128	"TO_DATE Function (IDS)" on page 4-138
TO_NUMBER function	Converts two numbers or two strings to DECIMAL	"TO_NUMBER Function" on page 4-139	"TO_NUMBER Function" on page 4-139
TODAY operator	Returns the current system date	"Constant Expressions" on page 4-70	"TODAY Operator" on page 4-75
TRIM function	Drops pad characters from a string argument	"String-Manipulation Functions" on page 4-141	"TRIM Function" on page 4-147
TRUNC function	Returns a truncated numeric or time value	"Algebraic Functions" on page 4-89	"TRUNC Function" on page 4-95
Unary minus (-) sign	Specifies a negative (< 0) numeric value	"Expression" on page 4-40	"Arithmetic Operators" on page 4-51
Unary plus (+) sign	Specifies a positive (> 0) numeric value .	"Expression" on page 4-40	"Arithmetic Operators" on page 4-51
UNITS operator	Convert an integer to an INTERVAL value	"Constant Expressions" on page 4-70	"UNITS Operator" on page 4-79
UPDATING Boolean operator	Returns 't' if triggering event is an UPDATE	"Trigger-Type Boolean Operator" on page 4-12	"Trigger-Type Boolean Operator" on page 4-12
UPPER function	Converts lowercase letters to uppercase	"Case-Conversion Functions" on page 4-156	"UPPER Function" on page 4-157
User-defined aggregate (IDS)	Aggregate that you define (as opposed to built-in aggregates that Dynamic Server provides)	"User-Defined Aggregates (IDS)" on page 4-175	"User-Defined Aggregates (IDS)" on page 4-175
User-defined function	Function that you write (as opposed to built-in functions that the database server provides)	"User-Defined Functions" on page 4-161	"User-Defined Functions" on page 4-161
USER operator	Returns the login name of the current user	"Constant Expressions" on page 4-70	"USER Operator" on page 4-72
Variable	Host or program variable that stores a value	"Syntax of SQL Expressions" on page 4-40	"Syntax of SQL Expressions" on page 4-40
VARIANCE function	Returns the variance for a set of values	"Aggregate Expressions" on page 4-165	"VARIANCE Function" on page 4-173
WEEKDAY function	Returns an integer code for the day of the week	"Time Functions" on page 4-128	"WEEKDAY Function" on page 4-131
YEAR function	Returns a 4-digit integer representing a year	"Time Functions" on page 4-128	"YEAR Function" on page 4-131
* symbol	See "Multiplication (*) operator"	"Syntax of SQL Expressions" on page 4-40	"Arithmetic Operators" on page 4-51

Name	Description	Syntax	Usage
+ symbol	See "Addition" and "Unary plus (+) sign"	"Syntax of SQL Expressions" on page 4-40	"Arithmetic Operators"
- symbol	See "Subtraction" and "Unary minus (-) sign"	"Syntax of SQL Expressions" on page 4-40	"Arithmetic Operators"
/ symbol	See "Division operator"	"Syntax of SQL Expressions" on page 4-40	"Arithmetic Operators"
:: symbols	See "Double-colon (::) cast operator"	"Cast Expressions (IDS)" on page 4-58	"Cast Expressions (IDS)" on page 4-58
symbol	See "Double-pipe () concatenation operator"	"Syntax of SQL Expressions" on page 4-40	"Concatenation Operator" on page 4-56
[<i>first, last</i>] symbols	See "Substring operator"	"Column Expressions" on page 4-59	"Using the Substring Operator" on page 4-63

Sections that follow describe the syntax and usage of each expression that appears in the preceding table.

Arithmetic Operators

Binary arithmetic operators can combine expressions that return numbers.

Arithmetic Operation	Arithmetic Operator	Operator Function	Arithmetic Operation	Arithmetic Operator	Operator Function
Addition	+	plus()	Multiplication	*	times()
Subtraction	-	minus()	Division	/	divide()

The following examples use binary arithmetic operators:

```
quantity * total_price
price * 2
COUNT(*) + 2
```

If you combine a DATETIME value with one or more INTERVAL values, all the fields of the INTERVAL value must be present in the DATETIME value; no implicit EXTEND function is performed. In addition, you cannot use YEAR to MONTH intervals with DAY to SECOND intervals. For additional information about binary arithmetic operators, see the *IBM Informix Guide to SQL: Reference*.

The binary arithmetic operators have associated operator functions, as the preceding table shows. Connecting two expressions with a binary operator is equivalent to invoking the associated operator function on the expressions. For example, the following two statements both select the product of the **total_price** column and 2. In the first statement, the * operator implicitly invokes the **times()** function.

```
SELECT (total_price * 2) FROM items
  WHERE order_num = 1001;
SELECT times(total_price, 2) FROM items
  WHERE order_num = 1001;
```

You cannot use arithmetic operators to combine expressions that use aggregate functions with column expressions.

The database server provides the operator functions associated with the relational operators for all built-in data types. You can define new versions of these operator functions to handle your own user-defined data types.

For more information, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

The database server also supports the following unary arithmetic operators.

Sign of Number	Unary Arithmetic Operator	Operator Function
Positive	+	positive()
Negative	-	negate()

The unary arithmetic operators have the associated operator functions that the preceding table shows. You can define new versions of these functions to handle your own user-defined data types. For more information on this topic, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

If any value that participates in an arithmetic expression is NULL, the value of the entire expression is NULL, as the following example shows:

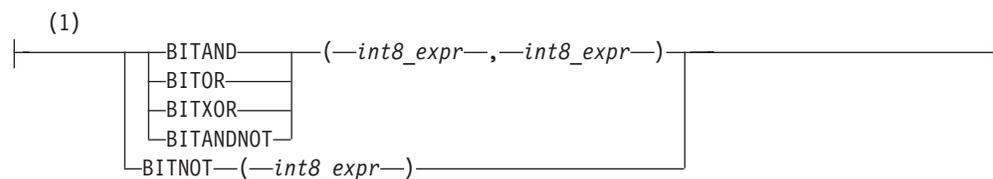
```
SELECT order_num, ship_charge/ship_weight FROM orders
WHERE order_num = 1023;
```

If either **ship_charge** or **ship_weight** is NULL, the value returned for the expression **ship_charge/ship_weight** is also NULL. If the NULL expression **ship_charge/ship_weight** is used in a condition, its truth value cannot be TRUE, and the condition is not satisfied (unless the NULL expression is an operand of the IS NULL operator).

Bitwise Logical Functions

Use the bitwise logical functions to perform named bit operations.

Bitwise Logical Functions:



Notes:

- 1 Informix extension

Element	Description	Restrictions	Syntax
<i>int8_expr</i>	Number expression that can be converted to an INT8 value	For BITNOT the maximum size is reduced by 1	"Expression" on page 4-40

The arguments to these functions can be any numeric data type that can be converted to the INT8 data type.

Except for **BITNOT**, which takes a single argument, these bitwise logical functions take two arguments that can be converted to an INT8 value.

If both arguments have the same integer types, the data type of the returned value is the same type as the arguments. If the two arguments are of different integer types, the returned value is the integer type with the greater precision. For example, if the first argument is of type INT, and the second argument is of type INT8, the returned value is of type INT8.

If the arguments are any other numeric type, such as DECIMAL, SMALLFLOAT, FLOAT, or MONEY, or some combination of those types, the returned data type is DECIMAL(32).

If using host variables, and the data types of the arguments are not known at prepare time, the data type INTEGER is assumed for both arguments, and the returned value is INTEGER. If, after prepare, at execution time, a different data type value is supplied for the host variable, Dynamic Server issues a -9750 error. To prevent such an occurrence, you can specify the host variable data type by using a cast, as in the following ESQL/C program fragment:

```
sprintf(query1, "  
    bitand( ?::int8, ?::int8) from mytab");  
EXEC SQL prepare selectq from :query;  
EXEC SQL declare select_cursor cursor for selectq;  
EXEC SQL open select_cursor  
    using :hostvar_int8_input1, :hostvar_int8_input2;  
  
EXEC SQL fetch select_cursor into :var_int8_output;
```

BITAND Function

The **BITAND** function takes two arguments. The arguments can be any number type value that can be converted to an INT8 value.

Fractional values are truncated before the bit operation. The result is the AND for the two arguments.

If both arguments have the same integer types, the data type of the returned value is the same type as the arguments. If the arguments are of different integer types (for example, INT and INT8), the data type with the greater precision is returned. If the arguments are any other numeric type, such as DECIMAL, SMALLFLOAT, FLOAT, or MONEY, or some combination of those types, the returned data type is DECIMAL(32).

The following example illustrates a query that calls the **BITAND** function:

```
select task_id, task_status,  
    decode(bitand(task_status,1), 1, ' Y', ' N') as task_a,  
    decode(bitand(task_status,2), 2, ' Y', ' N') as task_b,  
    decode(bitand(task_status,4), 4, ' Y', ' N') as task_c  
from tasks;
```

The following table shows the output of this SELECT statement.

task_id	task_status	task_a	task_b	task_c
100	1	Y	N	N
101	1	Y	N	N
102	2	N	Y	N
103	4	N	N	Y
104	6	N	Y	Y
105	3	Y	Y	N
106	5	Y	N	Y
107	7	Y	Y	Y

BITOR Function

The **BITOR** function takes two arguments. The arguments can be any number type value that can be converted to an INT8 value.

Fractional values are truncated before the bit operation. The result is the bitwise OR of its two arguments.

If both arguments have the same integer types, the data type of the returned value is the same type as the arguments. If the arguments are of different integer types (for example, INT and INT8), the returned type is the type with the greater precision. If the arguments are any other numeric type, such as DECIMAL, SMALLFLOAT, FLOAT, or MONEY, or some combination of those types, the returned data type is DECIMAL(32)

The following example illustrates a query that calls the **BITOR** function:

```
SELECT BITOR(8, 20) AS bitor FROM systables WHERE tabid = 1;
```

The following table shows the output of this SELECT statement.

bitor
28

BITXOR Function

The **BITXOR** function takes two arguments. The arguments can be any number type value that can be converted to an INT8 value.

Fractional values are truncated before the bit operation. The result is the bitwise XOR of its two arguments.

If both arguments have the same integer types, the data type of the returned value is the same type as the arguments. If the arguments are of different integer types (for example, INT and INT8), the returned type is the type with the greater precision. If the arguments are any other numeric type, such as DECIMAL, SMALLFLOAT, FLOAT, or MONEY, or some combination of those types, the returned data type is DECIMAL(32).

The following example illustrates a query that calls the **BITXOR** function:

```
SELECT BITXOR(41, 33) AS bitxor FROM systables WHERE tabid = 1;
```

The following table shows the output of this SELECT statement.

bitxor
8

This query calls the **BITXOR** function with negative arguments:

```
SELECT BITXOR(-20, -41) AS bitxor FROM systables WHERE tabid = 1;
```

The following table shows the output of this SELECT statement.

bitxor
59

BITANDNOT Function

The **BITANDNOT** function takes two arguments. The arguments can be any number type value that can be converted to an INT8 value.

Fractional values are truncated before the bit operation. The result is the same as **BITAND**(arg1, **BITNOT**(arg2)) for the two arguments.

If both arguments have the same integer types, the data type of the returned value is the same type as the arguments. If the arguments are of different integer types (for example, INT and INT8), the returned type is the type with the greater precision. If the arguments are any other numeric type, such as DECIMAL, SMALLFLOAT, FLOAT, or MONEY, or some combination of those types, the returned data type is DECIMAL(32).

The query in the following example calls the **BITANDNOT** function:

```
SELECT BITANDNOT(20,-20) AS bitandnot FROM systables WHERE tabid = 1;
```

The following table shows the output of this SELECT statement.

bitandnot
16

The following query calls the equivalent **BITAND** and **BITNOT** functions for the arguments in the previous example:

```
select bitand(20, bitnot(-20)) as bitandnot from systables
  where tabid = 1;
```

The following table shows the output of this SELECT statement.

bitandnot
16

BITNOT Function

The **BITNOT** function can take any number type value that is one less than the maximum INT8 value.

Fractional values are truncated before the bit operation. The result is the bitwise NOT of its argument.

The returned data type is the same type as the argument if the argument is SMALLINT, INT, BIGINT, or INT8. Otherwise the returned data type is DECIMAL(32).

The following query calls the **BITNOT** function:

```
SELECT BITNOT(8) AS bitnot FROM systables WHERE tabid = 1;
```

The following table shows the output of this SELECT statement.

bitnot
-9

The next query calls the **BITNOT** function with a negative argument:

```
SELECT BITNOT(-20) AS bitnot FROM systables WHERE tabid = 1;
```

The following table shows the output of this SELECT statement.

bitnot
19

Concatenation Operator

The concatenation operator is a binary operator, whose syntax is shown in the general diagram for an SQL “Expression” on page 4-40. You can use the concatenation operator (||) to concatenate two expressions that evaluate to character data types or to numeric data types. These examples show some possible concatenated expression combinations.

- The first example concatenates the **zipcode** column to the first three letters of the **Iname** column.
- The second example concatenates the suffix **.dbg** to the contents of a host variable called **file_variable**.
- The third example concatenates the value that the **TODAY** operator returns to the string **Date**.

```
Iname[1,3] || zipcode  
:file_variable || '.dbg'  
'Date:' || TODAY
```

You cannot use the concatenation operator in the following embedded-language statements:

- ALLOCATE COLLECTION
- ALLOCATE DESCRIPTOR
- ALLOCATE ROW
- CREATE FUNCTION FROM
- CREATE PROCEDURE FROM
- CREATE ROUTINE FROM
- DEALLOCATE COLLECTION

- DEALLOCATE DESCRIPTOR
- DEALLOCATE ROW DESCRIBE
- DESCRIBE INPUT
- EXECUTE
- FLUSH
- GET DESCRIPTOR
- GET DIAGNOSTICS
- PUT
- SET AUTOFREE
- SET CONNECTION
- SET DESCRIPTOR
- WHENEVER

Except as noted for the DECLARE and PREPARE statement, routines written in external languages, such as the Informix ESQL/C language, cannot use the concatenation operator in the following dynamic SQL statements:

- CLOSE
- DECLARE
- EXECUTE IMMEDIATE
- FETCH
- FREE
- OPEN
- PREPARE

Although input parameters of the DECLARE statement, such as a *cursor_id* specification, cannot be expressions that include the concatenation operator, Informix ESQL/C routines can use this operator in a SELECT, INSERT, EXECUTE FUNCTION, or EXECUTE PROCEDURE statement within the DECLARE statement.

Informix ESQL/C routines can use the concatenation operator in the text of the SQL statement or statements that you pass to the PREPARE statement.

In SPL routines, you can include the concatenation operator in an expression that specifies the text of the SQL statement that you pass to the EXECUTE IMMEDIATE statement or to the PREPARE statement, even if the calling context of the SPL routine is an Informix ESQL/C routine.

You cannot use the concatenation operator directly with user-defined data types, with complex or large-object data types, nor with operands that are not built-in character or number data types. You must explicitly cast UDTs or other unsupported data types to a built-in character or numeric data type before you can pass the result to the concatenation operator.

The data type of the result of a concatenation operation depends of the data types of the operands and on the length of the resulting string, using the return type promotion rules that the section Return Types from the CONCAT Function describes.

The concatenation operator (||) has an associated operator function called **CONCAT**. The **CONCAT** function cannot be overloaded.

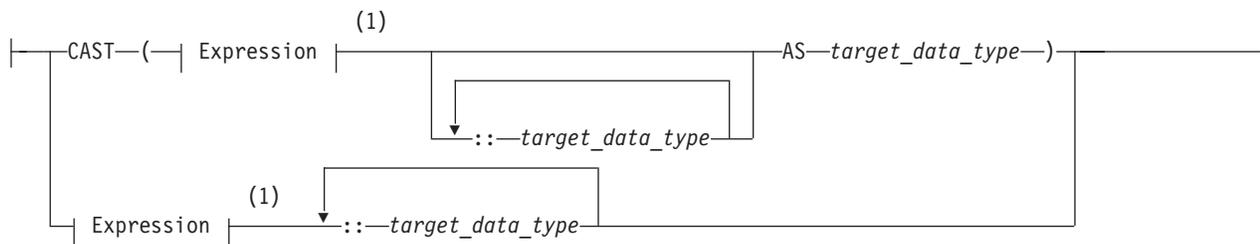
When you define a text-based UDT, you can define a **CONCAT** function to concatenate objects of that user-defined data type. For more information, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

Cast Expressions (IDS)

You can use the **CAST** and **AS** keywords or the double-colon cast operator (**::**) to cast an expression to another data type. Both the operator and the keywords invoke a cast from the data type of the expression to the specified target data type.

To invoke an explicit cast, you can use either the cast operator or the **CAST AS** keywords. If you use the cast operator or the **CAST** and **AS** keywords, but no explicit or implicit cast was defined to perform the conversion between two data types, the statement returns an error.

Cast Expressions:



Notes:

- 1 See "Expression" on page 4-40

Element	Description	Restrictions	Syntax
<i>target_data_type</i>	Data type returned by cast	See "Rules for the Target Data Type"	"Data Type " on page 4-21

Rules for the Target Data Type

The following rules restrict the *target data type* in cast expressions:

- The target data type must be either a built-in type, a user-defined type, or a named row type in the database.
- The target data type cannot be an unnamed row or a collection type.
- The target data type can be a BLOB data type under the following conditions:
 - The source expression (the expression to be cast to another data type) is a BYTE data type.
 - The source expression is a user-defined type and the user has defined a cast from the user-defined type to the BLOB type.
- The target data type can be a CLOB type under these conditions:
 - The source expression is a TEXT data type.
 - The source expression is a user-defined type and the user has defined a cast from the user-defined type to the CLOB type.
- You cannot cast a BLOB data type to a BYTE data type.
- You cannot cast a CLOB data type to a TEXT data type.
- An explicit or implicit cast must exist that can convert the data type of the source expression to the target data type.

Examples of Cast Expressions

The following examples show two different ways to convert the sum of x and y to a user-defined data type, `user_type`. The two methods produce identical results. Both require the existence of an explicit or implicit cast from the type returned by $(x + y)$ to the user-defined type:

```
CAST ((x + y) AS user_type)
(x + y)::user_type
```

The following examples show two different ways of finding the integer equivalent of the expression `expr`. Both require the existence of an implicit or explicit cast from the data type of `expr` to the `INTEGER` data type:

```
CAST (expr AS INTEGER)
expr::INTEGER
```

In the following example, the user casts a `BYTE` column to the `BLOB` type and copies the `BLOB` data to an operating-system file:

```
SELECT LOTOFILE(mybytecol::blob, 'fname', 'client')
FROM mytab
WHERE pkey = 12345;
```

In the following example, the user casts a `TEXT` column to a `CLOB` value and then updates a `CLOB` column in the same table to have the `CLOB` value derived from the `TEXT` column:

```
UPDATE newtab SET myclobcol = mytextcol::clob;
```

The Keyword NULL in Cast Expressions

Cast expressions can appear in the projection list, including expressions of the form `NULL::datatype`, where `datatype` is any data type known to the database:

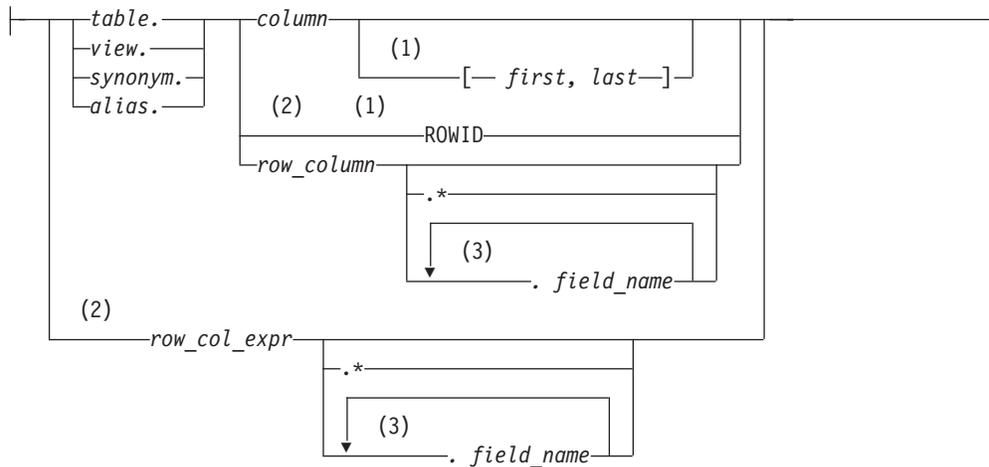
```
SELECT newtable.col0, null::int FROM newtable;
```

The keyword `NULL` has a global scope of reference within expressions. In SQL, the keyword `NULL` is the only syntactic mechanism for accessing a `NULL` value. Any attempt to redefine or restrict the global scope of the keyword `NULL` (for example, declaring an SPL variable called `null`) disables any cast expression that involves a `NULL` value. Make sure that the keyword `NULL` receives its global scope in all expression contexts.

Column Expressions

A *column expression* specifies a data value in a column in the database, or a substring of the value, or (for Dynamic Server only) a field within a `ROW`-type column. This is the syntax for column expressions.

Column Expressions:



Notes:

- 1 Informix extension
- 2 Dynamic Server only
- 3 Use path no more than three times

Element	Description	Restrictions	Syntax
<i>alias</i>	Temporary alternative name for a table or view, declared in the FROM clause of a query	Must return a string. Restrictions depend on the clause of the SELECT statement in which <i>alias</i> occurs	"Identifier " on page 5-21
<i>column</i>	Name of a column	Restrictions depend on the SQL statement where <i>column</i> occurs	"Identifier " on page 5-21
<i>field_name</i>	Name of a ROW field in the ROW column or ROW-column expression	Must be a member of the row that <i>row-column name</i> or <i>row_col_expr</i> or <i>field name</i> (for nested rows) specifies	"Identifier " on page 5-21
<i>first, last</i>	Integers indicating positions of first and last characters within <i>column</i>	The <i>column</i> must be of type CHAR, VARCHAR, NCHAR, NVARCHAR, BYTE, or TEXT, and $0 < first \leq last$	"Literal Number " on page 4-185
<i>row_col_expr</i>	Expression that returns ROW-type values	Must return a ROW data type	"Expression" on page 4-40
<i>row_column</i>	Name of a ROW-type column	Must be a named ROW data type or an unnamed ROW data type	"Identifier " on page 5-21
<i>synonym, table, view</i>	Table, view, or synonym (for the table or view) that contains <i>column</i>	Synonym and the table or view to which it points must exist	Database Object Name, p. "Database Object Name " on page 5-16

The following examples show column expressions:

company

items.price

cat_advert [1,15]

You must qualify the *column* name with a *table* name or *alias* whenever it is necessary to distinguish between columns that have the same name but are in different tables. The SELECT statements that the following example shows use

customer_num from the **customer** and **orders** tables. The first example precedes the column names with table names. The second example precedes the column names with table aliases.

```
SELECT * FROM customer, orders
    WHERE customer.customer_num = orders.customer_num;

SELECT * FROM customer c, orders o
    WHERE c.customer_num = o.customer_num;
```

Using Dot Notation

Dot notation (sometimes called the *membership operator*) allows you to qualify an SQL identifier with another SQL identifier of which it is a component. You separate the identifiers with the period (.) symbol. For example, you can qualify a column name with any of the following SQL identifiers:

- Table name: *table_name.column_name*
- View name: *view_name.column_name*
- Synonym name: *syn_name.column_name*

These forms of dot notation are called *column projections*.

You can also use dot notation to directly access the fields of a named or unnamed ROW column, as in the following example:

row-column name.field name

This use of dot notation is called a *field projection*. For example, suppose you have a column called **rect** with the following definition:

```
CREATE TABLE rectangles
(
    area float,
    rect ROW(x int, y int, length float, width float)
);
```

The following SELECT statement uses dot notation to access field **length** of the **rect** column:

```
SELECT rect.length FROM rectangles
    WHERE area = 64;
```

Selecting All Fields of a Column with Asterisk Notation: If you want to select all fields of a column that has a ROW type, you can specify the column name without dot notation. For example, you can select all fields of the **rect** column as follows:

```
SELECT rect FROM rectangles
    WHERE area = 64;
```

You can also use asterisk (*) notation to project all the fields of a column that has a ROW data type. For example, if you want to use asterisk notation to select all fields of the **rect** column, you can enter the following statement:

```
SELECT rect.* FROM rectangles
    WHERE area = 64;
```

Asterisk notation is easier than specifying each field of the **rect** column individually:

```
SELECT rect.x, rect.y, rect.length, rect.width
    FROM rectangles
    WHERE area = 64;
```

Asterisk notation for ROW fields is valid only in the projection list of a SELECT statement. It can specify all fields of a ROW-type column or the data that a ROW-column expression returns.

Asterisk notation is not necessary with ROW-type columns because you can specify the column name alone to project all of its fields. Asterisk notation is quite helpful, however, with ROW-type expressions such as subqueries and user-defined functions that return ROW-type values. For more information, see “Using Dot Notation with Row-Type Expressions.”

You can use asterisk notation with columns and expressions of ROW data types in the projection list of a SELECT statement only. You cannot use asterisk notation with columns and expressions of ROW type in any other clause of a SELECT statement.

Selecting Nested Fields: When the ROW type that defines a column itself contains other ROW types, the column contains nested fields. Use dot notation to access these nested fields within a column.

For example, assume that the **address** column of the **employee** table contains the fields: **street**, **city**, **state**, and **zip**. In addition, the **zip** field contains the nested fields: **z_code** and **z_suffix**. A query on the **zip** field returns values for the **z_code** and **z_suffix** fields. You can specify, however, that a query returns only specific nested fields. The following example shows how to use dot notation to construct a SELECT statement that returns rows for the **z_code** field of the **address** column only:

```
SELECT address.zip.z_code
FROM employee;
```

Rules of Precedence: The database server uses the following precedence rules to interpret dot notation:

1. schema *name_a* . table *name_b* . column *name_c* . field *name_d*
2. table *name_a* . column *name_b* . field *name_c* . field *name_d*
3. column *name_a* . field *name_b* . field *name_c* . field *name_d*

When the meaning of an identifier is ambiguous, the database server uses precedence rules to determine which database object the identifier specifies. Consider the following two tables:

```
CREATE TABLE b (c ROW(d INTEGER, e CHAR(2)));
CREATE TABLE c (d INTEGER);
```

In the following SELECT statement, the expression **c.d** references column **d** of table **c** (rather than field **d** of column **c** in table **b**) because a table identifier has a higher precedence than a column identifier:

```
SELECT *
FROM b,c
WHERE c.d = 10;
```

For more information about precedence rules and how to use dot notation with ROW columns, see the *IBM Informix Guide to SQL: Tutorial*.

Using Dot Notation with Row-Type Expressions: Besides specifying a column of a ROW data type, you can also use dot notation with any expression that evaluates to a ROW type. In an INSERT statement, for example, you can use dot notation in a subquery that returns a single row of values. Assume that you created a ROW type named **row_t**:

```
CREATE ROW TYPE row_t (part_id INT, amt INT);
```

Also assume that you created a typed table named **tab1** that is based on the **row_t** ROW type:

```
CREATE TABLE tab1 OF TYPE row_t;
```

Assume also that you inserted the following values into table **tab1**:

```
INSERT INTO tab1 VALUES (ROW(1,7));  
INSERT INTO tab1 VALUES (ROW(2,10));
```

Finally, assume that you created another table named **tab2**:

```
CREATE TABLE tab2 (colx INT);
```

Now you can use dot notation to insert the value from only the **part_id** column of table **tab1** into the **tab2** table:

```
INSERT INTO tab2  
VALUES ((SELECT t FROM tab1 t  
WHERE part_id = 1).part_id);
```

The asterisk form of dot notation is not necessary when you want to select all fields of a ROW-type column because you can specify the column name alone to select all of its fields. The asterisk form of dot notation can be quite helpful, however, when you use a subquery, as in the preceding example, or when you call a user-defined function to return ROW-type values.

Suppose that a user-defined function named **new_row** returns ROW-type values, and you want to call this function to insert the ROW-type values into a table. Asterisk notation makes it easy to specify that all the ROW-type values that the **new_row()** function returns are to be inserted into the table:

```
INSERT INTO mytab2 SELECT new_row (mycol).* FROM mytab1;
```

References to the fields of a ROW-type column or a ROW-type expression are not allowed in fragment expressions. A fragment expression is an expression that defines a table fragment or an index fragment in SQL statements like CREATE TABLE, CREATE INDEX, and ALTER FRAGMENT.

Using the Substring Operator

You can use the substring operator on CHAR, VARCHAR, NCHAR, NVARCHAR, LVARCHAR, BYTE, and TEXT columns to define a *column substring* as the portion of the column that is specified by the expression. After the identifier of a character column, when a pair of bracket ([]) symbols enclose a comma-separated pair of unsigned integers in which the *first* integer is greater than zero but not greater than the *last* integer, Dynamic Server interprets the brackets as the substring operator. The expression returns the *first* through *last* characters of the data value in the column, where *first* and *last* define a substring. For example, in the expression `cat_advert [6,15]`, the returned value is the 6th through 15th characters of column **cat_advert**.

In the default locale, if the data value occupies at least 15 bytes, this expression evaluates to a substring that includes ten bytes of the column value, but in a multibyte locale this expression returns a string of ten consecutive logical characters whose storage length might exceed 10 bytes, beginning with the sixth logical character. For more information on the GLS aspects of column substrings, see the *IBM Informix GLS User's Guide*.

In the following example, if a value in the **lname** column of the **customer** table is Greenburg, the following expression evaluates to burg:

```
lname[6,9]
```

A conditional expression can include a column expression that uses the substring operator ([*first*, *last*]), as in the following example:

```
SELECT lname FROM customer WHERE phone[5,7] = '356';
```

Here the quotation marks are required, to prevent the database server from applying a numeric filter to the digits in the criterion value.

See also the section “String-Manipulation Functions” on page 4-141, which describes two built-in SQL functions, **SUBSTR()** and **SUBSTRING()** that can specify a substring expression within an SQL statement.

Using Rowids (IDS)

In Dynamic Server, you can use the **rowid** column that is associated with a table row as a property of the row. The **rowid** column is essentially a hidden column in nonfragmented tables and in fragmented tables that were created with the WITH ROWIDS clause. The **rowid** column is unique for each row, but it is not necessarily sequential. It is recommended, however, that you use primary keys as an access method rather than exploiting the **rowid** column.

The following examples use the ROWID keyword in a SELECT statement:

```
SELECT *, ROWID FROM customer;
```

```
SELECT fname, ROWID FROM customer ORDER BY ROWID;
```

```
SELECT HEX(rowid) FROM customer WHERE customer_num = 106;
```

The last example shows how to get the page number (the first six digits after 0x) and the slot number (the last two digits) of the location of your row.

You cannot use the ROWID keyword in the select list of the Projection clause of a query that contains an aggregate function.

Using Smart Large Objects (IDS)

The SELECT, UPDATE, and INSERT statements do not manipulate the values of smart large objects directly. Instead, they use a *handle value*, which is a type of pointer, to access the BLOB or CLOB value, as follows:

- The SELECT statement returns a handle value to the BLOB or CLOB value that the projection list specifies. SELECT does not return the actual data for the BLOB or CLOB column that the projection list specifies. Instead, it returns a handle value to the column data.
- The INSERT and UPDATE statements do not send the actual data for the BLOB or CLOB column to the database server. Instead, they accept a handle value to this data as the value to be inserted or updated.

To access the data of a smart-large-object column, you must use one of the following application programming interfaces (APIs):

- From within an IBM Informix ESQL/C program, use the Informix ESQL/C library functions that access smart large objects. For more information, see the *IBM Informix ESQL/C Programmer's Manual*.

- From within a C program such as a DataBlade module, use the Client and Server API. For more information, see your *IBM Datablade Developers Kit User's Guide*.

You cannot use the name of a smart-large-object column in expressions that involve arithmetic operators. For example, operations such as addition or subtraction on the smart-large-object handle value have no meaning.

When you select a smart-large-object column, you can assign the handle value to any number of columns: all columns with the same handle value share the CLOB or BLOB value. This storage arrangement reduces the amount of disk space that the CLOB or BLOB value, but when several columns share the same smart-large-object value, the following conditions result:

- The chance of lock contention on a CLOB or BLOB column increases. If two columns share the same smart-large-object value, the data might be locked by either column that needs to access it.
- The CLOB or BLOB value can be updated from a number of points.

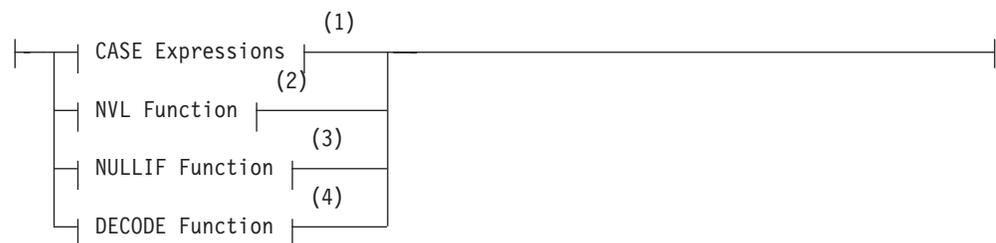
To remove these constraints, you can create separate copies of the BLOB or CLOB data for each column that needs to access it. You can use the **LOCOPY** function to create a copy of an existing smart large object.

You can also use the built-in functions **LOTOFILE**, **FILETOCLOB**, and **FILETOBLOB** to access smart-large-object values, as described in “Smart-Large-Object Functions (IDS)” on page 4-122. For more information on the BLOB and CLOB data types, see the *IBM Informix Guide to SQL: Reference*.

Conditional Expressions

Conditional expressions return values that depend on the outcome of conditional tests. This diagram shows the syntax for Conditional Expressions.

Conditional Expressions:



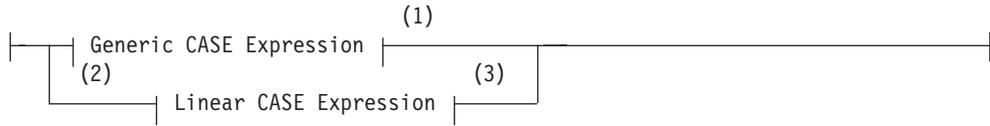
Notes:

- 1 See “CASE Expressions” on page 4-66
- 2 See “ NVL Function” on page 4-68
- 3 See “ NULLIF Function” on page 4-68
- 4 See “ DECODE Function” on page 4-69

CASE Expressions

The CASE expression allows an SQL statement such as the SELECT statement to return one of several possible results, depending on which of several conditions evaluates to true. The CASE expression has two forms: generic CASE expressions and linear CASE expressions.

CASE Expressions:



Notes:

- 1 See "Generic CASE Expressions"
- 2 Dynamic Server only
- 3 See "Linear CASE Expressions (IDS)" on page 4-67

You must include at least one WHEN clause in the CASE expression. Subsequent WHEN clauses and the ELSE clause are optional. You can use a generic or linear CASE expression wherever you can use a column expression in an SQL statement (for example, in the Projection clause a SELECT statement).

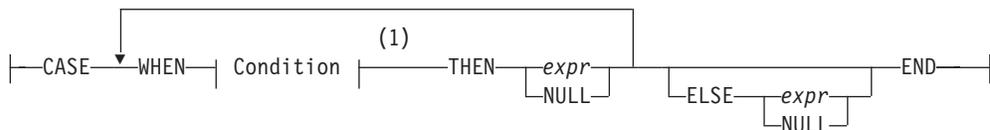
Expressions in the search condition or the result value expression can contain subqueries, and you can nest a CASE expression in another CASE expression. When a CASE expression appears in an aggregate expression, you cannot use aggregate functions in the CASE expression.

You can specify a trigger-type Boolean operator (DELETING, INSERTING, SELECTING, or UPDATING) as a condition in a CASE expression only within a trigger routine.

Generic CASE Expressions:

A generic CASE expression tests for a true condition in a WHEN clause. If it finds a true condition, it returns the result specified in the THEN clause.

Generic CASE Expression:



Notes:

- 1 See "Condition " on page 4-5

Element	Description	Restrictions	Syntax
<i>expr</i>	Expression that returns some data type	Data type of <i>expr</i> in a THEN clause must be compatible with data types of expressions in other THEN clauses	"Expression" on page 4-40

The database server processes the WHEN clauses in the order that they appear in the statement. If the search condition of a WHEN clause evaluates to TRUE, the database server uses the value of the corresponding THEN expression as the result, and stops processing the CASE expression.

If no WHEN condition evaluates to TRUE, the database server uses the ELSE expression as the overall result. If no WHEN condition evaluates to TRUE, and no ELSE clause was specified, the returned CASE expression value is NULL. You can use the IS NULL condition to handle NULL results. For information on how to handle NULL values, see “IS NULL and IS NOT NULL Conditions” on page 4-12.

The next example shows a generic CASE expression in the Projection clause.

In this example, the user retrieves the name and address of each customer as well as a calculated number that is based on the number of problems that exist for that customer:

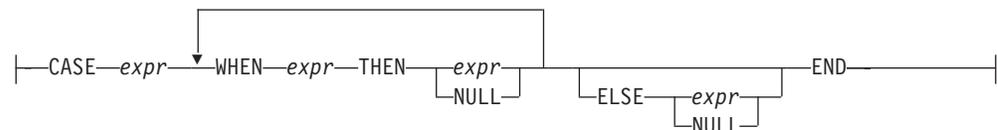
```
SELECT cust_name,
       CASE
         WHEN number_of_problems = 0
           THEN 100
         WHEN number_of_problems > 0 AND number_of_problems < 4
           THEN number_of_problems * 500
         WHEN number_of_problems >= 4 and number_of_problems <= 9
           THEN number_of_problems * 400
         ELSE
           (number_of_problems * 300) + 250
       END,
       cust_address
FROM custtab
```

In a generic CASE expression, all the results should be of the same data type, or they should evaluate to a common compatible data type. If the results in all the WHEN clauses are not of the same data type, or if they do not evaluate to values of mutually compatible types, an error occurs.

Linear CASE Expressions (IDS):

A linear CASE expression compares the value of the expression that follows the CASE keyword with an expression in a WHEN clause.

Linear CASE Expression:



Element	Description	Restrictions	Syntax
<i>expr</i>	Expression that returns a value of some data type	Data type of <i>expr</i> that follows the WHEN keyword must be compatible with data type of the expression that follows the CASE keyword. Data type of <i>expr</i> in the THEN clause must be compatible with data types of expressions in other THEN clauses.	“Expression” on page 4-40

The database server evaluates the expression that follows the CASE keyword, and then processes the WHEN clauses sequentially. If an expression after the WHEN keyword returns the same value as the expression that follows the CASE keyword,

the database server uses the value of the expression that follows the THEN keyword as the overall result of the CASE expression. Then the database server stops processing the CASE expression.

If none of the WHEN expressions return the same value as the expression that follows the CASE keyword, the database server uses the expression of the ELSE clause as the overall result of the CASE expression (or, if no ELSE clause was specified, the returned value of the CASE expression is NULL).

The next example shows a linear CASE expression in the projection list of the Projection clause of a SELECT statement. For each movie in a table of movie titles, the query returns the title, the cost, and the type of the movie. The statement uses a CASE expression to derive the type of each movie:

```
SELECT title, CASE movie_type
      WHEN 1 THEN 'HORROR'
      WHEN 2 THEN 'COMEDY'
      WHEN 3 THEN 'ROMANCE'
      WHEN 4 THEN 'WESTERN'
      ELSE 'UNCLASSIFIED'
END,
      our_cost FROM movie_titles;
```

In linear CASE expressions, the data types of WHEN clause expressions must be compatible with that of the expression that follows the CASE keyword.

NVL Function

The NVL expression returns different results, depending on whether its first argument evaluates to NULL.

NVL Function:

|—NVL—(—*expr1*—,—*expr2*—)|

Element	Description	Restrictions	Syntax
<i>expr1 expr2</i>	Expressions that return values of a compatible data type	Cannot be a host variable or a BYTE or TEXT data type	"Expression" on page 4-40

NVL evaluates *expression1*. If *expression1* is not NULL, then NVL returns the value of *expression1*. If *expression1* is NULL, NVL returns the value of *expression2*. The expressions *expression1* and *expression2* can be of any data type, as long as they can be cast to a common compatible data type.

Suppose that the **addr** column of the **employees** table has NULL values in some rows, and the user wants to be able to print the label Address unknown for these rows. The user enters the following SELECT statement to display the label Address unknown when the **addr** column has a NULL value:

```
SELECT fname, NVL (addr, 'Address unknown') AS address
      FROM employees;
```

NULLIF Function

The NULLIF expression returns different results, depending on whether its two arguments are equal.

NULLIF Function:

|—NULLIF—(—*expr1*—,—*expr2*—)|

Element	Description	Restrictions	Syntax
<i>expr1 expr2</i>	Expressions that return values of a compatible data type	Cannot be a BYTE or TEXT data type	“Expression” on page 4-40

NULLIF evaluates its two arguments, *expr1* and *expr2*.

- If their values are equal, then NULLIF returns NULL.
- If their values are not equal, then NULLIF returns *expr1*.

The *expr1* and *expr2* arguments can be of any data type for which a built-in comparison function exists, or any two data types that can be cast to a compatible data type that has a built-in comparison function.

The following example uses the NULLIF function to convert Boolean FALSE values ('f') to NULL values:

```
SELECT name, answer, NULLIF(answer, 'f') FROM booktab;
```

Here the first argument is a Boolean column expression that can have true ('t') or false ('f') values, and the second Boolean argument is always 'f' (for FALSE). For rows that have 'f' in the **answer** column, the value returned by the NULLIF function will be NULL (because the NULL value is returned when the arguments are equal). For rows that have 't' as the first argument, however, the value returned by NULLIF is always 't', because the two arguments cannot be equal when one is 't' and the other is 'f'; the first argument is returned when the two values are not equal.

DECODE Function

The DECODE expression is similar to the CASE expression in that it can print different results depending on the values found in a specified column.

DECODE Function:

|—DECODE—(—*expr*—,—*when_expr*—,—*then_expr*—,—*else_expr*—)|

Element	Description	Restrictions	Syntax
<i>expr, else_expr, then_expr, when_expr</i>	Expressions whose values and data types can be evaluated	Data types of <i>when_expr</i> and <i>expr</i> must be compatible, as must <i>then_expr</i> and <i>else_expr</i> . Value of <i>when_expr</i> cannot be a NULL.	“Expression” on page 4-40

The expressions *expr*, *when_expr*, and *then_expr* are required. DECODE evaluates *expr* and compares it to *when_expr*. If the value of *when_expr* matches the value of *expr*, then DECODE returns *then_expr*.

The expressions *when_expr* and *then_expr* are an expression pair, and you can specify any number of expression pairs in the DECODE function. In all cases,

DECODE compares the first member of the pair against *expr* and returns the second member of the pair if the first member matches *expr*.

If no expression matches *expr*, **DECODE** returns *else_expr*. If no expression matches *expr* and you specified no *else_expr*, then **DECODE** returns NULL.

You can specify any data type for the arguments, but two restrictions exist:

- All instances of *when_expr* must have the same data type, or a common compatible type must exist. All instances of *when_expr* must also have the same (or a compatible) data type as *expr*.
- All instances of *then_expr* must have the same data type, or a common compatible type must exist. All instances of *then_expr* must also have the same (or a compatible) data type as *else_expr*.

Suppose that a user wants to convert descriptive values in the **evaluation** column of the **students** table to numeric values in the output. The following table shows the contents of the **students** table.

firstname	evaluation	firstname	evaluation
Edward	Great	Mary	Good
Joe	Not done	Jim	Poor

The user now enters a query with the **DECODE** function to convert the descriptive values in the **evaluation** column to numeric equivalents:

```
SELECT firstname, DECODE(evaluation,
    'Poor', 0,
    'Fair', 25,
    'Good', 50,
    'Very Good', 75,
    'Great', 100,
    -1) as grade
FROM students;
```

The following table shows the output of this SELECT statement.

firstname	evaluation	firstname	evaluation
Edward	100	Mary	50
Joe	-1	Jim	0

Constant Expressions

Certain expressions that return a fixed value are called *constant expressions*. These include variant function operators that read the system clock, but that are valid in contexts where literal constants are also valid. Among these expressions are the following operators (or *system constants*) whose returned values are determined at runtime:

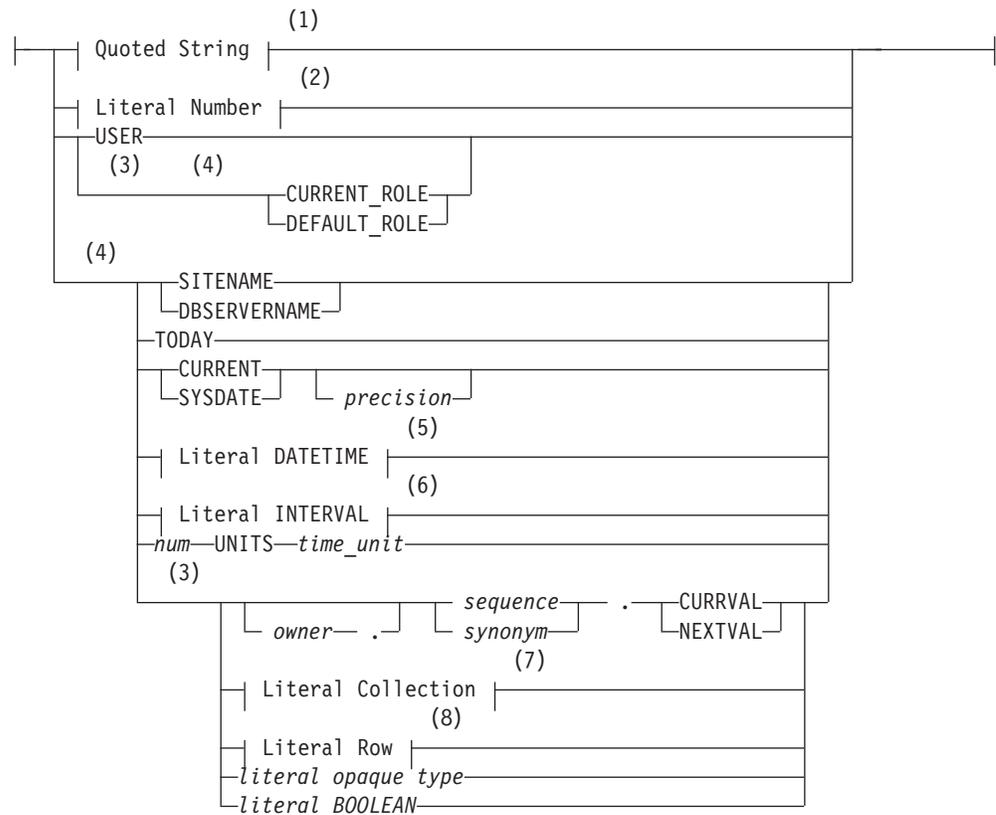
- **CURRENT** returns the current time and date from the system clock.
- **CURRENT_ROLE** returns the name of the role, if any, whose privileges are enabled for the current user.
- **DEFAULT_ROLE** returns the name of the role, if any, that is the default role for the current user.
- **DBSERVERNAME** returns the name of the current database server.

- **SITENAME** is a synonym for **DBSERVERNAME**.
- **SYSDATE** reads the DATETIME value from the system clock like the **CURRENT** operator, but has a different default precision.
- **TODAY** returns the current calendar date from the system clock.
- **USER** returns the login name (also called the *authorization identifier*) of the current user.

Besides these operators, the term *constant expression* can also refer to a quoted string, to a literal value, or to the **UNITS** operator with its operands.

The Constant Expression segment has the following syntax.

Constant Expressions:



Notes:

- 1 See "Quoted String" on page 4-189
- 2 See "Literal Number " on page 4-185
- 3 Dynamic Server only
- 4 Informix extension
- 5 See "Literal DATETIME " on page 4-181
- 6 See "Literal INTERVAL" on page 4-183
- 7 See "Literal Collection" on page 4-178
- 8 See "Literal Row" on page 4-186

Element	Description	Restrictions	Syntax
<i>literal Boolean</i>	Literal representation of a BOOLEAN value	Must be either <i>t</i> (TRUE) or <i>f</i> (FALSE)	“Quoted String” on page 4-189
<i>literal opaque type</i>	Literal representation of value of an opaque data type	Must be recognized by the input support function of opaque type	Defined by UDT developer
<i>num</i>	How many of specified time units. See “UNITS Operator” on page 4-79.	If <i>num</i> is not an integer, the fractional part is truncated	“Literal Number ” on page 4-185
<i>owner</i>	Name of the owner of sequence	Must own sequence	“ Owner Name ” on page 5-45
<i>precision</i>	Precision of the returned DATETIME expression	On Windows systems the maximum scale of seconds is FRACTION(3).	“ DATETIME Field Qualifier” on page 4-38
<i>sequence</i>	Name of a sequence	Must exist in current database	“Identifier ” on page 5-21
<i>synonym</i>	Synonym for the name of a sequence	Must exist in current database	“Identifier ” on page 5-21
<i>time_unit</i>	Keyword to specify time unit: YEAR, MONTH, DAY, HOUR, MINUTE, SECOND, or FRACTION	Must be one of the keywords at left. Case insensitive but cannot be enclosed within quotes	See the Restrictions column.

Quoted String

The following examples show quoted strings as expressions:

```
SELECT 'The first name is ', fname FROM customer;

INSERT INTO manufact VALUES ('SPS', 'SuperSport');

UPDATE cust_calls SET res_dtime = '2007-1-1 10:45'
WHERE customer_num = 120 AND call_code = 'B';
```

For more information, see “Quoted String” on page 4-189.

Literal Number

The following examples show literal numbers as expressions:

```
INSERT INTO items VALUES (4, 35, 52, 'HRO', 12, 4.00);

INSERT INTO acreage VALUES (4, 5.2e4);

SELECT unit_price + 5 FROM stock;

SELECT -1 * balance FROM accounts;
```

For more information, see “Literal Number ” on page 4-185.

USER Operator

The **USER** operator returns a string containing the login name (also called the authorization identifier) of the current user who is running the process.

The following statements show how you might use the **USER** operator:

```
INSERT INTO cust_calls VALUES
(221,CURRENT,USER,'B','Decimal point off', NULL, NULL);
```

```
SELECT * FROM cust_calls WHERE user_id = USER;

UPDATE cust_calls SET user_id = USER WHERE customer_num = 220;
```

The **USER** operator does not change the lettercase of a user ID. If you use **USER** in an expression and the current user is **Robertm**, the **USER** operator returns **Robertm**, not **robertm** or **ROBERTM**.

If you specify **USER** as a default column value, *column* must be of type CHAR, VARCHAR, NCHAR, NVARCHAR, or (in Dynamic Server) LVARCHAR.

If you specify **USER** as the default value for a column, the size of *column* should not be less than 32 bytes. You risk getting an error during operations such as INSERT or ALTER TABLE if the column length is too small to store the default value.

In an ANSI-compliant database, if you do not enclose the *owner* name in quotation marks, the name of the table owner is stored as uppercase letters. If you use the **USER** operator as part of a condition, you must be sure that the way the *user* name is stored matches what the **USER** operator returns with respect to lettercase.

CURRENT_ROLE Operator

The **CURRENT_ROLE** operator returns a string that contains the name of the currently enabled role of the user who is running the session. This *role* was either set in the session explicitly, using the SET ROLE statement, or else implicitly as a default role when the current user connected to the database. If the user holds no role, or if no role that was granted to the user is currently enabled, **CURRENT_ROLE** returns a NULL value. If the user has been granted no role individually, but a default role has been granted to PUBLIC, and this default role has been explicitly or implicitly enabled, **CURRENT_ROLE** returns the name of this default role.

The next statement shows how you might use the **CURRENT_ROLE** operator:

```
select CURRENT_ROLE FROM systables WHERE tabid = 1;
```

The **CURRENT_ROLE** operator does not change the lettercase of the identifier of a role. If you use **CURRENT_ROLE** in an expression and your current role is **Czarina**, the **CURRENT_ROLE** operator returns **Czarina**, not **czarina**.

If you specify **CURRENT_ROLE** as the default value for a column, the column must have a CHAR, VARCHAR, LVARCHAR, NCHAR, or NVARCHAR data type. Because the name of a role is an authorization identifier, truncation might occur if the column length is less than 32 bytes.

DEFAULT_ROLE Operator

The **DEFAULT_ROLE** operator evaluates to a string that contains the name of the default role that has been granted to the user who is running the session. This default role need not be currently enabled, but it must not have been revoked since the most recent GRANT DEFAULT ROLE statement that referenced the user or PUBLIC in the TO clause.

If has no default role is explicitly defined for the current user, but PUBLIC has a default role, **DEFAULT_ROLE** returns the default role of PUBLIC.

If the user has no default role, or if the default role that was most recently granted to the user explicitly, or as PUBLIC, was subsequently revoked by the REVOKE DEFAULT ROLE statement, **DEFAULT_ROLE** returns a NULL value. If the user has been granted no default role individually, but a default role has been granted to PUBLIC, the **DEFAULT_ROLE** operator returns the name of this default role. If no default role is currently defined for the user nor for PUBLIC, however, **DEFAULT_ROLE** returns NULL.

The SET ROLE statement has no effect on the **DEFAULT_ROLE** operator, but any access privileges of the default role are not necessarily available to the user if SET ROLE has activated some other role, or if SET ROLE specified NULL or NONE as the current role of the user.

The next statements show how you might use the **DEFAULT_ROLE** operator:

```
select DEFAULT_ROLE from systables where tabid = 1;
```

DEFAULT_ROLE does not change the lettercase of the identifier of a role.

If you specify **DEFAULT_ROLE** as the default value for a column, the column must have a CHAR, VARCHAR, LVARCHAR, NCHAR, or NVARCHAR data type. Because the name of a role is an authorization identifier, truncation might occur if the column width is less than 32 bytes. (See “ Owner Name ” on page 5-45 for the syntax of authorization identifiers.)

DBSERVERNAME and SITENAME Operators

The **DBSERVERNAME** operator returns the database server name, as defined in the ONCONFIG file for the installation where the current database resides, or as specified in the **INFORMIXSERVER** environment variable. The two operators, **DBSERVERNAME** and **SITENAME**, are synonymous. You can use the **DBSERVERNAME** operator to determine the location of a table, to put information into a table, or to extract information from a table. You can insert **DBSERVERNAME** into a simple character field or use it as a default value for a column.

If you specify **DBSERVERNAME** as a default column value, the column must have a CHAR, VARCHAR, LVARCHAR, NCHAR, or NVARCHAR data type.

If you specify **DBSERVERNAME** or **SITENAME** as the default value for a column, the size of the column should be at least 128 bytes long. You risk getting an error message during INSERT and ALTER TABLE operations if the length of the column is too small to store the default value.

In the following example, the first statement returns the name of the database server where the **customer** table resides. Because the query is not restricted with a WHERE clause, it returns **DBSERVERNAME** for every row in the table. If you add the DISTINCT keyword to the SELECT clause, the query returns **DBSERVERNAME** once. The second statement adds a row that contains the current site name to a table. The third statement returns all the rows that have the site name of the current system in **site_col**. The last statement changes the company name in the **customer** table to the current system name:

```
SELECT DBSERVERNAME FROM customer;
```

```
INSERT INTO host_tab VALUES ('1', DBSERVERNAME);
```

```
SELECT * FROM host_tab WHERE site_col = DBSERVERNAME;

UPDATE customer SET company = DBSERVERNAME
  WHERE customer_num = 120;
```

TODAY Operator

Use the **TODAY** operator to return the system date as a DATE data type. If you specify **TODAY** as a default column value, the column must be a DATE column.

The following examples show how you might use the **TODAY** operator in an INSERT, UPDATE, or SELECT statement:

```
UPDATE orders (order_date) SET order_date = TODAY
  WHERE order_num = 1005;

INSERT INTO orders VALUES
  (0, TODAY, 120, NULL, N, '1AUE217', NULL, NULL, NULL, NULL);

SELECT * FROM orders WHERE ship_date = TODAY;
```

For Extended Parallel Server, the value returned by **TODAY** defaults to the date from the system clock-calendar, but you can specify instead the date from some other time zone. You can specify a time zone by using the SET ENVIRONMENT CLIENT_TZ statement of SQL (or by setting the **IBM_XPS_PARAMS** environment variable) to specify an offset from Greenwich Mean Time (GMT).

For code examples of setting non-default time zones, see “CURRENT Operator.” For information on how to set **IBM_XPS_PARAMS**, see the chapter on environment variables in the *IBM Informix Guide to SQL: Reference*.

CURRENT Operator

The **CURRENT** operator returns a DATETIME value with the date and time of day, showing the current instant.

If you do not specify a DATETIME qualifier, the default qualifier is YEAR TO FRACTION(3). The USEOSTIME configuration parameter specifies whether or not the database server uses subsecond precision when it obtains the current time from the operating system. For more information on the USEOSTIME configuration parameter, see your *IBM Informix Administrator's Reference*.

You can use **CURRENT** in any context where a literal DATETIME is valid. (See “Literal DATETIME ” on page 4-181). If you specify **CURRENT** as the default value for a column, it must be a DATETIME column and the qualifier of **CURRENT** must match the column qualifier, as the following example shows:

```
CREATE TABLE new_acct (co11 INT, co12 DATETIME YEAR TO DAY
  DEFAULT CURRENT YEAR TO DAY);
```

CURRENT is always evaluated in the database server where the current database is located. If the current database is in a remote database server, the returned value is from the remote host.

SQL is not a procedural language, and **CURRENT** might not execute in the lexical order of its position in a statement. You should not use **CURRENT** to mark the start, the end, nor a specific point in the execution of an SQL statement.

If you use the **CURRENT** operator in more than once in a single statement, identical values might be returned by each instance of **CURRENT**. You cannot rely on **CURRENT** to return distinct values each time it executes.

The returned value is based on the system clock and is fixed when the SQL statement that specifies **CURRENT** starts execution. For example, any call to **CURRENT** from inside the SPL function that an **EXECUTE FUNCTION** (or **EXECUTE PROCEDURE**) statement invokes returns the value of the system clock when the SPL function starts.

On UNIX and Linux[®] systems, the precision of the value returned by the **CURRENT** operator is determined by its **DATETIME** Qualifier, which can range from a single time unit (such as **MONTH TO MONTH**) up to **YEAR TO FRACTION (5)**. The system clock on Windows, however, returns only millisecond precision. Even if you specify "FRACTION(5)" in the **DATETIME** Qualifier, the **CURRENT** operator on Windows supports no greater than "FRACTION(3)" precision.

If your platform does not provide a system call that returns the current time with subsecond precision, **CURRENT** returns a zero for the **FRACTION** field.

In the following example, the first statement uses **CURRENT** in a **WHERE** condition. The second statement uses **CURRENT** as an argument to the **DAY** function. The last query selects rows whose **call_dtime** value is within a range from the beginning of 2007 to the current instant:

```
DELETE FROM cust_calls WHERE res_dtime < CURRENT YEAR TO MINUTE;

SELECT * FROM orders WHERE DAY(ord_date) < DAY(CURRENT);

SELECT * FROM cust_calls WHERE call_dtime
    BETWEEN '2007-1-1 00:00:00' AND CURRENT;
```

For more information, see " **DATETIME** Field Qualifier" on page 4-38.

SYSDATE Operator

The **SYSDATE** operator returns the current **DATETIME** value from the system clock. **SYSDATE** is identical to the **CURRENT** operator, except that the default precision of **SYSDATE** is **YEAR TO FRACTION(5)**, while the default precision of **CURRENT** is **YEAR TO FRACTION(3)**.

On Windows platforms that do not support a *seconds* scale greater than **FRACTION(3)**, **SYSDATE** is in effect a synonym for the **CURRENT** operator.

You can use **SYSDATE** in any context where the **CURRENT** operator is valid.

The SQL statements in the following example use the **SYSDATE** operator to specify the default values for two **DATETIME** columns of a database table, and to insert a new row into the table:

```
CREATE TABLE tab1 (
    id SERIAL,
    value CHAR(20),
    time1 DATETIME YEAR TO FRACTION(5) DEFAULT SYSDATE,
    time2 DATETIME YEAR TO SECOND DEFAULT SYSDATE YEAR TO SECOND
);

INSERT INTO tab1 VALUES (0, 'description', SYSDATE, SYSDATE);
```

The following query accesses the table that was created in the previous example:
SELECT SYSDATE AS sysdate, * FROM tab1;

The results are sensitive to the date and time when the INSERT and SELECT statements are issued, but the query could return these values on September 23, 2007:

```
sysdate  2007-09-23 21:30:23.00000
id        1
value    description
time1    2007-09-23 21:29:27.00000
time2    2007-09-23 21:29:27
```

The next query accesses the same table, using **SYSDATE** in the WHERE clause as an argument to the **DAY** function:

```
SELECT *, DAY(time1) AS day FROM tab1
      WHERE DAY(time1) = DAY(SYSDATE);
```

The query could return these values on September 23, 2007:

```
id        1
value    description
time1    2007-09-23 21:29:27.00000
time2    2007-09-23 21:29:27
day      23
```

Only Dynamic Server supports **SYSDATE**. Except for its name and its default precision, the description of the **CURRENT** operator in this document also describes the **SYSDATE** operator.

Non-default Time Zones for **CURRENT** or **TODAY** (XPS)

For Extended Parallel Server, the time and date returned by **CURRENT** (like the date returned by **TODAY**) defaults to the value from the system clock-calendar, but you can specify instead that **CURRENT** return the time and date from some other time zone. You can specify a time zone by using the SET ENVIRONMENT CLIENT_TZ statement of SQL (or by setting the **IBM_XPS_PARAMS** environment variable) to specify an offset from Greenwich Mean Time (GMT).

Suppose that a table in an Extended Parallel Server instance has this schema:

```
create table testtab (intcol integer,
                    dtcol      date,
                    datecol    datetime year to second);
```

Assume also that the **IBM_XPS_PARAMS** setting for CLIENT_TZ has the value +1, that the server time zone is at GMT+4 hours, that the current server date is 2005-05-15, and the current server time is 1:10:39. If the system clocks for client and server are both based on the same value for GMT, this implies a client time zone of GMT+1 hour, or 3 hours earlier than the time zone of the server.

Now you enter the following SQL statements:

```
INSERT INTO testtab VALUES (1, TODAY, CURRENT);
SELECT * FROM testtab WHERE intcol = 1;
```

The query returns 1 (from column **intcol**) and the following **TODAY** and **CURRENT** values:

```
2005-05-15          01:10:39.000 05/15/2005
```

Now assume that the client issues the following statements:

```

SET ENVIRONMENT CLIENT_TZ "-06:30";
INSERT INTO testtab VALUES (3, TODAY, CURRENT);
SELECT * FROM testtab WHERE intcol = 2;

```

The query returns 2 (from column **intcol**) and the following **TODAY** and **CURRENT** values:

```

2005-05-14    14:40:39 05/14/2005

```

This value is 10.5 hours earlier than the result from the first example.

Note that values that have been stored into a database table are not affected by the **CLIENT_TZ** specification, as the following example illustrates. The **CLIENT_TZ** setting is used only when evaluating the **CURRENT** and **TODAY** operators as specific **DATETIME** and **DATE** values.

```

SET ENVIRONMENT CLIENT_TZ ON;
SELECT * FROM testtab WHERE intcol = 1;

```

The query returns 1 (from column **intcol**) and the following **TODAY** and **CURRENT** values:

```

2005-05-15    01:10:39 05/15/2005

```

It is also important to note that the **INSERT** statements above do not evaluate to the time on the client system. Instead, they are based on the time of the database server, as offset to the time zone of the client system. These two times might differ, if the server and client system clocks are not synchronized correctly.

Literal DATETIME

The following examples show literal **DATETIME** as an expression:

```

SELECT DATETIME (2007-12-6) YEAR TO DAY FROM customer;

```

```

UPDATE cust_calls SET res_dtime = DATETIME (2008-07-07 10:40)
    YEAR TO MINUTE
WHERE customer_num = 110
AND call_dtime = DATETIME (2008-07-07 10:24) YEAR TO MINUTE;

```

```

SELECT * FROM cust_calls
WHERE call_dtime
    = DATETIME (2008-12-25 00:00:00) YEAR TO SECOND;

```

For more information, see “Literal **DATETIME** ” on page 4-181.

Literal INTERVAL

The following examples each use a literal **INTERVAL** as an expression:

```

INSERT INTO manufact VALUES ('CAT', 'Catwalk Sports',
    INTERVAL (16) DAY TO DAY);

```

```

SELECT lead_time + INTERVAL (5) DAY TO DAY FROM manufact;

```

The second example adds five days to each value of **lead_time** selected from the **manufact** table.

For more information, see “Literal **INTERVAL**” on page 4-183.

UNITS Operator

The UNITS operator specifies an INTERVAL value whose precision includes only one time unit. You can use UNITS in arithmetic expressions that increase or decrease one of the time units in an INTERVAL or DATETIME value.

If the *num* operand is not an integer, it is truncated to the largest whole number that is the same as (or nearer to zero than) the specified value when the database server evaluates the expression.

In the following example, the first SELECT statement uses the UNITS operator to select all the **manufacturer.lead_time** values, increased by five days. The second SELECT statement finds all the calls that were placed more than 30 days ago.

If the expression in the WHERE clause returns a value greater than 99 (maximum number of days), the query fails. The last statement increases the lead time for the ANZA manufacturer by two days:

```
SELECT lead_time + 5 UNITS DAY FROM manufact;
```

```
SELECT * FROM cust_calls WHERE (TODAY - call_dtime) > 30 UNITS DAY;
```

```
UPDATE manufact SET lead_time = 2 UNITS DAY + lead_time  
WHERE manu_code = 'ANZ';
```

NEXTVAL and CURRVAL Operators (IDS)

You can access the value of a sequence using the **NEXTVAL** or **CURRVAL** operators in SQL statements. You must qualify **NEXTVAL** or **CURRVAL** with the name (or synonym) of a sequence object that exists in the same database, using the format *sequence.NEXTVAL* or *sequence.CURRVAL*. An expression can also qualify *sequence* by the *owner* name, as in **zelaine.myseq.CURRVAL**. You can specify the SQL identifier of *sequence* or a valid synonym, if one exists.

In an ANSI-compliant database, you must qualify the name of the *sequence* with the name of its owner (*owner.sequence*) if you are not the owner.

To use **NEXTVAL** or **CURRVAL** with a sequence, you must have the Select privilege on the sequence or have the DBA privilege on the database. For information about sequence-level privileges, see the “GRANT ” on page 2-395 statement.

Using NEXTVAL:

To access a sequence for the first time, you must refer to *sequence.NEXTVAL* before you can refer to *sequence.CURRVAL*. The first reference to **NEXTVAL** returns the initial value of the sequence. Each subsequent reference to **NEXTVAL** increments the value of the sequence by the defined *step* and returns a new incremented value of the sequence.

You can increment a given sequence only once within a single SQL statement. Even if you specify *sequence.NEXTVAL* more than once within a single statement, the sequence is incremented only once, so that every occurrence of *sequence.NEXTVAL* in the same SQL statement returns the same value.

Except for the case of multiple occurrences within the same statement, every *sequence.NEXTVAL* expression increments the *sequence*, regardless of whether you subsequently commit or roll back the current transaction.

If you specify *sequence.NEXTVAL* in a transaction that is ultimately rolled back, some sequence numbers might be skipped.

Using CURRVAL:

Any reference to **CURRVAL** returns the current value of the specified sequence, which is the value that your last reference to **NEXTVAL** returned. After you generate a new value with **NEXTVAL**, you can continue to access that value using **CURRVAL**, regardless of whether another user increments the sequence.

If both *sequence.CURRVAL* and *sequence.NEXTVAL* occur in an SQL statement, the sequence is incremented only once. In this case, each *sequence.CURRVAL* and *sequence.NEXTVAL* expression returns the same value, regardless of the order of *sequence.CURRVAL* and *sequence.NEXTVAL* within the statement.

Concurrent Access to a Sequence:

A sequence always generates unique values within a database without perceptible waiting or locking, even when multiple users refer to the same sequence concurrently. When multiple users use **NEXTVAL** to increment the sequence, each user generates a unique value that other users cannot see.

When multiple users concurrently increment the same sequence, gaps occur between the values that each user sees. For example, one user might generate a series of values, such as 1, 4, 6, and 8, from a sequence, while another user concurrently generates the values 2, 3, 5, and 7 from the same sequence object.

Restrictions:

NEXTVAL and **CURRVAL** are valid only in SQL statements, not directly in SPL statements. (But SQL statements that use **NEXTVAL** and **CURRVAL** can be used in SPL routines.) The following restrictions apply to these operators in SQL statements:

- You must have Select privilege on the *sequence*.
- In a CREATE TABLE or ALTER TABLE statement, you cannot specify **NEXTVAL** or **CURRVAL** in the following contexts:
 - In the Default clause
 - In the definition of a check constraint.
- In a SELECT statement, you cannot specify **NEXTVAL** or **CURRVAL** in the following contexts:
 - In the projection list when the DISTINCT keyword is used
 - In the WHERE, GROUP BY, or ORDER BY clauses
 - In a subquery
 - When the UNION operator combines SELECT statements.
- You also cannot specify **NEXTVAL** or **CURRVAL** in these contexts:
 - In fragmentation expressions
 - In reference to a remote sequence object in another database.

Examples:

In the following examples, it is assumed that no other user is concurrently accessing the sequence and that the user executes the statements consecutively.

These examples are based on the following sequence object and table:

```
CREATE SEQUENCE seq_2
  INCREMENT BY 1 START WITH 1
  MAXVALUE 30 MINVALUE 0
  NOCYCLE CACHE 10 ORDER;
```

```
CREATE TABLE tab1 (col1 int, col2 int);
INSERT INTO tab1 VALUES (0, 0);
```

You can use **NEXTVAL** (or **CURRVAL**) in the Values clause of an INSERT statement, as the following example shows:

```
INSERT INTO tab1 (col1, col2)
  VALUES (seq_2.NEXTVAL, seq_2.NEXTVAL);
```

In the previous example, the database server inserts an incremented value (or the first value of the sequence, which is 1) into the **col1** and **col2** columns of the table.

You can use **NEXTVAL** (or **CURRVAL**) in the SET clause of the UPDATE statement, as the following example shows:

```
UPDATE tab1
  SET col2 = seq_2.NEXTVAL
  WHERE col1 = 1;
```

In the previous example, the incremented value of the **seq_2** sequence, which is 2, replaces the value in **col2** where **col1** is equal to 1.

The following example shows how you can use **NEXTVAL** and **CURRVAL** in the Projection clause of the SELECT statement:

```
SELECT seq_2.CURRVAL, seq_2.NEXTVAL FROM tab1;
```

In the previous example, the database server returns two rows of incremented values, 3 and 4, from both the **CURRVAL** and **NEXTVAL** expressions. For the first row of **tab1**, the database server returns the incremented value 3 for **CURRVAL** and **NEXTVAL**; for the second row of **tab1**, it returns the incremented value 4.

For more examples on how to use **NEXTVAL** and **CURRVAL**, see the *IBM Informix Guide to SQL: Tutorial*.

Literal Row (IDS)

The syntax for a literal representation of the value of a named or unnamed ROW data type is described in the section “Literal Row” on page 4-186. The following examples show literal rows as expressions:

```
INSERT INTO employee VALUES
  (ROW('103 Baker St', 'San Francisco',
       'CA', 94500));
```

```
UPDATE rectangles
  SET rect = ROW(8, 3, 7, 20)
  WHERE area = 140;
```

```
EXEC SQL update table(:a_row)
  set x=0, y=0, length=10, width=20;
```

```
SELECT row_col FROM tab_b
  WHERE ROW(17, 'abc') IN (row_col);
```

For the syntax of expressions that evaluate to field values of a ROW data type, see “ROW Constructors” on page 4-82.

The format of the value for each field must be compatible with the data type of the ROW field to which it is assigned.

You can use any kind of expression as a value with a ROW constructor, including literals, functions, and variables. The following examples show the use of different types of expressions with ROW constructors to specify values:

```
ROW(5, 6.77, 'HMO')
```

```
ROW(col1.name, 45000)
```

```
ROW('john davis', TODAY)
```

```
ROW(USER, SITENAME)
```

The following statement uses literal numbers and quoted strings with ROW constructors to insert values into **col1** and **col2** of the **new_tab** table:

```
INSERT INTO new_tab
VALUES
(
ROW(32, 65)::row_t,
ROW('CA', 34)
);
```

When you use a ROW constructor to generate values for a named ROW type, you must explicitly cast the ROW value to the appropriate named ROW type. The cast is necessary to generate a value of the named ROW type. To cast the ROW value as a named ROW type, you can use the cast operator (**::**) or the **CAST AS** keywords, as the following examples show:

```
ROW(4,5)::row_t
CAST (ROW(3,4) AS row_t)
```

You can use a ROW constructor to generate ROW type values in **INSERT**, **UPDATE**, and **SELECT** statements. In the next example, the **WHERE** clause of a **SELECT** statement specifies a ROW type value that is cast as type **person_t**:

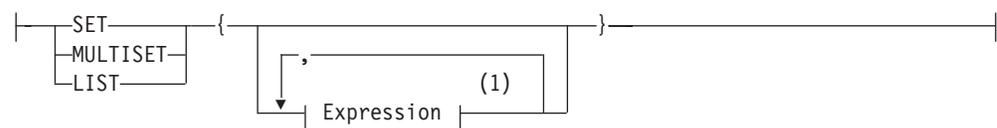
```
SELECT * FROM person_tab
WHERE col1 = ROW('charlie','hunter')::person_t;
```

For more information on using ROW constructors in **INSERT** and **UPDATE** statements, see the **INSERT** and **UPDATE** statements in this document. For information on named ROW types, see the **CREATE ROW TYPE** statement. For information on unnamed ROW types, see the discussion of the ROW data type in the *IBM Informix Guide to SQL: Reference*. For task-oriented information on named ROW types and unnamed ROW types, see the *IBM Informix Database Design and Implementation Guide*.

Collection Constructors

Use a collection constructor to specify values for a collection column.

Collection Constructors:



Notes:

- 1 See "Expression" on page 4-40

You can use collection constructors in the WHERE clause of the SELECT statement and the VALUES clause of the INSERT statement. You can also pass collection constructors to UDRs.

This table differentiates the types of collections that you can construct.

Keyword	Description
SET	Indicates a collection of elements with the following qualities: <ul style="list-style-type: none"> • The collection must contain unique values. • Elements have no specific order associated with them.
MULTISET	Indicates a collection of elements with the following qualities: <ul style="list-style-type: none"> • The collection can contain duplicate values. • Elements have no specific order associated with them.
LIST	Indicates a collection of elements with the following qualities: <ul style="list-style-type: none"> • The collection can contain duplicate values. • Elements have ordered positions.

The element type of the collection can be any built-in or extended data type. You can use any kind of expression with a collection constructor, including literals, functions, and variables.

When you use a collection constructor with a list of expressions, the database server evaluates each expression to its equivalent literal form and uses the literal values to construct the collection.

You specify an empty collection with a set of empty braces ({ }).

Elements of a collection cannot be NULL. If a collection element evaluates to a NULL value, the database server returns an error.

The element type of each expression must all be exactly the same data type. To accomplish this, cast the entire collection constructor expression to a collection type, or cast individual element expressions to the same type. If the database server cannot determine that the collection type and the element types are homogeneous, then the collection constructor returns an error. In the case of host variables, this determination is made at bind time when the client declares the element type of the host variable.

An exception to this restriction can occur when some elements of a collection are VARCHAR data types but others are longer than 255 bytes. Here the collection constructor can assign a CHAR(*n*) type to all elements, for *n* the length in bytes of the longest element. (But see "Collection Data Types" on page 4-35 for an example based on this exception, where the user avoids fixed-length CHAR elements by an explicit cast to the LVARCHAR data type.)

Examples of Collection Constructors:

The following example shows that you can construct a collection with various expressions, if the resulting values are of the same data type:

```

CREATE FUNCTION f (a int) RETURNS int;
    RETURN a+1;
END FUNCTION;
CREATE TABLE tab1 (x SET(INT NOT NULL));
INSERT INTO tab1 VALUES
(
SET{10,
    1+2+3,
    f(10)-f(2),
    Sqrt(100) +POW(2,3),
    (SELECT tabid FROM systables WHERE tabname = 'sysusers'),
    'T'::BOOLEAN::INT}
);
SELECT * FROM tab1 WHERE
    x=SET{10,
        1+2+3,
        f(10)-f(2),
        Sqrt(100) +POW(2,3),
        (SELECT tabid FROM systables WHERE tabname = 'sysusers'),
        'T'::BOOLEAN::INT}
};

```

This assumes that a cast from BOOLEAN to INT exists. (For a more restrictive syntax to specify collection values , see “Literal Collection” on page 4-178.)

NULL Keyword

The NULL keyword is valid in most contexts where you can specify a value. What it specifies, however, is the absence of any value (or an unknown or missing value).

NULL Keyword:

Within SQL, the keyword NULL is the only syntactic mechanism for accessing a NULL value. NULL is not equivalent to zero, nor to any specific value. In ascending ORDER BY operations, NULL values precede any non-NULL value; in descending sorts, NULL values follow any non-NULL value. In GROUP BY operations, all NULL values are grouped together. (Such groups might in fact be logically heterogeneous, if they include missing or unknown values.)

The keyword NULL is a global symbol in the syntactic context of expressions, meaning that its scope of reference is global.

Every data type, whether built-in or user-defined, can represent a NULL value. IBM Informix Dynamic Server supports cast expressions in the projection list. This means that users can write expressions of the form `NULL::datatype`, in which *datatype* is any data type known to the database server.

IBM Informix Dynamic Server supports the typed NULL keyword in general expressions. NULL alone in these scenarios results in a -201 syntax error. As a result, if null is defined as a column name or a procedure name, it must be referenced with a table alias. Otherwise, it returns a -201 syntax error. The behavior is summarized in the following examples and results:

```

create table tab1 (a int, null int);
create table tab2 (a int, b int);

```

Table 4-1. NULL behavior

Statement	Result
select null from tab1 where a = 1	-201 syntax error
select * from tab1 where null = a	-201 syntax error
select * from tab1 where tab1.null = a	Valid syntax
select * from tab1 where a = null	-201 syntax error
select * from tab2 where a = null	-201 syntax error
select * from tab2 where null = a	-201 syntax error
select * from tab2 where null = a	-201 syntax error
select NULL::int from tab1	Valid syntax
select NULL::int from tab1	Valid syntax
select 1 + NULL::int from tab1	Valid syntax
select 1 + NULL::int from tab2	Valid syntax
select NULL::int + 1 from tab1	Valid syntax

IBM Informix Dynamic Server prohibits the redefinition of NULL, because allowing such definition would restrict the global scope of the NULL keyword. For this reason, any mechanism that restricts the global scope or redefines the scope of the keyword NULL will syntactically disable any cast expression involving a NULL value. You must ensure that the occurrence of the keyword NULL receives its global scope in all expression contexts.

For example, consider the following SQL code:

```
CREATE TABLE newtable
(
  null int
);

SELECT null, null::int FROM newtable;
```

The CREATE TABLE statement is valid, because the column identifiers have a scope of reference that is restricted to the table definition; they can be accessed only within the scope of a table.

The SELECT statement in the example, however, poses some syntactic ambiguities. Does the identifier **null** appearing in the projection list refer to the global keyword NULL, or does it refer to the column identifier **null** that was declared in the CREATE TABLE statement?

- If the identifier **null** is interpreted as the column name, the global scope of cast expressions with the NULL keyword will be restricted.
- If the identifier **null** is interpreted as the NULL keyword, the SELECT statement must generate a syntactic error for the first occurrence of **null** because the NULL keyword can appear only as a cast expression in the projection list.

A SELECT statement of the following form is valid because the NULL column of **newtable** is qualified with the table name:

```
SELECT newtable.null, null::int FROM newtable;
```

More involved syntactic ambiguities arise in the context of an SPL routine that has a variable named **null**. An example follows:

```
CREATE FUNCTION nulltest() RETURNING INT;  
  DEFINE a INT;  
  DEFINE null INT;  
  DEFINE b INT;  
  LET a = 5;  
  LET null = 7;  
  LET b = null;  
  RETURN b;  
END FUNCTION;  
  
EXECUTE FUNCTION nulltest();
```

When the preceding function executes in DB-Access, in the expressions of the LET statement, the identifier **null** is treated as the keyword NULL. The function returns a NULL value instead of 7.

Using **null** as a variable of an SPL routine would restrict the use of a NULL value in the body of the SPL routine. Therefore, the preceding SPL code is not valid, and causes IBM Informix Dynamic Server to return the following error:

```
-947 Declaration of an SPL variable named 'null' conflicts  
with SQL NULL value.
```

In ESQL/C, you should use an indicator variable if there is the possibility that a SELECT statement will return a NULL value.

Function Expressions

A function expression can return one or more values from built-in SQL functions or from user-defined functions, as the following diagram shows.

Function Expressions:

(1)	Algebraic Functions	(2)
	CARDINALITY Function	(3)
	SQLCODE Function	(4)
	DBINFO Function	(5)
(6)	Encryption and Decryption Functions	(7)
	Exponential and Logarithmic Functions	(8)
	HEX Function	(9)
	Length Functions	(10)
(6)	Security Label Support Functions	(11)
	Smart-Large-Object Functions	(11)
	Time Functions	(12)
	Trigonometric Functions	(13)
	String-Manipulation Functions	(14)
	IFX_ALLOW_NEWLINE Function	(15)
	User-Defined Functions	(16)

Notes:

- 1 Informix extension
- 2 See "Algebraic Functions" on page 4-89
- 3 See "CARDINALITY Function (IDS)" on page 4-98
- 4 See "SQLCODE Function (SPL)" on page 4-98
- 5 See "DBINFO Function" on page 4-99
- 6 Dynamic Server only
- 7 See "Encryption and Decryption Functions" on page 4-107
- 8 See "Exponential and Logarithmic Functions" on page 4-116
- 9 See "Aggregate Expressions" on page 4-165
- 10 See "Length Functions" on page 4-118
- 11 See "Smart-Large-Object Functions (IDS)" on page 4-122
- 12 See "Time Functions" on page 4-128
- 13 See "Trigonometric Functions" on page 4-139
- 14 See "String-Manipulation Functions" on page 4-141
- 15 See "IFX_ALLOW_NEWLINE Function" on page 4-160
- 16 See "User-Defined Functions" on page 4-161

The following examples show function expressions:

```

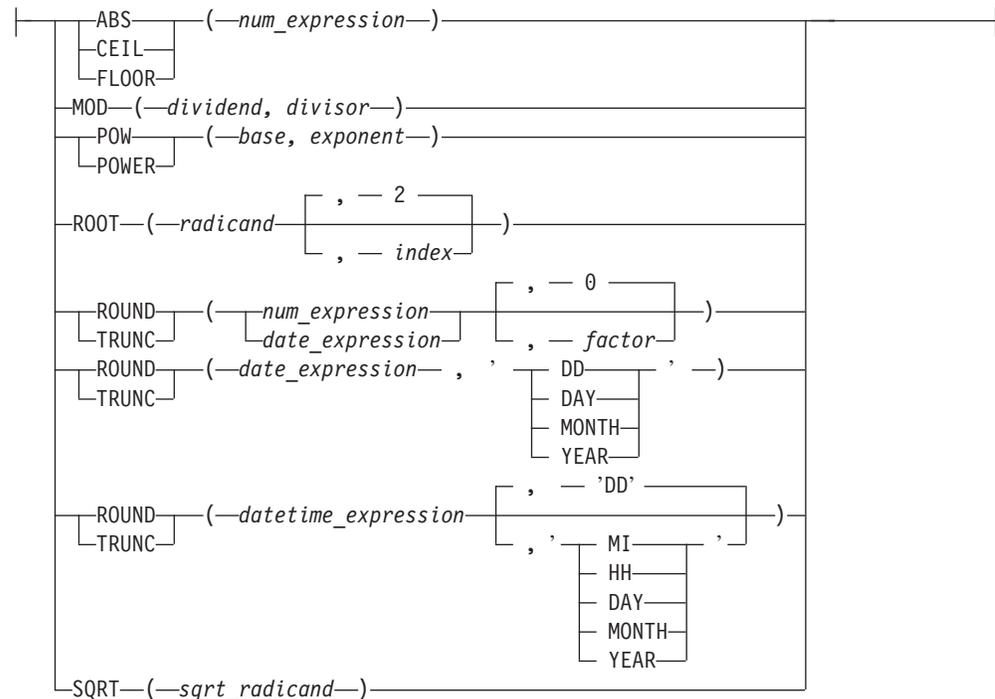
EXTEND (call_dtime, YEAR TO SECOND)
HEX (LENGTH(123))
MDY (12, 7, 1900 + cur_yr)
TAN (radians)
DATE (365/2)
ABS (-32)
LENGTH ('abc') + LENGTH (pvar)
EXP (3)
HEX (customer_num)
MOD (10,3)

```

Algebraic Functions

Algebraic functions take one or more arguments of numeric data types. Besides supporting numeric arguments, the **CEIL** and **FLOOR** functions can also take character string arguments that can be converted to DECIMAL values, and the **ROUND** and **TRUNC** functions can also take DATE or DATETIME arguments.

Algebraic Functions:



Element	Description	Restrictions	Syntax
<i>base</i>	Value to be raised to the power specified in <i>exponent</i>	Must return a real number	"Expression" on page 4-40
<i>date_expression</i>	Expression that evaluates to (or is cast to) a DATE value	Must return a DATE value	"Expression" on page 4-40
<i>datetime_expression</i>	Expression that evaluates to (or is cast to) a DATETIME value	Must return a DATETIME value	"Expression" on page 4-40
<i>dividend</i>	Value to be divided by <i>divisor</i>	A real number	"Expression" on page 4-40

Element	Description	Restrictions	Syntax
<i>divisor</i>	Value by which to divide <i>dividend</i>	A nonzero real number	"Expression" on page 4-40
<i>exponent</i>	Power to which to raise <i>base</i>	A real number	"Expression" on page 4-40
<i>factor</i>	Number of significant digits to replace with zero in the returned value. Default is to return the rounded or truncated integer part of the first argument.	Integer in range +32 to -32. Positive or unsigned values are applied to the right of the decimal point, and negative values are applied to the left.	"Literal Number " on page 4-185
<i>index</i>	Root to extract. The default is 2.	A nonzero real number	"Expression" on page 4-40
<i>num_expression</i>	Expression that evaluates to (or is cast to) a numeric value	A real number	"Expression" on page 4-40
<i>radicand</i>	Value whose root is to be returned	A real number	"Expression" on page 4-40
<i>sqrt_radicand</i>	Number with a real square root	A nonnegative real number	"Expression" on page 4-40

ABS Function:

The **ABS** function returns the absolute value of its numeric argument, returning the same data type as its argument. The query in the following example returns all orders for which a **ship_charge** greater than \$20 was paid in cash (+) or as store credit (-) .

```
SELECT order_num, customer_num, ship_charge
FROM orders WHERE ABS(ship_charge) > 20;
```

CEIL Function:

The **CEIL** function takes as its argument a numeric expression, or a string that can be converted to a DECIMAL data type, and returns the DECIMAL(32) representation of the smallest integer that is greater than or equal to its single argument. The following query returns 33 as the smallest integer that is larger than or equal to the **CEIL** argument of 32.3:

```
SELECT CEIL(32.3) FROM systables WHERE tabid = 1;
```

The next example returns -32 as the smallest integer that is larger than or equal to the **CEIL** argument of -32.3 :

```
SELECT CEIL(-32.3) FROM systables WHERE tabid = 1;
```

FLOOR Function:

The **FLOOR** function takes as its argument a numeric expression, or a string that can be converted to a DECIMAL data type, and returns the DECIMAL(32) representation of the largest integer that is smaller than or equal to its single argument. The following query returns 32 as the largest integer that is smaller than or equal to the **FLOOR** argument of 32.3:

```
SELECT FLOOR(32.3) FROM systables WHERE tabid = 1;
```

The next example returns -33 as the largest integer that is smaller than or equal to the **FLOOR** argument of -32.3 :

```
SELECT FLOOR(-32.3) FROM systables WHERE tabid = 1;
```

These examples illustrate how the **FLOOR** and **CEIL** functions provide upper and lower bounds that differ by 1 when they have the same argument that has a nonzero fractional part. For an integer argument, **FLOOR** and **CEIL** return the same **DECIMAL(32)** representation of their argument.

MOD Function:

The **MOD** function takes as arguments two real number operands, and returns the remainder from integer division of the integer part of the first argument (the *dividend*) by the integer part of the second argument (the *divisor*). The value returned is an **INT** data type (or **INT8** for remainders outside the range of **INT**). The quotient and any fractional part of the remainder are discarded. The *divisor* cannot be 0. Thus, **MOD** (*x*,*y*) returns *y* (modulo *x*). Make sure that any variable that receives the result is of a data type that can store the returned value.

This example tests to see if the current date is within a 30-day billing cycle:
`SELECT MOD(TODAY - MDY(1,1,YEAR(TODAY)),30) FROM orders;`

POW Function:

The **POW** function raises its first numeric argument, the *base*, to the power of its second numeric argument, the *exponent*. The returned value is a **FLOAT** data type. The following example returns all rows from the **circles** table in which the **radius** column value implies an area less than 1,000 square units, using an approximation to pi with a scale of 4:

```
SELECT * FROM circles WHERE (3.1416 * POW(radius,2)) < 1000;
```

The function identifier **POWER** is a synonym for **POW**.

To use *e*, the base of natural logarithms, see “**EXP Function**” on page 4-117.

ROOT Function:

The **ROOT** function extracts a positive real root value, returned as a **FLOAT** data type, from its first numeric expression argument, the *radicand*. If you specify a second numeric argument, the *index*, which cannot be zero, then the returned value to the power *index* is equal (within rounding error) to the *radicand* argument. If only the *radicand* argument is supplied, 2 is the default *index* value. You cannot specify zero as the value of *index*.

The first **SELECT** statement in the following example, which uses the default *index* value of 2, returns the positive square root of the literal number 9. The second example returns the cube root of the literal number 64.

```
SELECT ROOT(9) FROM angles;           -- square root of 9
SELECT ROOT(64,3) FROM angles;       -- cube root of 64
```

SQRT Function:

The **SQRT** function returns the positive square root of its argument, which must be a non-negative numeric expression. The following example returns the square root of 9 for each row of the **angles** table:

```
SELECT SQRT(9) FROM angles;
```

The **SQRT** function is equivalent to **ROOT(x)**.

ROUND Function:

The **ROUND** function can reduce the precision of its first numeric, DATE, or DATETIME argument, and returns the rounded value. If the first argument is neither a number nor a point in time, it must be cast to a numeric, DATE, or DATETIME data type.

The **ROUND** function resembles the **TRUNC** function, but differs in how it treats any portion of its first argument that is smaller than the least significant digit or time unit within the precision that its second argument specifies.

- If the absolute value of this portion is greater than half of the smallest unit within the precision, the value of that digit or time unit is incremented by 1 in the value returned by **ROUND**. If this portion is less than half of a unit, however, it is discarded, and only the digits or time units of the first argument within the specified or default precision are returned.
- The **TRUNC** function, in contrast, replaces with zero any digits less than the specified precision for numeric expressions, and for DATE or DATETIME expressions replaces any time units smaller than the specified format string with 1 for *month* or *day* time units, or with zero for time units smaller than *day*.

The **ROUND** function can accept an optional second argument that specifies the precision of the returned value. The syntax and semantics of the second argument depend on whether the first argument is a number expression, a DATETIME expression, or DATE expression.

- When the first argument is a numeric expression, the returned value is a DECIMAL and the second argument can be an integer in the range from -32 to +32 inclusive, specifying the position (relative to the decimal point) of the last significant digit of the returned value. If you omit the *factor* specification when the first argument is numeric, **ROUND** returns the value of the first argument rounded to a scale of zero, or to the units place.

Positive-digit values specify rounding to the right of the decimal point; negative-digit values specify rounding to the left of the decimal point, as Figure 4-1 shows.



Figure 4-1. Examples of negative, zero, and positive rounding factors

The following example uses the **ROUND** function with a column expression as its first argument and no second argument, so that the numeric expression is rounded to a scale of zero. This query returns the order number and rounded total price of items whose total price (rounded to the default scale of zero decimal places) is equal to \$124.00.

```
SELECT order_num , ROUND(total_price) FROM items
WHERE ROUND(total_price) = 124.00;
```

If you use a MONEY data type as the argument for the **ROUND** function and you round to an explicit or default scale of zero, the returned value is represented with .00 as the fractional part. The SELECT statement in the following example rounds 125.46 and a MONEY column value. The query returns 125 and a rounded price in the form xxx.00 for each row in the **items** table.

```
SELECT ROUND(125.46), ROUND(total_price) FROM items;
```

- When the first argument to **ROUND** is a DATETIME expression, the returned value is a DATETIME YEAR TO MINUTE data type and the second argument must be a quoted string that specifies the smallest significant time unit in the returned value. If you omit the second argument, the default format string is 'DD', specifying the nearest day, with the hour and minute rounded to 00:00. The following format strings are valid as the second argument:

Table 4-2. Format strings for DATETIME arguments to the ROUND function

Format String	Effect on Returned DATETIME Value
'YEAR'	Rounded to the beginning of the nearest year, with dates after June 30 rounded up to the next year. The <i>month</i> , <i>day</i> , <i>hour</i> , and <i>minute</i> values round to -01-01 00:00.
'MONTH'	Rounded to the beginning of the nearest month. Dates after the 15th are rounded up to the next month. The <i>day</i> , <i>hour</i> , and <i>minute</i> values round to 01 00:00.
'DD'	Rounded to the beginning (00:00 = midnight) of the nearest day. DATETIME values later than 12:00 noon are rounded up to the next day.
'DAY'	Rounded to the beginning of the nearest Sunday. Dates that fall on Wednesday, Thursday, Friday, or Saturday are rounded up to the next Sunday.
'HH'	Rounded to the beginning of the nearest hour. Time of day values with <i>minute:second</i> later than 29:59 are rounded up to the next hour. Minutes round to zero.
'MI'	Rounded to the beginning of the nearest minute. Time of day values with <i>second</i> later than 30 are rounded up to the next minute.

If you omit the format string specification after an initial DATETIME expression argument, the returned value is the value of the first argument rounded to the nearest day, as if you had specified 'DD' as the format string.

Examples that follow use the **ROUND** function with a column expression that returns a DATETIME YEAR TO FRACTION(5) value in a SELECT statement. In these queries, table **mytab** has only a single row, and in that row the value of **mytab.col_dt** is 2006-12-07 14:30:12.12300.

The following query specifies 'YEAR' as the DATETIME format string:

```
SELECT ROUND(col_dt, 'YEAR') FROM mytab;
```

The value returned is 2007-01-01 00:00.

The next query resembles the previous query, but casts the returned value to a DATE data type:

```
SELECT ROUND(col_dt, 'YEAR')::DATE FROM mytab;
```

The value returned is 01/01/2007.

This example specifies 'MONTH' as the DATETIME format string:

```
SELECT ROUND(col_dt, 'MONTH') FROM mytab;
```

The value returned is 2006-12-01 00:00.

This example rounds the DATETIME expression to YEAR TO HOUR precision:

```
SELECT ROUND(col_dt, 'HH') FROM mytab;
```

The value returned is 2006-12-07 15:00.

- When the first argument is a DATE expression, the returned value is also a DATE data type if the second argument is a quoted string that specifies the smallest time unit in the returned value. These are the same format strings as for rounding DATETIME values, except that 'HH' and 'MI' are not valid for dates. There is no default format string for rounding DATE arguments.

To return formatted DATE values, you must specify one of the following quoted strings as the second argument to the **ROUND** function:

Table 4-3. Format strings for DATE arguments to the ROUND function

Format String	Effect on Returned DATE Value
'YEAR'	Rounded to the beginning of the nearest year. Dates after June 30 are rounded up to the next year. The <i>month</i> and <i>day</i> values each round to 01.
'MONTH'	Rounded to the beginning of the nearest month. Dates after the 15th are rounded up to the next month. The returned <i>day</i> value is 01.
'DD'	The DATE value of the first <i>date_expression</i> argument is returned.
'DAY'	The value is rounded to the nearest Sunday. If the first argument is a Sunday, that date is returned. Dates that fall on Wednesday, Thursday, Friday, or Saturday are rounded up to the next Sunday.

If you specify no format string as the second argument when the first argument is a DATE data type, no format string takes effect as the default. No error is issued, but the first argument is treated as numeric expression that evaluates to an integer, rather than as a DATE value. Dynamic Server stores DATE values internally as the integer count of days since 31 December 1899. For dates in the 21st century, integer equivalents to DATE values are 5-digit integers, ranging between approximately 37,000 and 74,000.

For example, the query `SELECT ROUND(TODAY) FROM systables` provides no format string for a DATE expression, and returns the integer 39538 if the query is issued on 1 April 2008.

If you apply a numeric format specification as the second argument, nonnegative numbers have no effect on DATE values, but the following example rounds the last two digits of the returned value to zero:

```
SELECT ROUND(TODAY, -2) FROM systables;
```

For applications where integer-format dates like 39500 are unhelpful, you can use the 'YEAR', 'MONTH', 'DAY', or 'DD' format strings as the second argument to the **ROUND** function to prevent the DATE argument from being processed as if it were a number expression. On 1 April 2008, the following query returns the DATE value 04/01/2008 if MDY4/ is the **DBDATE** environment variable setting:

```
SELECT ROUND(TODAY, 'DD') FROM systables WHERE tabid = 1;
```

In the following example, a query is issued on Tuesday, April 3, 2007:

```
SELECT ROUND(TODAY, 'DAY') FROM mytab;
```

The returned value is 04/01/2007, the current date rounded to the nearest Sunday.

If you are using a host variable to store a rounded point-in-time value in dynamic SQL, and the data type of the first argument is not known at prepare time, Dynamic Server assumes that a DATETIME data type is the first argument to the **ROUND** function and returns a DATETIME YEAR TO MINUTE rounded value. At

execution time, after the statement is prepared, error -9750 is issued if a DATE value is supplied for the host variable. To prevent this error, you can specify the data type for the host variable by using a cast, as in this program fragment.

```

sprintf(query1, ",
      "select round( ?::date, 'DAY') from mytab");
EXEC SQL prepare selectq from :query;
EXEC SQL declare select_cursor cursor for selectq;
EXEC SQL open select_cursor using :hostvar_date_input;

EXEC SQL fetch select_cursor into :var_date_output;

```

TRUNC Function:

The **TRUNC** function can reduce the precision of its first numeric, DATE, or DATETIME argument by returning the truncated value. If the first argument is neither a number nor a point in time, it must be cast to a numeric, DATE, or DATETIME data type.

The **TRUNC** function resembles the **ROUND** function, but truncates (rather than rounds to the nearest whole number) any portion of its first argument that is smaller than the least significant digit or time unit within the precision that its second argument specifies.

- For numeric expressions, **TRUNC** replaces with zero any digits less than the specified precision.
- For DATE or DATETIME expressions, **TRUNC** replaces any time units smaller than the format specification with 1 for *month* or *day* time units, or with 0 for time units smaller than *day*.

The **TRUNC** function can accept an optional second argument that specifies the precision of the returned value.

- When the first argument is a numeric expression, the second argument must be an integer in the range from -32 to +32 inclusive, specifying the position (relative to the decimal point) of the last significant digit of the returned value. If you omit the *factor* specification when the first argument is numeric, **TRUNC** returns the value of the first argument truncated to a scale of zero, or to the units place. Positive digit values specify truncation to the right of the decimal point; negative digit values specify truncation to the left, as Figure 4-2 shows.

Expression:		2 4 5 3 6 . 8 7 4 6
TRUNC (24536.8746, -2) =24500		↓ ↓ ↓
TRUNC (24536.8746, 0) = 24536		↓ ↓ ↓
TRUNC (24536.8746, 2) = 24536.87		-2 0 2

Figure 4-2. Examples of negative, zero, and positive truncation factors

The following example calls the **TRUNC** function with a column expression that returns a numeric value in a SELECT statement. This statement displays the order number and truncated total price of items whose total price (truncated to the default scale of zero decimal places) is equal to \$124.00.

```

SELECT order_num , TRUNC(total_price) FROM items
WHERE TRUNC(total_price) = 124.00;

```

If a MONEY data type is the argument in a call to the **TRUNC** function that specifies a scale of zero, the fractional part becomes .00 in the returned value.

For example, the following SELECT statement truncates 125.46 and a MONEY column value. It returns 125 and a truncated price in the form xxx.00 for each row in the **items** table.

```
SELECT TRUNC(125.46), TRUNC(total_price) FROM items;
```

- When the first argument to **TRUNC** is a DATETIME expression, the second argument must be a quoted string that specifies the smallest significant time unit in the returned value. Only the following format strings are valid as the second argument:

Table 4-4. Format strings for DATETIME arguments to the TRUNC function

Format String	Effect on Returned Value
'YEAR'	Truncated to the beginning of the year. The <i>month</i> , <i>day</i> , <i>hour</i> , and <i>minute</i> values truncate to 01-01 00:00.
'MONTH'	Truncated to the beginning of the first day of the month. The <i>hour</i> and <i>minute</i> values round to 00:00.
'DD'	Truncated to the beginning (00:00 = midnight) of the same day.
'DAY'	If the first argument is a Sunday, midnight (00:00) on that date is returned. For any other day of the week, midnight on the previous Sunday is returned.
'HH'	Truncated to the beginning of the hour. The <i>minute</i> value truncates to zero.
'MI'	Truncated to the beginning of the nearest minute. As for all of these format strings, time units smaller than <i>minute</i> are discarded.

If you omit the format string specification after an initial DATETIME expression argument, the returned value is the value of the first argument truncated to the day, as if you had specified 'DD' as the format string.

Examples that follow invoke the **TRUNC** function with a column expression that returns a DATETIME YEAR TO FRACTION(5) value in a SELECT statement. In these examples, table **mytab** has only a single row, and in that row the value of **mytab.col_dt** is 2006-12-07 14:30:12.12300.

This query specifies 'YEAR' as the DATETIME format string:

```
SELECT TRUNC(col_dt, 'YEAR') FROM mytab;
```

The value returned is 2006-01-01 00:00.

The next query resembles the previous query, but casts the truncated value to a DATE data type:

```
SELECT TRUNC(col_dt, 'YEAR')::DATE FROM mytab;
```

The value returned is 01/01/2006.

This example specifies 'MONTH' as the DATETIME format string:

```
SELECT TRUNC(col_dt, 'MONTH') FROM mytab;
```

The value returned is 2006-12-01 00:00.

The following example truncates the DATETIME expression to YEAR TO HOUR precision:

```
SELECT TRUNC(col_dt, 'HH') FROM mytab;
```

The value returned is 2006-12-07 14:00.

- When the first argument is a DATE expression, the second argument should generally be a quoted string that specifies the smallest time unit in the returned value. These are the same format strings as for truncating DATETIME values,

except that 'HH' and 'MI' are not valid for dates, and there is no default format string for truncating DATE expression arguments.

To return formatted DATE values, you must use one of the following quoted strings as the second argument to the **TRUNC** function:

Table 4-5. Format strings for DATE arguments to the TRUNC function

Format String	Effect on Returned Value
'YEAR'	Truncated to the beginning of the year. The <i>month</i> and <i>day</i> values are each 01.
'MONTH'	Truncated to the beginning of the month. The <i>day</i> value is 01.
'DD'	The DATE value of the first <i>date_expression</i> argument is returned.
'DAY'	If the first argument is a Sunday, that date is returned. For any other day of the week, the date of the previous Sunday is returned.

If you specify no format string as the second argument when the first argument is a DATE data type, no format string takes effect as the default. No error is issued, but the first argument is treated as numeric expression that evaluates to an integer, rather than as a DATE value. Dynamic Server stores DATE values internally as the integer count of days since 31 December 1899.

For example, the query `SELECT ROUND(TODAY) FROM systables` provides no format string for a DATE expression, and returns the integer 39538 if the query is issued on 1 April 2008.

If you apply a numeric format specification as the second argument, nonnegative numbers have no effect on DATE values, but the following example rounds the last two digits of the returned value to zero:

```
SELECT TRUNC(TODAY, -2) FROM systables;
```

For applications where integer dates like 39500 are unhelpful, use the 'YEAR', 'MONTH', 'DAY', or 'DD' format strings as the second argument to the **TRUNC** function, to prevent the DATE expression from being processed as if it were a number expression. On 1 April 2008, the following query returns the DATE value 04/01/2008 if MDY4/ is the setting of the **DBDATE** environment variable:

```
SELECT TRUNC(TODAY, 'DD') FROM systables;
```

If you are using a host variable to store a truncated point-in-time value in dynamic SQL, and the data type of the first argument is not known at prepare time, Dynamic Server assumes that a DATETIME data type is the first argument to the **TRUNC** function and returns a DATETIME YEAR TO MINUTE truncated value. At execution time, after the statement is prepared, error -9750 is issued if a DATE value is supplied for the host variable. To prevent this error, you can specify the data type for the host variable by using a cast, as in this program fragment.

```
sprintf(query2, "%s",
"select trunc( ?::date, 'DAY') from mytab");
EXEC SQL prepare selectq from :query2;
EXEC SQL declare select_cursor cursor for selectq;
EXEC SQL open select_cursor using :hostvar_date_input;

EXEC SQL fetch select_cursor into :var_date_output;
```

Note that the **TRUNC** function name is based on a use of the English word "truncate" that is different from its meaning in the TRUNCATE statement of SQL.

The **TRUNC** function replaces the value of its first argument with another value that has a smaller precision or the same precision. The **TRUNCATE** statement deletes all of the rows from a database table, without dropping the table schema.

CARDINALITY Function (IDS)

The **CARDINALITY** function returns the number of elements in a collection column (SET, MULTISET, LIST).

CARDINALITY Function:

|—CARDINALITY—(— $\left. \begin{array}{l} \text{collection_col} \\ \text{collection_var} \end{array} \right\}$ —)——|

Element	Description	Restrictions	Syntax
<i>collection_col</i>	An existing collection column	Must be declared as a collection data type	"Identifier " on page 5-21
<i>collection_var</i>	Host or program collection variable	Must be declared as a collection data type	Language specific

Suppose that the **set_col** SET column contains the following value:

```
{3, 7, 9, 16, 0}
```

The following SELECT statement returns 5 as the number of elements in the **set_col** column:

```
SELECT CARDINALITY(set_col)
FROM table1;
```

If the collection contains duplicate elements, **CARDINALITY** counts each individual element.

SQLCODE Function (SPL)

The **SQLCODE** function takes no arguments, but returns to its calling context the value of **sqlca.sqlcode** for the most recently executed SQL statement (whether static or dynamic) that the current SPL routine has executed:

SQLCODE:

|—SQLCODE——|

You can use **SQLCODE** in expressions within SPL routines to identify the state of a dynamic cursor. This built-in function is useful in error handling and in contexts such as determining whether a query or function call has returned no rows, or when a cursor has reached the last row of the active set, or other conditions when SPL program control should exit from a loop.

The following SPL program fragment illustrates the use of **SQLCODE** to detect the end of the active set of a cursor within a WHILE loop.

```
CREATE PROCEDURE ...
...
DEFINE myc1 ...
...
PREPARE p FOR "SELECT c1 FROM t1";
```


Element	Description	Restrictions	Syntax
<i>table</i>	Table for which to display the dbspace name or coserver ID corresponding to each row (XPS only) or containing an integer <i>column</i> of UTC values.	Must match the name of a table in the FROM clause of the query	"Identifier " on page 5-21
<i>tblspace_num</i>	Tblspace number (partition number) of a table	Must exist in the partnum column of the sysables table for the database	"Literal Number " on page 4-185
<i>utc_value</i>	A UTC value to be converted to the DATETIME equivalent	Must be a non-negative integer number of seconds since 1970-01-01 00:00:00+00:00	"Literal Number " on page 4-185

DBINFO Options:

The **DBINFO** function is actually a set of functions that return different types of information about the database. To invoke each function, specify a particular option after the **DBINFO** keyword. You can use any **DBINFO** option anywhere within SQL statements and within UDRs.

The following table shows the categories of database and database server information that Dynamic Server and Extended Parallel Server can retrieve with valid **DBINFO** options.

- The **Arguments** column shows the parentheses-delimited argument list of each valid **DBINFO** option.
- The **Information Returned** column shows the type of database information that the **Arguments** option retrieves.
- The **Page** column shows where you can find more information about the **Arguments** option.

Arguments	Information Returned	Page
('dbhostname')	The host name of the database server to which a client application is connected	"Using the 'dbhostname' Option" on page 4-104
('dbspace' <i>tblspace_num</i>)	The name of a dbspace corresponding to a tblspace number	"Using the ('dbspace', <i>tblspace_num</i>) Option" on page 4-101
('get_tz')	The time zone of the session, \$TZ , as specified as a string by the client.	" Using the 'get_tz' Option" on page 4-105
('serial8')	The last SERIAL8 value inserted in a table	" Using the 'serial8' Option and 'bigserial' Option" on page 4-105
('bigserial')	The last BIGSERIAL value inserted in a table	" Using the 'serial8' Option and 'bigserial' Option" on page 4-105
('sessionid')	The session ID number of the current session	"Using the 'sessionid' Option" on page 4-103
('sqlca.sqlerrd1')	The last SERIAL value inserted in a table	"Using the 'sqlca.sqlerrd1' Option" on page 4-102

Arguments	Information Returned	Page
('sqlca.sqlerrd2')	The number of rows processed by SELECT, INSERT, DELETE, UPDATE, EXECUTE PROCEDURE, and EXECUTE FUNCTION statements	"Using the 'sqlca.sqlerrd2' Option" on page 4-102
('utc_current')	The current UTC time value (as an integer number of seconds since 1970-01-01 00:00:00+00:00) when the SQL statement began to execute.	"Using the 'utc_current' Option" on page 4-106
('utc_to_datetime', <i>table.column</i>)	The DATETIME value corresponding to a specified integer column containing a UTC time value (as an integer number of seconds since 1970-01-01 00:00:00+00:00).	"Using the 'utc_to_datetime' Option" on page 4-106
('utc_to_datetime', <i>utc_value</i>)	The DATETIME value corresponding to a specified UTC time value (as an integer number of seconds since 1970-01-01 00:00:00+00:00).	"Using the 'utc_to_datetime' Option" on page 4-106
('version', ' <i>parameter</i> ')	Type of the database server and its release version to which the client application is connected. (The call to DBINFO fails with an error if no <i>parameter</i> specifies a format for the version information.)	"Using the 'version' Option" on page 4-104

The following additional **DBINFO** options are valid only for Extended Parallel Server:

Arguments (XPS only)	Information Returned	Page
('coserverid')	The coserver ID of the coserver to which the user who entered the query is connected	"Using the 'coserverid' Option with No Other Arguments (XPS)" on page 4-106
('coserverid' <i>table.column</i> 'currentrow')	The coserver ID of the coserver where each row of a specified table is located	"Using the 'coserverid' Option Followed by Table and Column Names (XPS)" on page 4-107
('dbspace' <i>Ttable.column</i> 'currentrow')	The name of the dbspace where each row of a specified table is located	"Using the 'dbspace' Option Followed by Table and Column Names (XPS)" on page 4-107

Using the ('dbspace', *tblspace_num*) Option: The '**dbspace**' option returns a character string that contains the name of the dbspace that corresponds to a *tblspace_num*. You must supply an additional parameter, either *tblspace_num* or an expression that evaluates to *tblspace_num*. The following example uses the '**dbspace**' option. First, it queries the **systables** system catalog table to determine the *tblspace_num* for the table customer, then it executes the function to determine the *dbspace* name.

```
SELECT tabname, partnum FROM systables
  where tabname = 'customer';
```

If the query returns a partition number of 1048892, you insert that value into the second argument to find which dbSPACE contains the **customer** table, as the following example shows:

```
SELECT DBINFO ('dbSPACE', 1048892) FROM systables
  where tabname = 'customer';
```

If the table for which you want to know the dbSPACE name is fragmented, you must query the **sysfragments** system catalog table to find out the tblSPACE number of each table fragment. Then you must supply each tblSPACE number in a separate **DBINFO** query to find out all the dbSPACES across which a table is fragmented.

Using the 'sqlca.sqlerrd1' Option: The **'sqlca.sqlerrd1'** option returns a single integer that provides the last serial value that is inserted into a table. To ensure valid results, use this option immediately following a singleton INSERT statement that inserts a single row with a serial value into a table.

Tip: To obtain the value of the last SERIAL8 value that is inserted into a table, use the **'serial8'** option of **DBINFO**. For more information, see "Using the 'serial8' Option and 'bigserial' Option" on page 4-105.

The following example uses the **'sqlca.sqlerrd1'** option:

```
EXEC SQL create table fst_tab (ordernum serial, partnum int);
EXEC SQL create table sec_tab (ordernum serial);
EXEC SQL insert into fst_tab VALUES (0,1);
EXEC SQL insert into fst_tab VALUES (0,4);
EXEC SQL insert into fst_tab VALUES (0,6);
EXEC SQL insert into sec_tab values (dbinfo('sqlca.sqlerrd1'));
```

This example inserts a row that contains a primary-key serial value into the **fst_tab** table, and then uses the **DBINFO** function to insert the same serial value into the **sec_tab** table. The value that the **DBINFO** function returns is the serial value of the last row that is inserted into **fst_tab**.

For more information about the SQL Communications Area (SQLCA) data structure, within which **sqlca.sqlerrd1** is a field, see the *IBM Informix Guide to SQL: Tutorial*.

Using the 'sqlca.sqlerrd2' Option: The **'sqlca.sqlerrd2'** option returns a single integer that provides the number of rows that SELECT, INSERT, DELETE, UPDATE, EXECUTE PROCEDURE, and EXECUTE FUNCTION statements processed. To ensure valid results, use this option after SELECT, EXECUTE PROCEDURE, and EXECUTE FUNCTION statements have completed executing. In addition, to ensure valid results when you use this option within cursors, make sure that all rows are fetched before the cursors are closed.

The following example shows an SPL routine that uses the **'sqlca.sqlerrd2'** option to determine the number of rows that are deleted from a table:

```
CREATE FUNCTION del_rows (pnumb INT)
RETURNING INT;

DEFINE nrows INT;

DELETE FROM fst_tab WHERE part_number = pnumb;
```

```

LET nrows = DBINFO('sqlca.sqlerrd2');
RETURN nrows;

END FUNCTION;

```

For more information about the SQL Communications Area (SQLCA) data structure, within which **sqlca.sqlerrd2** is a field, see the *IBM Informix Guide to SQL: Tutorial*.

Using the 'sessionid' Option: The **'sessionid'** option of the **DBINFO** function returns the session ID of your current session. When a client application makes a connection to the database server, the database server starts a session with the client and assigns a session ID for the client. The session ID serves as a unique identifier for a given connection between a client and a database server.

The database server stores the value of the session ID in a data structure in shared memory that is called the *session control block*. The session control block for a given session also includes the user ID, the process ID of the client, the name of the host computer, and a variety of status flags.

When you specify the **'sessionid'** option, the database server retrieves the session ID of your current session from the session control block and returns this value to you as an integer. Some of the System-Monitoring Interface (SMI) tables in the **sysmaster** database include a column for session IDs, so you can use the session ID that the **DBINFO** function obtained to extract information about your own session from these SMI tables. For further information on the session control block, see the *IBM Informix Administrator's Guide*. For further information on the **sysmaster** database and the SMI tables, see the *IBM Informix Administrator's Reference*.

In the following example, the user specifies the **DBINFO** function in a **SELECT** statement to obtain the value of the current session ID. The user poses this query against the **systables** system catalog table and uses a **WHERE** clause to limit the query result to a single row.

```

SELECT DBINFO('sessionid') AS my_sessionid
FROM systables
WHERE tabname = 'systables';

```

In the preceding example, the **SELECT** statement queries against the **systables** system catalog table. You can, however, obtain the session ID of the current session by querying against any system catalog table or user table in the database. For example, you can enter the following query to obtain the session ID of your current session:

```

SELECT DBINFO('sessionid') AS user_sessionid
FROM customer
WHERE customer_num = 101;

```

You can use the **DBINFO 'sessionid'** option not only in SQL statements but also in SPL routines. The following example shows an SPL function that returns the value of the current session ID to the calling program or routine:

```

CREATE FUNCTION get_sess()
RETURNING INT;
RETURN DBINFO('sessionid');
END FUNCTION;

```

Using the 'dbhostname' Option: You can use the 'dbhostname' option to retrieve the host name of the database server to which a database client is connected. This option retrieves the physical computer name of the computer on which the database server is running.

In the following example, the user enters the 'dbhostname' option of DBINFO in a SELECT statement to retrieve the host name of the database server to which DB–Access is connected:

```
SELECT DBINFO('dbhostname')
      FROM systables
      WHERE tabid = 1;
```

The following table shows the result of this query.

(constant)
rd_lab1

Using the 'version' Option: You can use the 'version' option of the DBINFO function to retrieve information from the message log about the type of database server (and about its release version) against which the client application is running. You must include a 'parameter' specification after the 'version' option to indicate which part of the version string you want to retrieve.

If after 'version' you specify 'full' as the parameter value, DBINFO returns the complete version string, which is the same value that the -V option of the oninit utility displays. The following table lists all the valid parameter arguments to DBINFO that can retrieve version information about the database server:

- The **Arguments** column shows the parentheses-delimited argument list of each valid DBINFO ('version', 'parameter') combination.
- The **Part of Version String Returned** column shows which part of the version string each Arguments list returns.
- The **Example of Returned Value** column shows gives an example of what is returned by each value of parameter for the Arguments option.

Each example returns part of the complete version string Dynamic Server Version 11.10.UC1.

Arguments	Part of Version String Returned	Example of Returned Value
('version', 'server-type')	Type of database server	Dynamic Server
('version', 'major')	Major version number of the current database server version	11
('version', 'minor')	Minor version number of the current database server version	10
('version', 'os')	Operating-system identifier within the version string:	U
	T = Windows	
	U = UNIX 32-bit running on a 32-bit operating system	
	H = UNIX 32-bit running on a 64-bit operating system	
	F = UNIX 64-bit running on a 64-bit operating system	
('version', 'level')	Interim release level of the current database server version	C1

Arguments	Part of Version String Returned	Example of Returned Value
('version', 'full')	Complete version string as it would be returned by oninit -V	Dynamic Server, Version 11.10.UC1

Important: Not all UNIX environments fit the word-length descriptions of operating- system (os) codes in the preceding table. For example, some U versions can run on 64-bit operating systems. Similarly, some F versions can run on operating systems with 32-bit kernels that support 64-bit applications.

The following example shows how to use the **'version'** option of **DBINFO** in a **SELECT** statement to retrieve the major version number of the database server that the DB-Access client is connected to:

```
SELECT DBINFO('version', 'major')
FROM systables
WHERE tabid = 1;
```

The following table shows the result of this query.

(constant)
7

Using the 'serial8' Option and 'bigserial' Option: The **'serial8'** option and **'bigserial'** options return a single integer that provides the last **SERIAL8** or **BIGSERIAL** value that was inserted into a table. To ensure valid results, use this option immediately following an **INSERT** statement that inserts a **SERIAL8** or **BIGSERIAL** value.

Tip: To obtain the value of the last **SERIAL** value that is inserted into a table, use the **'sqlca.sqlerrd1'** option of **DBINFO()**. For more information, see "Using the 'sqlca.sqlerrd1' Option" on page 4-102.

The following example uses the **'serial8'** option:

```
EXEC SQL create table fst_tab
(ordernum serial8, partnum int);
EXEC SQL create table sec_tab (ordernum serial8);

EXEC SQL insert into fst_tab VALUES (0,1);
EXEC SQL insert into fst_tab VALUES (0,4);
EXEC SQL insert into fst_tab VALUES (0,6);

EXEC SQL insert into sec_tab
select dbinfo('serial8')
from fst_tab where partnum = 6;
```

This example inserts a row that contains a primary-key **SERIAL8** value into the **fst_tab** table and then uses the **DBINFO** function to insert the same **SERIAL8** value into the **sec_tab** table. The value that the **DBINFO** function returns is the **SERIAL8** value of the last row that is inserted into **fst_tab**. The subquery in the last line contains a **WHERE** clause so that a single value is returned.

Using the 'get_tz' Option: The **'get_tz'** option returns the **\$TZ** string that shows the time zone of the current session.

The following example uses the **'get_tz'** option in a query of the **cust_calls** table of the **stores_demo** database:

```
EXEC SQL select first call_dtime, dbinfo('get_tz')
      from cust_calls where customer_num = 106;
```

This example returns a string value of the session time zone and the first **call_dtime** value in the **cust_calls** table for which the **customer_num** value is 106.

Using the 'utc_current' Option: The **'utc_current'** option returns the current value of Coordinated Universal Time (UTC) as an integer value that shows the number of seconds that have elapsed between 1970-01-01 00:00:00+00:00) and when the current SQL statement began to execute. Unlike Universal Time (UT), which calculates the duration of seconds from the earth's rotation, UTC uses seconds of a fixed length, based on high-precision atomic clocks.

Because of variation in the earth's gradually diminishing rotation rate, intercalary *leap seconds* are introduced from time to time in UTC to reduce discrepancies with UT time. By default, Dynamic Server ignores leap seconds in DATETIME and INTERVAL arithmetic. When Dynamic Server is supported by an operating system that takes leap seconds into account, however, the leap seconds are reflected in subsequent DATETIME and INTERVAL operations after the operating system adjusts the system clock for leap seconds.

Using the 'utc_to_datetime' Option: The **'utc_to_datetime'** option casts to a DATETIME value its last argument, which must be an integer expression representing a Coordinated Universal Time (UTC) value. The expression can be either a literal integer or the name of an integer column that stores UTC values, as shown in the following examples:

```
DBINFO ('utc_to_datetime', 1299912999 )
```

```
DBINFO ('utc_to_datetime', timesheet.utc_checkin )
```

Because SMALLINT columns cannot store UTC values more recent than 9 January 1970, the **utc_checkin** column in the second example must be an INT, BIGINT or INT8 data type for **DBINFO** to return a UTC value from the current millennium.

Using the 'coserverid' Option with No Other Arguments (XPS): The **'coserverid'** option with no other arguments returns a single integer that corresponds to the coserver ID of the coserver to which the user who entered the query is connected.

Suppose that you use the following statement to create the **mytab** table:

```
CREATE TABLE mytab (mycol INT)
      FRAGMENT BY EXPRESSION
      mycol < 5 in rootdbs.1
      mycol > 5 in rootdbs.2;
```

Further, suppose that the dbspace named **rootdbs.1** resides on coserver 1, and the dbspace named **rootdbs.2** resides on coserver 2. Also suppose that you use the following statements to insert rows into the **mytab** table:

```
INSERT INTO mytab VALUES ('1');
INSERT INTO mytab VALUES ('6');
```

Finally, suppose that you are logged on to coserver 1 when you make the following query, which displays the values of all columns in the row where the value of the **mycol** column is 1. This query also displays the coserver ID of the coserver to which you are logged on when you enter the query:

```
SELECT *, DBINFO ('coserverid') AS cid
      FROM mytab
      WHERE mycol = 1;
```

The following table shows the result of this query.

mycol	cid
1	1

Using the 'coserverid' Option Followed by Table and Column Names (XPS):

Use the 'coserverid' option followed by the table name and column name and the 'currentrow' string to find out the coserver ID where each row in a specified table is located. This option is especially useful when you fragment a table across multiple coservers.

In the following example, the user asks to see all columns and rows of the **mytab** table as well as the coserver ID of the coserver where each row resides. For a description of the **mytab** table, see "Using the 'coserverid' Option with No Other Arguments (XPS)" on page 4-106:

```
SELECT *, DBINFO ('coserverid', mytab.mycol, 'currentrow')
AS dbsp
FROM mytab;
```

The following table shows the result of this query.

mycol	cid
1	1
6	2

The *column* argument to the **DBINFO** function can be any column in the specified *table*.

Using the 'dbspace' Option Followed by Table and Column Names (XPS): Use the 'dbspace' option followed by the table name and column name and the 'currentrow' string to find out the name of the dbspace where each row in a specified table is located. This option is especially useful when you fragment a table across multiple dbspaces.

In the following example, the user asks to see all columns and rows of the **mytab** table as well as the name of the dbspace where each row resides. For a description of the **mytab** table, see "Using the 'coserverid' Option with No Other Arguments (XPS)" on page 4-106.

```
SELECT *, DBINFO ('dbspace', mytab.mycol, 'currentrow')
AS dbsp
FROM mytab;
```

The following table shows the result of this query.

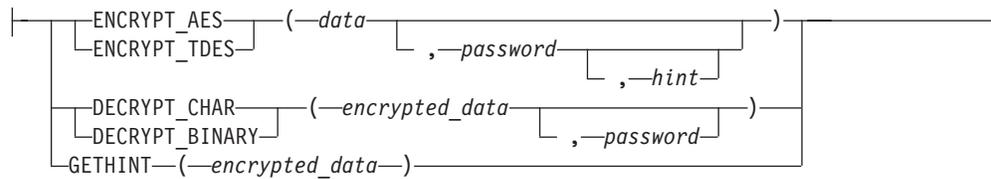
mycol	cid
1	rootdbs.1
6	rootdbs.2

The *column* arguments to **DBINFO** can be any columns in the specified *table*.

Encryption and Decryption Functions

Dynamic Server supports built-in encryption and decryption functions. The encryption functions **ENCRYPT_AES** and **ENCRYPT_TDES** each returns an *encrypted_data* value that encrypts the *data* argument. Conversely, decryption functions **DECRYPT_CHAR** and **DECRYPT_BINARY** return a plain-text *data* value from the *encrypted_data* argument. Use this syntax to call these functions:

Encryption and Decryption Functions:



Element	Description	Restrictions	Syntax
<i>data</i>	A plain text character string, variable, or large object of type BLOB or CLOB to be encrypted	Must be a character or BLOB data type	“Expression” on page 4-40
<i>encrypted_data</i>	A character string or variable containing output from ENCRYPT_AES or from ENCRYPT_TDES	Decryption requires the encryption password	“Expression” on page 4-40
<i>hint</i>	A character string that you define here. Default is the value from the WITH HINT clause of the SET ENCRYPTION statement that defined password.	No more than 32 bytes	“Quoted String” on page 4-189
<i>password</i>	A character string that the encryption function defines. Default is the session password value defined by the SET ENCRYPTION statement	At least 6 bytes, but no more than 128 bytes	“Quoted String” on page 4-189

You can invoke these encryption and decryption functions from within DML statements or with the EXECUTE FUNCTION statement.

For distributed operations over a network, all participating database servers must support these (or equivalent) functions. If the network is not secure, the DBSA must enable the *encryption communication support module* (ENCCSM) to provide data encryption between the database server and client systems, in order to avoid transmitting passwords as plain text.

Encryption or decryption calls slow the performance of the SQL statement within which these functions are invoked, but have no effect on other statements. However, if you store encrypted data in a column that is an index key, or in a column on which a constraint is defined, Dynamic Server cannot enforce the constraint, and DML statements cannot use the index.

You cannot encrypt the security label in a column of type IDSSECURITY label.

Column Level and Cell Level Encryption: The encryption and decryption functions can support two ways of using data encryption features, namely *column level* and *cell level* encryption.

- Column level encryption means that all values in a given column are encrypted with the same password (which can be a word or phrase), the same cipher, and the same cipher mode.

Users of this form of encryption should consider not using the *hint* feature of these functions, but instead store a mnemonic hint for remembering the password in some other location. Otherwise, the same *hint* will occupy disk space in every row that contains an encrypted value.

- Cell level encryption means that within a column of encrypted data many different passwords (or different ciphers or cipher modes) are used.

This use of encryption is also called *row-column level* or *set-column level* encryption. Compared to column-level encryption, this makes the task of data management more complex, because if different passwords are required for

decrypting different rows of the same table, it is not possible to write a single SELECT statement to fetch all the decrypted data. In some situations, however, individual users may need this technique to protect personal data.

To protect data security and confidentiality, the database server does not store information in the system catalog to indicate whether any table (or any column or row) includes encrypted data. Similarly, the logical logs of Dynamic Server do not record SET ENCRYPTION statements, nor calls to encryption or decryption functions. (The Trusted Facility feature for secure auditing, however, can use the 'STEP' audit-event mnemonic to record execution of the SET ENCRYPTION statement, and can use the 'CRPT' audit-event mnemonic to record successful or unsuccessful calls to DECRYPT_CHAR or DECRYPT_BINARY.)

The Password and Hint Specifications: The SET ENCRYPTION statement or an encryption function can define a password and hint for the current session. The *password* must be specified as a character expression that returns at least 6 bytes, but no more than 128. The optional *hint* is specified as a character expression that returns no more than 32 bytes.

The purpose of the *hint* is to help users to remember the *password*. When you call ENCRYPT_AES or ENCRYPT_TDES with a *hint* argument, it is encrypted and embedded in the *encrypted_data*, from which GETHINT can retrieve it. But if you define *hint* as NULL, or omit *hint* when SET ENCRYPTION specified no default *hint* for the *session password*, no *hint* is embedded in the *encrypted_data*.

The *password* used for encryption and decryption is either the *password* argument to the function, or if you omit this argument, it is the *session password* specified in the last SET ENCRYPTION statement executed before you invoke the function.

The DECRYPT_CHAR, DECRYPT_BINARY, or GETHINT function call fails with an error if the *encrypted_data* argument is not in an encrypted format, or if the *password* argument to a decryption function is omitted when no *session password* value was set by SET ENCRYPTION. An error also results if the *password* used for decryption is not the same *password* used for encryption.

Encryption key management, which is critical to the secure operation of the database, is delegated entirely to the application. This implementation means that the *password* itself is not stored in the database. Without help from the user through the application, the database server cannot decrypt the encrypted data.

If you invoke any of these functions from a UDR, you might prefer to set a *session password* in the SET ENCRYPTION statement. Otherwise, *password* will be visible to users who can view the **sysprocbody.data** column in the system catalog.

Data Types, Encoding, and Size of Encrypted Values: The *data* and corresponding *encrypted_data* arguments can be of any character type (CHAR, LVARCHAR, NCHAR, NVARCHAR, or VARCHAR), or can be a smart large object of type BLOB or CLOB. (Use CLOB in place of TEXT, which these functions do not support.)

Corresponding *data* and *encrypted_data* values that the encryption or decryption functions return have the same character, BLOB, or CLOB data type, except in cases where encryption of a VARCHAR or NVARCHAR string would return an overflow error. For operations on CHAR, LVARCHAR, NCHAR, NVARCHAR, or VARCHAR data and *encrypted_data* values, the encryption and decryption

functions follow the type promotion rules of **CONCAT** and other string manipulation functions in the data type of their return value, as in the following examples.

- If the **VARCHAR** *data* argument to **ENCRYPT_TDES** (with no hint) is a 200 byte string, then Dynamic Server automatically promotes the returned value to an **LVARCHAR** data type, because the encrypted value exceeds the 255 byte limit for **VARCHAR** objects.
- If the **NVARCHAR** *encrypted_data* argument to **ENCRYPT_AES** (with a hint) is a string 200 bytes long, then Dynamic Server automatically promotes the returned value to an **NCHAR** data type, because the encrypted value exceeds the 255 byte limit for **NVARCHAR** objects.

For more information about return type promotion for character strings that the encryption, decryption, and certain other string-manipulation functions return, see Return Types from the **CONCAT** Function. The first table in that section describes data or *encrypted_data* arguments that are not smart blobs. (For smart blob arguments, the return type is a **BLOB** or **CLOB** object.)

The encryption or decryption function call returns overflow error -881, however, if the return value exceeds the 32,767-byte limit for **CHAR**, **NCHAR**, or **LVARCHAR** strings. To avoid this error, use **BLOB** or **CLOB** objects as the data or *encrypted_data* argument, rather than a character data type, when the encryption or decryption operation requires an argument or a return value that might be larger than 32Kb.

Except for original *data* of **BLOB** or **CLOB** data types, the *encrypted_data* value is encoded in BASE64 format. An encrypted value requires more space than the corresponding plain text, because the database must also store the information (except for the encryption key) that is needed for decryption. If a *hint* is used, it adds to the length of *encrypted_data*.

The BASE64 encoding scheme stores 6 bits of input data as 8 bits of output. To encode N bytes of data, BASE64 requires at least $((4N+3)/3)$ bytes of storage, where (/) represents integer division. Padding and headers can increase BASE64 storage requirements above this $((4N+3)/3)$ ratio. “Example of Column Level Encryption” on page 4-111 lists formulae to estimate the size of data values encrypted in BASE64 format. It typically requires changes to the schema of an existing table that will store BASE64 format encrypted data, especially if a hint will also be stored.

The following table shows how the data type of the input string corresponds to the data type of the value that **ENCRYPT_AES** or **ENCRYPT_TDES** returns:

Table 4-6. Data Types for ENCRYPT_AES and ENCRYPT_TDES Functions

Plain Text Data Type	Encrypted Data Type	Decryption Function
CHAR	CHAR	DECRYPT_CHAR
NCHAR	NCHAR	DECRYPT_CHAR
VARCHAR	VARCHAR or CHAR	DECRYPT_CHAR
NVARCHAR	NVARCHAR or NCHAR	DECRYPT_CHAR
LVARCHAR	LVARCHAR	DECRYPT_CHAR
BLOB	BLOB	DECRYPT_BINARY
CLOB	BLOB	DECRYPT_CHAR

Columns of type VARCHAR and NVARCHAR store no more than 255 bytes. If the *data* string is too long for these data types to store both the encrypted data and encryption overhead, then the value returned by the encryption function is automatically changed from VARCHAR or NVARCHAR into a fixed CHAR or NCHAR value, with no trailing blanks in the encoded encrypted value.

Encrypted values of type BLOB or CLOB are not in BASE64 encoding format, and their size increase after encryption is independent of the original data size. For BLOB or CLOB values, the encrypted size (in bytes) has the following formula, where N is the original size of the plain text, and H is the size of the unencrypted hint string, if encryption is performed by **ENCRYPT_TDES**:

$N + H + 24$ bytes.

For BLOB or CLOB values that **ENCRYPT_AES** encrypts, the overhead is larger:

$N + H + 32$ bytes.

Example of Column Level Encryption:

The following example illustrates how to use the built-in encryption and decryption functions of Dynamic Server to create and use a table that stores encrypted credit card numbers in a column that has a character data type.

For purposes of this example, assume that the plain text of the values to be encrypted consists of strings of 16 digits. Because encryption functions support character data types, these values are stored in a CHAR column rather than in an INT, BIGINT, or INT8 column.

Calculating storage requirements for encrypted data:

The **LENGTH** function provides a convenient way to calculate the storage requirements of encrypted data directly:

```
EXECUTE FUNCTION LENGTH(ENCRYPT_TDES("1234567890123456", "simple password"));
```

This returns 55.

```
EXECUTE FUNCTION LENGTH(ENCRYPT_TDES("1234567890123456", "simple password",  
"12345678901234567890123456789012"));
```

This returns 107.

```
EXECUTE FUNCTION LENGTH(ENCRYPT_AES("1234567890123456", "simple password"));
```

This returns 67.

```
EXECUTE FUNCTION LENGTH(ENCRYPT_AES("1234567890123456", "simple password",  
"12345678901234567890123456789012"));
```

This returns 119.

The required storage size for encrypted data is sensitive to three factors:

- N, the number of bytes in the plain text
- whether or not a hint is provided
- which encryption function you use (**ENCRYPT_TDES** or **ENCRYPT_AES**)

The following formulae describe the four possible cases, and are not simplified:

- Encryption by **ENCRYPT_TDES()** with no hint:
Encrypted size = $(4 \times ((8 \times (N + 8)/8) + 10)/3) + 11$
- Encryption by **ENCRYPT_AES()** with no hint:

$$\text{Encrypted size} = (4 \times ((16 \times ((N + 16)/16) + 10)/3) + 11)$$

- Encryption by `ENCRYPT_TDES()` with a hint:

$$\text{Encrypted size} = (4 \times ((8 \times ((N + 8)/8) + 50)/3) + 11)$$

- Encryption by `ENCRYPT_AES()` with no hint:

$$\text{Encrypted size} = (4 \times ((16 \times ((N + 16)/16) + 50)/3) + 11)$$

The integer division (/) returns an integer quotient and discards any remainder.

Based on these formulae, the following table shows the encrypted size (in bytes) for selected ranges of values of N:

N	ENCRYPT_TDES No Hint	ENCRYPT_AES No Hint	ENCRYPT_TDES With Hint	ENCRYPT_AES With Hint
1 to 7	35	43	87	99
8 to 15	43	43	99	99
16 to 23	55	67	107	119
24 to 31	67	67	119	119
32 to 39	75	87	131	139
40 to 47	87	87	139	139
100	163	171	215	227
200	299	299	355	355
500	695	707	747	759

If the column size is smaller than the data size returned by encryption functions, the encrypted value is truncated when it is inserted. In this case, it will not be possible to decrypt the data, because the header will indicate that the length should be longer than the data value that the column contains.

These formulae and the values returned by the `LENGTH` function, however, indicate that the table schema in the next example can store the encrypted form of 16-digit credit card numbers (with a hint).

Implementing column-level encryption:

The following steps create a table from which a user who knows the password can retrieve rows that include one column of encrypted data.

1. Create a database table containing at least one column of type BLOB, CLOB, or a character data type of sufficient length to store the encrypted values. For example, the following statement creates a table called `customer` in which the column `creditcard` can store encrypted credit card numbers:

```
CREATE TABLE customer (id CHAR(20), creditcard CHAR(107));
```

2. Specify a password (and optional hint) and insert encrypted data:

```
SET ENCRYPTION PASSWORD 'credit card number is encrypted'  
WITH HINT 'Why is this difficult to read?';  
INSERT INTO customer VALUES ('Alice',  
    encrypt_tdes('1234567890123456'));  
INSERT INTO customer VALUES ('Bob',  
    encrypt_tdes('2345678901234567'));
```

3. Query the encrypted data, using a decryption function:

```
SELECT id, DECRYPT_CHAR(creditcard,  
    'credit card number is encrypted') FROM customer;
```

The following query call a decryption function in the WHERE clause, using the *session password* default, rather than an explicit *password* argument:

```
SELECT id FROM customer
WHERE DECRYPT_CHAR(creditcard) = '2345678901234567';
```

Column level encryption offers the coding convenience of passing the implicit *session password* for all rows with encrypted columns, and in multiple encryption and decryption function calls in the same SQL statement. Confidentiality of the data, however, requires users who know the password on encrypted columns to avoid compromising its secrecy. Triggers and UDRs, for example, should always use the *session password*, rather than explicit *password* arguments if they invoke the encryption or decryption functions.

The DBSA can manage highly confidential data with column level encryption. Dynamic Server does not, however, prevent users with sufficient privileges from entering data encrypted by some other password into a table whose other rows use the designated column level encryption password.

DECRYPT_CHAR Function

The **DECRYPT_CHAR** function accepts as its first argument an *encrypted_data* character string that can have any character type (CHAR, LVARCHAR, NCHAR, NVARCHAR, or VARCHAR). You must specify a *password* as its second argument, unless the SET ENCRYPTION statement has specified for this session the same *session password* by which the first argument was encrypted.

The **DECRYPT_CHAR** function also accepts as its first argument an *encrypted_data* large object of type BLOB or CLOB. You must specify a password as its second argument, unless the SET ENCRYPTION statement has specified as the default for this session the same *password* by which the first argument was encrypted. If the call to **DECRYPT_CHAR** is successful, it returns a CLOB large object that contains the plain text version of the *encrypted_data* argument.

If the call to **DECRYPT_CHAR** with an encrypted string argument is successful, it returns a character string that contains the plain text version of the *encrypted_data* argument. The following example returns a character string containing a decrypted value from the **ssid** column of the **engineers** table for the row whose **empno** value is 287:

```
SELECT DECRYPT_CHAR (ssid) FROM engineers WHERE empno = 287;
```

If the first argument to **DECRYPT_CHAR** is not an encrypted value, or if the second argument (or the default *password* specified by SET ENCRYPTION) is not the *password* that was used when the first argument was encrypted, Dynamic Server issues an error, and the call to **DECRYPT_CHAR** fails. (See the description of the “GETHINT Function” on page 4-116 for one possible action to take when you cannot remember the *password* that was used for encryption.)

Do not use **DECRYPT_CHAR** (or any other decryption function) to create a functional index on an encrypted column. This would store the decrypted values as plain text data in the database, defeating the purpose of encryption.

For additional information about using data encryption in column values of Dynamic Server databases, see “Encryption and Decryption Functions” on page 4-107, and “SET ENCRYPTION PASSWORD” on page 2-611.

DECRYPT_BINARY Function

The **DECRYPT_BINARY** function accepts as its first argument an *encrypted_data* large object of type BLOB or CLOB. You must specify a *password* as its second argument, unless the SET ENCRYPTION statement has specified as the default for this session the same *password* by which the first argument was encrypted.

If the call to **DECRYPT_BINARY** is successful, it returns a BLOB or CLOB large object that contains the plain text version of the *encrypted_data* argument. The decrypted BLOB or CLOB object is temporarily stored in the default sbspace that the SBSPACENAME configuration parameter setting specifies.

If the first argument to **DECRYPT_BINARY** is an encrypted value of a character data type, Dynamic Server invokes the **DECRYPT_CHAR** function and attempts to decrypt the specified value.

If the first argument to **DECRYPT_BINARY** is not an encrypted value, or if the second argument (or the default *password* specified by SET ENCRYPTION) is not the *password* that was used when the first argument was encrypted, Dynamic Server issues an error, and the call to **DECRYPT_BINARY** fails. (See the description of the “GETHINT Function” on page 4-116 for one possible action to take when you cannot remember the *password* that was used for encryption.)

Do not use **DECRYPT_BINARY** (or any other decryption function) to create a functional index on an encrypted column. This would store the decrypted values as plain text data in the database, defeating the purpose of encryption.

For additional information about using data encryption in column values of Dynamic Server databases, see “Encryption and Decryption Functions” on page 4-107, and “SET ENCRYPTION PASSWORD” on page 2-611.

ENCRYPT_AES Function

The **ENCRYPT_AES** function returns an encrypted value that it derives by applying the AES (Advanced Encryption Standard) algorithm to its first argument, which must be an unencrypted character expression or a smart large object (that is, a BLOB or CLOB data type). A character argument can have a length of up to 32640 bytes if an explicit or default *hint* is used, or 32672 bytes if no *hint* (or a NULL *hint*) is specified. Theoretical size limits on BLOB or CLOB arguments are many orders of magnitude larger, but practical limits might be imposed by your hardware, or by time required for encryption and decryption. The encrypted BLOB or CLOB object is temporarily stored in the default sbspace that the SBSPACENAME configuration parameter specifies.

You must specify a *password* as its second argument, unless a SET ENCRYPTION statement has specified a *session password*, which the database server uses by default if you omit the second argument. If a *session password* has been set, any *password* that you specify overrides the *session password* for the returned value of this function call. The explicit or default *password* will also be required for any subsequent decryption of the returned encrypted value. A valid *password* must have at least 6 bytes but no more than 128.

You can optionally specify a *hint* as the third argument. If the SET ENCRYPTION statement specified a default *hint* for this session, and you specify no *hint*, that default *hint* is stored in an encrypted form within the returned value. Any *hint* that you specify overrides the default *hint*. A valid *hint* can be no longer than 32 bytes.

You can use consecutive quotation marks (") to specify a NULL *hint*. If you specify an explicit *hint*, you must also specify an explicit *password*.

The purpose of the *hint* is to help users to remember the *password*. For example, if the *password* is "buggy," you might define the hint as "whip." Neither string is restricted to a single word, but the size of the *hint* contributes to the size of the returned value. If you subsequently cannot remember the *hint*, use the returned value from `ENCRYPT_AES` as the argument to `GETHINT` to retrieve the *hint*.

The following example calls `ENCRYPT_AES` from the `VALUES` clause of an `INSERT` statement that stores in `tab1` a plain-text string and an *encrypted_data* value that `ENCRYPT_AES` returns from its 12-byte first argument. Here `SET ENCRYPTION` defines a *session password* and *hint* that are used as default second and third arguments to the `ENCRYPT_AES` function:

```
EXEC SQL SET ENCRYPTION PASSWORD 'CHARYBDIS' WITH HINT 'messina';
EXEC SQL INSERT INTO tab1 VALUES ('abcd', ENCRYPT_AES("111-222-3333"));
```

The call to `ENCRYPT_AES` fails with an error if the *password* argument is omitted when no *session password* has been set, or if the length of an explicit *password* argument is shorter than 6 bytes or longer than 128 bytes.

In some contexts, an error is issued if the encrypted returned value is too large to be stored by the data type that receives it.

For additional information about using data encryption in column values of Dynamic Server databases, see "Encryption and Decryption Functions" on page 4-107, and "SET ENCRYPTION PASSWORD" on page 2-611.

ENCRYPT_TDES Function

The `ENCRYPT_TDES` function returns a value that is the result of encrypting a character expression, or a BLOB or CLOB value, by applying the TDES (Triple Data Encryption Standard, which is sometimes also called DES3) algorithm to its first argument. This algorithm is slower than the AES algorithm that is used by the `ENCRYPT_AES` function, but is considered somewhat more secure. The disk space required as encryption overhead resembles that of `ENCRYPT_AES`, but is somewhat smaller because of the smaller block size of `ENCRYPT_TDES`. (See "Calculating storage requirements for encrypted data" "Calculating storage requirements for encrypted data" on page 4-111 for a discussion of how to estimate the size of encrypted character strings.) For BLOB or CLOB values, the encrypted object is temporarily stored in the default sbspace that the `SBSPECNAME` configuration parameter specifies.

Those differences in performance, tamper-resistance, and in the returned *encrypted_data* size that the previous paragraph lists are the practical differences between the `ENCRYPT_TDES` and `ENCRYPT_AES` functions, which otherwise follow the same rules, defaults, and restrictions that appear in the description of `ENCRYPT_AES` on the previous page in regard to the following features:

- The required first argument (the plain text *data* value to be encrypted)
- The explicit or default second argument (the *password* string that must also be an argument to `DECRYPT_CHAR` or `DECRYPT_BINARY` to decrypt the returned *encrypted_data* value). This must be specified unless a default *session password* has been set by the `SET ENCRYPTION` statement
- The optional third argument (the *hint* value) that might assist users who forget the *password*. If you subsequently cannot remember an explicit or default *hint*

that was defined for *password*, you can use the returned value from **ENCRYPT_TDES** as the argument to **GETHINT** to retrieve the *hint*.

The following example calls **ENCRYPT_TDES** from the SET clause of an UPDATE statement. Here the *session password* is 'PERSEPHONE' and the *hint* string is "pomegranate", with column **colU** of table **tabU** the *data* argument. Because the WHERE clause condition of "1=1" is true for all rows of **tabU**, the effect of this statement is to replace every plain text **colU** value with encrypted strings returned by the algorithm that **ENCRYPT_TDES** implements:

```
EXEC SQL SET ENCRYPTION PASSWORD 'PERSEPHONE' WITH HINT 'pomegranate';
EXEC SQL UPDATE tabU SET colU = ENCRYPT_TDES (colU) WHERE 1=1;
```

This example assumes that the character data type of **colU** is of sufficient size to store the new encrypted values without truncation. (A more cautious example might execute an appropriate ALTER TABLE statement before the UPDATE.)

For additional information about using data encryption in column values of Dynamic Server databases, see "Encryption and Decryption Functions" on page 4-107, and "SET ENCRYPTION PASSWORD" on page 2-611.

GETHINT Function

The **GETHINT** function returns a character string that a previously executed SET ENCRYPTION PASSWORD statement defined for the *password* that was used when *encrypted_data* was encrypted by the **ENCRYPT_AES** function or by the **ENCRYPT_TDES** function. This *hint* string typically provides information that helps the user to specify the *password* needed to return the plain text version of *encrypted_data* with the **DECRYPT_CHAR** or **DECRYPT_BINARY** decryption function. The *hint* string, however, should not be the same as the *password*. In the following example, a query returns the *hint* string into a host variable called **myhint**:

```
EXEC SQL SELECT GETHINT(creditcard) INTO :myhint
FROM customer WHERE id = :myid;
```

An error is returned, rather than a *hint* string, if the *encrypted_data* argument to the **GETHINT** function is not an encrypted string or an encrypted large object.

For additional information about using data encryption in column values of Dynamic Server databases, see "Encryption and Decryption Functions" on page 4-107, and "SET ENCRYPTION PASSWORD" on page 2-611.

Exponential and Logarithmic Functions

Exponential and logarithmic functions take at least one argument and return a FLOAT data type.

Exponential and Logarithmic Functions:

```
|-----|
|  EXP  (---float_expression---) |-----|
|  LOGN  (---float_expression---) |
|  LOG10  (---float_expression---) |
```

Element	Description	Restrictions	Syntax
<i>float_expression</i>	An argument to the EXP , LOGN , or LOG10 functions. For the meaning of <i>float_expression</i> in these functions, see the individual heading for each function on the pages that follow.	The domain is the set of real numbers, and the range is the set of positive real numbers	“Expression” on page 4-40

EXP Function:

The **EXP** function returns the exponent of a numeric expression. The following example returns the exponent of 3 for each row of the **angles** table:

```
SELECT EXP(3) FROM angles;
```

For this function, the base is always *e*, the base of natural logarithms, as the following example shows:

```
e=exp(1)=2.718281828459
```

When you want to use the base of natural logarithms as the base value, use the **EXP** function. If you want to specify a particular value to raise to a specific power, see the “**POW Function**” on page 4-91.

LOG10 Function:

The **LOG10** function returns the log of a value to base 10. The following example returns the log base 10 of distance for each row of the **travel** table:

```
SELECT LOG10(distance) + 1 digits FROM travel;
```

LOG10 Function:

The **LOG10** function returns the log of a value to base 10. The following example returns the log base 10 of distance for each row of the **travel** table:

```
SELECT LOG10(distance) + 1 digits FROM travel;
```

LOGN Function:

The **LOGN** function returns the natural logarithm of a numeric argument. This value is the inverse of the exponential value. The following query returns the natural log of **population** for each row of the **history** table:

```
SELECT LOGN(population) FROM history WHERE country='US'
ORDER BY date;
```

HEX Function

The **HEX** function returns the hexadecimal encoding of an integer expression.

HEX Function:

```
|—HEX—(—int_expression—)|
```

Element	Description	Restrictions	Syntax
<i>int_expression</i>	Expression for which you want the hexadecimal equivalent	Must be a literal integer or some other expression that returns an integer	“Expression” on page 4-40

The next example displays the data type and column length of the columns of the **orders** table in hexadecimal format. For MONEY and DECIMAL columns, you can then determine the precision and scale from the lowest and next-to-the-lowest bytes. For VARCHAR and NVARCHAR columns, you can determine the minimum space and maximum space from the lowest and next-to-the-lowest bytes. For more information about encoded information, see the *IBM Informix Guide to SQL: Reference*.

```
SELECT colname, coltype, HEX(collength)
   FROM syscolumns C, systables T
   WHERE C.tabid = T.tabid AND T.tabname = 'orders';
```

The following example lists the names of all the tables in the current database and their corresponding tblspace number in hexadecimal format.

```
SELECT tabname, HEX(partnum) FROM systables;
```

The two most significant bytes in the hexadecimal number constitute the dbspace number. They identify the table in **oncheck** output (in Dynamic Server) and in **onutil check** output (in Extended Parallel Server).

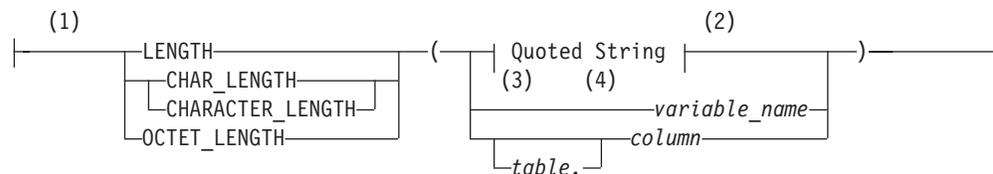
The **HEX** function can operate on an expression, as the next example shows:

```
SELECT HEX(order_num + 1) FROM orders;
```

Length Functions

Use length functions to determine the length of a column, string, or variable.

Length Functions:



Notes:

- 1 Informix extension
- 2 See "Quoted String" on page 4-189
- 3 ESQL/C
- 4 Stored Procedure Language

Element	Description	Restrictions	Syntax
<i>column</i>	Name of a column in <i>table</i>	Must have a character data type	"Identifier " on page 5-21
<i>table</i>	Name of the table in which the specified column occurs	Must exist	"Identifier " on page 5-21
<i>variable</i>	Host variable or SPL variable that contains a character string	Must have a character data type	See language-specific rules for names.

Each of these functions has a distinct purpose:

- **LENGTH**
- **OCTET_LENGTH**
- **CHAR_LENGTH** (also known as **CHARACTER_LENGTH**)

LENGTH Function:

The **LENGTH** function returns the number of bytes in a character column, not including any trailing blank spaces. For **BYTE** or **TEXT** columns, **LENGTH** returns the full number of bytes, including any trailing blank spaces.

In Informix ESQL/C, **LENGTH** can also return the length of a character variable.

The next example illustrates the use of the **LENGTH** function:

```
SELECT customer_num, LENGTH(fname) + LENGTH(lname),
       LENGTH('How many bytes is this?')
FROM customer WHERE LENGTH(company) > 10;
```

See also the discussion of **LENGTH** in the *IBM Informix GLS User's Guide*.

OCTET_LENGTH Function:

The **OCTET_LENGTH** returns the number of bytes in a character column, including any trailing spaces. See also the *IBM Informix GLS User's Guide*.

CHAR_LENGTH Function:

The **CHAR_LENGTH** function (also called **CHARACTER_LENGTH**) returns the number of logical characters (which can be distinct from the number of bytes in some East Asian locales) in a character column value. For a discussion of this function, see the *IBM Informix GLS User's Guide*.

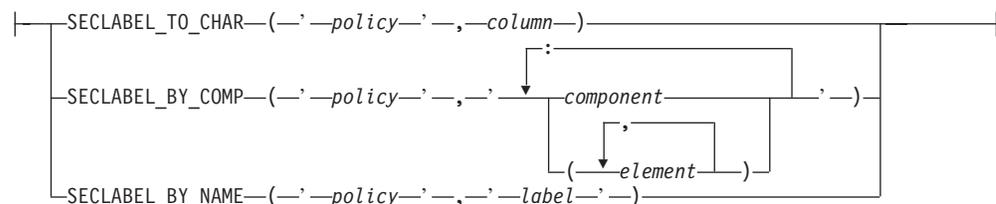
Security Label Support Functions (IDS)

The security label support functions enable users to manipulate security labels. A security label can be referenced in three different ways:

- A name, as declared in the **CREATE SECURITY LABEL** or **RENAME SECURITY LABEL** statement.
- A list of values for each component of the security policy of the security label.
- An internal encoded value that the **IDSSECURITYLABEL** data type stores.

These functions can convert between the various forms of a security label. They are typically used to specify a label in DML operations on data rows that are secured by label-based access control (LBAC). In these operations, however, the security label support functions do not provide any more access to protected data than is already provided by the security credentials of the user who invokes the function.

Security Label Support Function:



Element	Description	Restrictions	Syntax
<i>column</i>	A column of type IDSSECURITYLABEL	Must exist and must store a label of the <i>policy</i>	“Identifier ” on page 5-21

Element	Description	Restrictions	Syntax
<i>component</i>	Value of a component of the <i>policy</i>	Must exist and must be a component of the <i>policy</i>	“Quoted String” on page 4-189
<i>element</i>	Value of an element within a list of values of the <i>component</i>	Must exist and must be elements of a single component of the <i>policy</i>	“Quoted String” on page 4-189
<i>label</i>	Identifier of the security label whose value the function returns	Must exist and must be a label of the <i>policy</i>	“Quoted String” on page 4-189
<i>policy</i>	The security policy supported by the security label whose value the function returns	Must exist and must be the security policy that secures the table	“Quoted String” on page 4-189

These functions return a security label of the specified security *policy*. They can be used within DML statements that reference a protected database table, but they can also evaluate to a security labels in other calling contexts. Each of these functions requires a different argument list:

- **SECLABEL_TO_CHAR** requires the security policy name and an expression that returns a `IDSSECURITYLABEL` object, such as the name of a column of that data type.
- **SECLABEL_BY_COMP** requires the security policy name and the values of the individual components of the security label.
- **SECLABEL_BY_NAME** requires the names of the security policy and of the security label.

SECLABEL_BY_NAME Function:

The **SECLABEL_BY_NAME** function enables users to provide a security label directly by specifying its name.

The following `INSERT` statement inserts a row into table **T1**, which is protected by the security policy called '**MegaCorp**'. The `VALUES` clause of the `INSERT` statement provides the security label '**mylabel**' for the row to be inserted by using the **SECLABEL_BY_NAME** function.

```
INSERT INTO T1 VALUES (SECLABEL_BY_NAME ('MegaCorp', 'mylabel'), 1, 'xyz');
```

The success of this **SECLABEL_BY_NAME** function call does not guarantee success of the `INSERT` operation in this example, because whether or not the user has sufficient security credentials to insert the label **mylabel** into the row is subject to the **IDSLBACWRITE** rules of the **MegaCorp** security policy.

SECLABEL_BY_COMP Function:

The **SECLABEL_BY_COMP** function returns an `IDSSECURITYLABEL` object, which is a security label in its internal encoded string format. This function enables users to provide a security label directly by specifying its component values.

If a security label component requires multiple values, then such multiple values can be specified by putting those values between parenthesis as in (`value_1, value_2, ...`). When a component in a particular security label needs to be empty, it can be specified by putting nothing between an opening and a closing parenthesis, as in (). Because the blank space (ASCII 32) is a valid character in an element value for a security component, any blank space appearing in the security label string is treated as part of the element value for that component.

The security label string is limited to a maximum of 32 kilobytes. An error is returned if the string length exceeds this limit.

The following INSERT statement inserts a row into table **T1** which is protected by the security policy called **'MegaCorp'** that has three components: **'level'**, **'compartments'**, and **'groups'**. Here the user provides a security label for the row to be inserted by specifying the **SECLABEL_BY_COMP** function. The security label in this example has the value **'VP'** for the **level** component, the value **'Marketing'** for the **compartments** component, and the value **'West'** for the **groups** component. In the arguments to **SECLABEL_BY_COMP**, colon symbols separate these security component element values, and quotation marks delimit the list of component values of the security label.

```
INSERT INTO T1
VALUES (SECLABEL_BY_COMP ('MegaCorp', 'VP:Marketing:West'), 1, 'xyz');
```

In the next example, the INSERT statement inserts a row in table **T1** which is protected by the same **MegaCorp** security policy, which has the same three components as in the previous example: **level**, **compartments**, and **groups**. The user provides the security label for the row to be inserted by specifying the policy name and a list of security component elements as arguments to the **SECLABEL_BY_COMP** function. Here the security label has the value **'Director'** for the **level** component, the values **'HR'** and **'Finance'** for the **compartments** component, and the value **'East'** for the **groups** component.

```
INSERT INTO T1
VALUES (SECLABEL_BY_COMP ('MegaCorp', 'Director:(HR,Finance):East'), 1, 'xyz');
```

The following example inserts a row into table **T1** which is protected by the **MegaCorp** security policy, whose three components are **level**, **compartments**, and **groups**. The **SECLABEL_BY_COMP** function specifies the security label for the row to be inserted. The security label in this example has the value **'CEO'** for **level component**, the empty set for the **compartments** component, and the value **'EntireRegion'** for the **groups** component.

```
INSERT INTO T1
VALUES (SECLABEL_BY_COMP ('MegaCorp', 'CEO:():EntireRegion'), 3, 'abc');
```

As in all of these examples, the success of the **SECLABEL_BY_COMP** function call does not guarantee the success of the INSERT statement, because the security credentials of the user are first compared to the security label that protects table **T1**, using the **IDSLBACRWRITE** rules of the **MegaCorp** security policy, before the database server allows or denies write access for inserting the new row.

SECLABEL_TO_CHAR Function:

The **SECLABEL_TO_CHAR** function returns a security label in the security label string format.

The security credentials of the user executing this function can affect the output of the function. An element of a security label component is not included in the output if the user does not have read access to that element. A user has read access to an element if the security credentials of the user provide read access to data that is protected by a security label containing only that element and no other elements.

For the rule set **IDSLBACRULES**, only components of type **TREE** can contain elements to which a user does not have read access to a subset of elements. For other types of component, if any element blocks read access, then the user cannot

read the row at all. Thus, only security components of type TREE can have a subset of security component elements excluded in this way.

For example, if the TREE type component of the security label of a user is {A} and the TREE type component of a row security label is {A, B}, then only component A is returned, and the user is not aware that B existed in the row security label. If the user holds an exemption on the IDSLBACREADTREE rule, however, the returned security components are both A and B.

In the next example, the **MegaCorp** security policy has a security label called **mylabel** that consists of a **level** component whose value is 'Director', and a **compartments** component with the values 'HR' and 'Finance.' A user to whom 'mylabel' was granted has inserted a row with that security label into table T1. In this context, the security label string returned by the SECLABEL_TO_CHAR function in the following SELECT statement on T1 is as follows.

```
SELECT SECLABEL_TO_CHAR ('MegaCorp', C1) FROM T1;
```

Row returned:

```
'Director:(HR,Finance)'
```

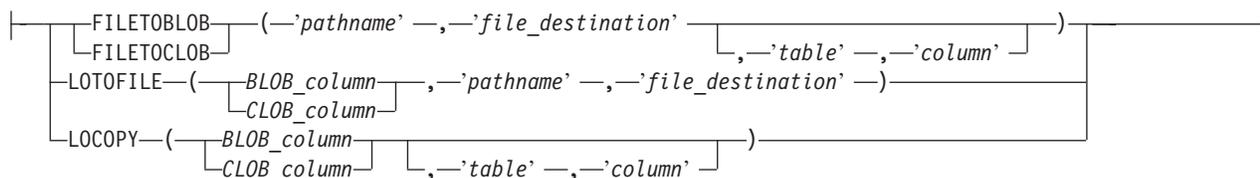
The success of this query implies that the SECLABEL_TO_CHAR function succeeded, and that the security credentials of the user were sufficient, according to the IDSLBACREAD rules of the **MegaCorp** security policy, for the database server to allow read access to the values of the security policy name and of the security label components."

The security label string is limited to a maximum size of 32 kilobytes. If the length of the security label string to be returned exceeds this upper limit, a warning is issued, and a truncated 32 kilobyte string is returned.

Smart-Large-Object Functions (IDS)

The smart-large-object functions support BLOB and CLOB data types:

Smart-Large-Object Functions:



Element	Description	Restrictions	Syntax
BLOB_column, CLOB_column	A column of type BLOB; a column of type CLOB	The column data type must be BLOB or CLOB	"Identifier " on page 5-21
column	Column within table for the copy of the BLOB or CLOB value	Must have CLOB or BLOB as its data type	"Quoted String" on page 4-189
file_destination	The system on which to put or get the smart large object	The only valid values are the strings 'server' or 'client'	"Quoted String" on page 4-189
pathname	Directory path and filename to locate the smart large object	No more than 256 bytes. Must exist on file_destination system. See also "Pathnames with Commas" on page 4-124.	"Quoted String" on page 4-189

Element	Description	Restrictions	Syntax
<i>table</i>	Table containing <i>column</i> for the copy of the BLOB or CLOB value	A comma (not a period) separates the ' <i>table</i> ' and ' <i>column</i> ' arguments	"Quoted String" on page 4-189

FILETOBLOB and FILETOCLOB Functions:

The **FILETOBLOB** function creates a BLOB value for data that is stored in a specified operating-system file. Similarly, the **FILETOCLOB** function creates a CLOB value for a data value that is stored in an operating-system file.

These functions determine the operating-system file to use from the following parameters:

- The *pathname* parameter identifies the directory path and name of the source file.
- The *file destination* parameter identifies the computer, 'client' or 'server', on which this file resides:
 - Set *file destination* to 'client' to identify the client computer as the location of the source file. The *pathname* can be either a full pathname or relative to the current directory.
 - Set *file destination* to 'server' to identify the server computer as the location of the source file. The *pathname* must be a full pathname.

The *table* and *column* parameters are optional:

- If you omit *table* and *column*, the **FILETOBLOB** function creates a BLOB value with the system-specified storage defaults, and the **FILETOCLOB** function creates a CLOB value with the system-specified storage defaults. These functions obtain the system-specific storage characteristics from either the ONCONFIG file or the sbspace. For more information on system-specified storage defaults, see the *IBM Informix Administrator's Guide*.
- If you specify *table* and *column*, the **FILETOBLOB** and **FILETOCLOB** functions use the storage characteristics from the specified column for the BLOB or CLOB value that they create.

The **FILETOBLOB** function returns a handle value (a pointer) to the new BLOB value. Similarly, **FILETOCLOB** returns a handle value to the new CLOB value. Neither function actually copies the smart-large-object value into a database column. You must assign the BLOB or CLOB value to the appropriate column.

The **FILETOCLOB** function performs any code-set conversion that might be required when it copies the file from the client or server computer to the database.

The following INSERT statement uses the **FILETOCLOB** function to create a CLOB value from the value in the **smith.rsm** file:

```
INSERT INTO candidate (cand_num, cand_lname, resume)
VALUES (2, 'Smith', FILETOCLOB('smith.rsm', 'client'));
```

In the preceding example, the **FILETOCLOB** function reads the **smith.rsm** file in the current directory on the client computer and returns a handle value to a CLOB value that contains the data in this file. Because the **FILETOCLOB** function does not specify a table and column name, this new CLOB value has the system-specified storage characteristics. The INSERT statement then assigns this CLOB value to the **resume** column in the **candidate** table.

The following INSERT statement uses the **FILETOBLOB** function to create a BLOB value from the value in the **photos.xxx** file on the local database server, and insert that value into the **election2008** table of the **rdb** database, which is another database of the local database server:

```
INSERT INTO rdb@:election2008 (cand_pic)
VALUES (FILETOBLOB('C:\tmp\photos.xxx', 'server',
'candidate', 'cand_photo'));
```

In the preceding example, the **FILETOBLOB** function reads the **photos.xxx** file in the specified directory on the local database server and returns a handle value to a BLOB value that contains the data in this file. The INSERT statement then assigns this BLOB value to the **cand_pic** column in the **election2008** table in the **rdb** database of the local database server. This new BLOB value has the storage characteristics of the **cand_photo** column in the **candidate** table in the local database.

In the following example, the new BLOB value has the storage characteristics of the **cand_pix** column in the **election96** table in the **rdb2** database, where **rdb1** and **rdb2** are databases of the local Dynamic Server instance:

```
INSERT INTO rdb1:election2008 (cand_pic)
VALUES (FILETOBLOB('C:\tmp\photos.xxx', 'server',
'rdb2:election96', 'cand_pix'));
```

When you qualify the **FILETOBLOB** or **FILETOCLOB** function with the name of a remote database and a remote database server, the *pathname* and the *file destination* become relative to the remote database server.

When you specify server as the file destination, as the following example shows, the **FILETOBLOB** function looks for the source file (in this case, **photos.xxx**) on the remote database server:

```
INSERT INTO rdb@rserv:election (cand_pic)
VALUES (rdb@rserv:FILETOBLOB('C:\tmp\photos.xxx', 'server'));
```

When you specify client as the file destination, however, as in the following example, the **FILETOBLOB** function looks for the source file (in this case, **photos.xxx**) on the local client computer:

```
INSERT INTO rdb@rserv:election (cand_pic)
VALUES (rdb@rserv:FILETOBLOB('photos.xxx', 'client'));
```

Pathnames with Commas: If a comma (,) symbol is within the *pathname* of the function, the database server expects the pathname to have the following format:
"offset, length, pathname"

For pathnames that contain a comma, you must also specify an offset and length, as in the following example:

```
FILETOBLOB("0,-1,/tmp/blob,x","server");
```

The first term in the quoted *pathname* string is an *offset* of 0, which instructs the database server to begin reading at the start of the file.

The second term is a *length* of -1, which instructs the database server to continue reading until the end of the entire file.

The third term is the */tmp/blob,x pathname*, specifying which file to read. (Notice the comma symbol that precedes the *x*.)

Because the *pathname* includes a comma, the comma-separated *offset* and *length* specifications are necessary in this example to avoid an error when **FILETOBLOB** is called. You do not need to specify *offset* and *length* for pathnames that include no comma, but including 0,-1, as the initial characters of the pathname string avoids this error for any valid pathname.

LOTOFILE Function:

The **LOTOFILE** function copies a smart large object to an operating-system file. The first parameter specifies the BLOB or CLOB column to copy. The function determines what file to create from the following parameters:

- The *pathname* identifies the directory path and the source file name.
- The *file destination* identifies the computer, 'client' or 'server', on which this file resides:
 - Set *file destination* to 'client' to identify the client computer as the location of the source file. The *pathname* can be either a full pathname or a path relative to the current directory.
 - Set *file destination* to 'server' to identify the server computer as the location of the source file. The full pathname is required.

By default, the **LOTOFILE** function generates a filename of the form:

file.hex_id

In this format, *file* is the filename you specify in *pathname* and *hex_id* is the unique hexadecimal smart-large-object identifier. The maximum number of digits for a smart-large-object identifier is 17. Most smart large objects, however, would have an identifier with fewer digits.

For example, suppose that you specify a UNIX *pathname* value as follows:

`'/tmp/resume'`

If the CLOB column has the identifier **203b2**, then **LOTOFILE** creates the file:

`/tmp/resume.203b2`

For another example, suppose that you specify a Windows *pathname* value as follows:

`'C:\tmp\resume'`

If the CLOB column has an identifier of **203b2**, the **LOTOFILE** function would create the file:

`C:\tmp\resume.203b2`

To change the default filename, you can specify the following wildcards in the filename of the *pathname*:

- One or more contiguous question mark (?) characters in the filename can generate a unique filename.

The **LOTOFILE** function replaces each question mark with a hexadecimal digit from the identifier of the BLOB or CLOB column.

For example, suppose that you specify a UNIX *pathname* value as follows:

`'/tmp/resume??.txt'`

The **LOTOFILE** function puts 2 digits of the hexadecimal identifier into the name. If the CLOB column has an identifier of **203b2**, the **LOTOFILE** function would create the file:

```
/tmp/resume20.txt
```

If you specify more than 17 question marks, **LOTOFILE** ignores them.

- An exclamation (!) point at the end of the filename indicates that the filename does not need to be unique.

For example, suppose that you specify a Windows pathname value as follows:

```
'C:\tmp\resume.txt!'
```

The **LOTOFILE** function does not use the smart-large-object identifier in the filename, so it generates the following file:

```
C:\tmp\resume.txt
```

If the filename that you specify already exists, **LOTOFILE** returns an error.

The **LOTOFILE** function performs any code-set conversion that might be required when it copies a CLOB value from the database to a file on the client or server computer.

When you qualify **LOTOFILE** with the name of a remote database and a remote database server, the BLOB or CLOB column, the *pathname*, and the *file destination* become relative to the remote database server.

When you specify *server* as the file destination, as in the next example, the **LOTOFILE** function copies the smart large object from the remote database server to a source file in the specified directory on the remote database server:

```
rdb@rserv:LOTOFILE(blob_col, 'C:\tmp\photo.gif!', 'server')
```

If you specify *client* as the file destination, as in the following example, the **LOTOFILE** function copies the smart large object from the remote database server to a source file in the specified directory on the local client computer:

```
rdb@rserv:LOTOFILE(clob_col, 'C:\tmp\essay.txt!', 'client')
```

LOCOPY Function:

The **LOCOPY** function creates a copy of a smart large object. The first parameter specifies the BLOB or CLOB column to copy. The *table* and *column* parameters are optional.

- If you omit *table* and *column*, the **LOCOPY** function creates a smart large object with system-specified storage defaults and copies the data in the BLOB or CLOB column into it.

The **LOCOPY** function obtains the system-specific storage defaults from either the ONCONFIG file or the sbspace. For more information on system-specified storage defaults, see the *IBM Informix Administrator's Guide*.

- When you specify *table* and *column*, the **LOCOPY** function uses the storage characteristics from the specified *column* for the BLOB or CLOB value that it creates.

The **LOCOPY** function returns a handle value (a pointer) to the new BLOB or CLOB value. This function does *not* actually store the new smart-large-object value into a column in the database. You must assign the BLOB or CLOB value to the appropriate column.

The following Informix ESQL/C code fragment copies the CLOB value in the **resume** column of the **candidate** table to the **resume** column of the **interview** table:

```

/* Insert a new row in the interviews table and get the
 * resulting SERIAL value (from sqlca.sqlerrd[1])
 */
EXEC SQL insert into interviews (intrv_num, intrv_time)
      values (0, '09:30');
intrv_num = sqlca.sqlerrd[1];

/* Update this interviews row with the candidate number
 * and resume from the candidate table. Use LOCOPY to
 * create a copy of the CLOB value in the resume column
 * of the candidate table.
 */
EXEC SQL update interviews
      SET (cand_num, resume) =
          (SELECT cand_num,
              LOCOPY(resume, 'candidate', 'resume')
           FROM candidate
           WHERE cand_lname = 'Haven')
      WHERE intrv_num = :intrv_num;

```

In the preceding example, the **LOCOPY** function returns a handle value for the copy of the CLOB **resume** column in the **candidate** table. Because the **LOCOPY** function specifies a table and column name, this new CLOB value has the storage characteristics of this **resume** column. If you omit the table (**candidate**) and column (**resume**) names, the **LOCOPY** function uses the system-defined storage defaults for the new CLOB value. The UPDATE statement then assigns this new CLOB value to the **resume** column in the **interviews** table.

In the following example, the **LOCOPY** function executes on the local database and returns a handle value on the local server for the copy of the BLOB **cand_pic** column in the **election2008** table in **rdb**, which is another database of the local database server. The INSERT statement then assigns this new BLOB value to the **cand_photo** column in the local **candidate** table.

```

INSERT INTO candidate (cand_photo)
      SELECT LOCOPY(cand_pic) FROM rdb:election2008;

```

When the **LOCOPY** function executes on the same database server as the original BLOB or CLOB column in a distributed query, it produces two copies of the BLOB or CLOB value, one in the remote database and the other in the local database, as the following two examples show.

In the first example, the **LOCOPY** function executes on the remote **rdb** database and returns a handle value in the remote database for the copy of the BLOB **cand_pic** column in the remote **election2008** table. The INSERT statement then assigns this new BLOB value to the **cand_photo** column in the local **candidate** table:

```

INSERT INTO candidate (cand_photo)
      SELECT rdb:LOCOPY(cand_pic)
           FROM rdb:election2008;

```

In the second example, the **LOCOPY** function executes on the local database and returns a handle value on the local database for the copy of the BLOB **cand_photo** column in the local **candidate** table. The INSERT statement then assigns this new BLOB value to the **cand_pic** column in the **election2008** table in the remote **rdb** database:

```

INSERT INTO rdb:election2008 (cand_pic)
      SELECT LOCOPY(cand_photo) FROM candidate;

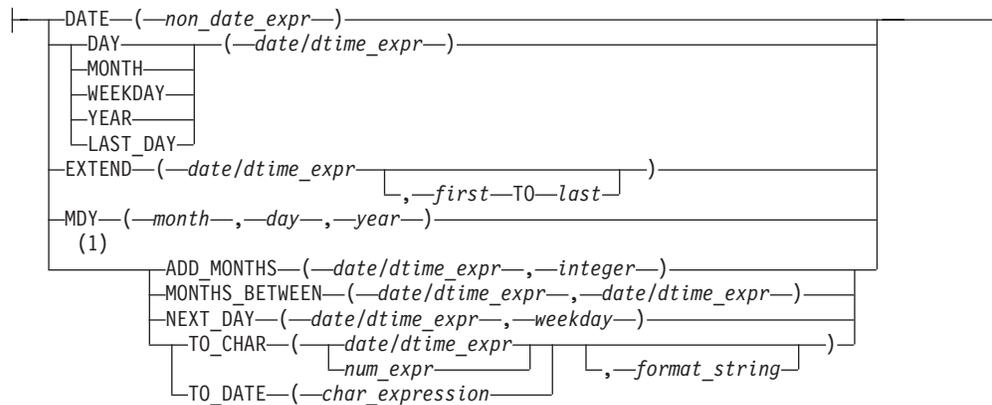
```

The BLOB and CLOB arguments of the built-in **LOCOPY** function are built-in opaque data types. These can be values returned by cross-database DML operations or by cross-database function calls, but built-in opaque types do not support distributed operations across database server instances. If the local database and the **rdb** database are databases of different Dynamic Server instances, the INSERT statements in the previous two examples fail with error -999.

Time Functions

The time functions of Dynamic Server accept DATE or DATETIME arguments, or character representation of a DATE or DATETIME value. They typically return DATE or DATETIME values, or convert information that they extract from DATE or DATETIME values into character strings. See also the descriptions of the **ROUND** and **TRUNC** functions, which can change the precision of DATE or DATETIME values, in the section “Algebraic Functions” on page 4-89.

Time Functions:



Notes:

- 1 Dynamic Server only

Element	Description	Restrictions	Syntax
<i>char_expression</i>	Expression to be converted to a DATE or DATETIME value	Must be a literal, host variable, expression, or column of a character data type	“Expression” on page 4-40
<i>date/dtime_expr</i>	Expression that returns a DATE or DATETIME value	Can be host variable, expression, column, or constant.	“Expression” on page 4-40
<i>day</i>	Expression that returns the number of a day of the month	Must return integer > 0 but no greater than the number of days in the specified month	“Expression” on page 4-40
<i>first</i>	Largest time unit in the result. If you omit <i>first</i> and <i>last</i> , the default <i>first</i> is YEAR.	Must be a DATETIME qualifier keyword that specifies a time unit no smaller than <i>last</i>	“ DATETIME Field Qualifier” on page 4-38
<i>format_string</i>	String that contains a format mask for the first argument	Must be a character data type that specifies a valid format. Can be a column, host variable, expression, or constant	“Quoted String” on page 4-189
<i>integer</i>	Expression that specifies a whole number of months	Must evaluate to positive or negative integer	“Expression” on page 4-40
<i>last</i>	Smallest time unit in the result	Must be a DATETIME qualifier keyword that specifies a time unit no smaller than <i>first</i>	“ DATETIME Field Qualifier” on page 4-38

Element	Description	Restrictions	Syntax
<i>month</i>	Expression that represents the number of the month	Must evaluate to an integer in the range from 1 to 12, inclusive	"Expression" on page 4-40
<i>non_date_expr</i>	Expression that represents a value to be converted to a DATE data type	Typically an expression that returns a CHAR, DATETIME, or INTEGER value that can be converted to a DATE data type	"Expression" on page 4-40
<i>num_expr</i>	Expression that evaluates to a real number	Must return a numeric data type	"Expression" on page 4-40
<i>weekday</i>	Abbreviated name of a day of the week	A character data type containing a valid abbreviation for a day of the week	"Quoted String" on page 4-189
<i>year</i>	Number expression that represents a year	Must evaluate to a 4-digit integer. You cannot use a 2-digit abbreviation.	"Expression" on page 4-40

ADD_MONTHS Function:

The **ADD_MONTHS** function takes a DATETIME or DATE expression as its first argument, and requires a second integer argument, specifying the number of months to add to the first argument value. The second argument can be positive or negative.

The value returned is the sum of the DATE or DATETIME value of the first argument and an INTERVAL UNITS MONTH value that is based on the number of months that the second argument specifies.

The returned data type depends on the data type of the first argument:

- If the first argument evaluates to a DATE value, **ADD_MONTHS** returns a DATE value.
- If the first argument evaluates to a DATETIME value, **ADD_MONTHS** returns a DATETIME YEAR TO FRACTION(5) value, with the same values for time units smaller than *day* as in the first argument.

If the *day* and *month* time units in the first argument specify the last day of the month, or if the resulting month has fewer days than the *day* in the first argument, then the returned value is the last day of the resulting month. Otherwise, the returned value has the same day of the month as the first argument.

The returned value can be in a different year, if the resulting month is later than December (or for negative second arguments, earlier than January) of the year in the first argument.

The following query calls the **ADD_MONTHS** function twice in the Projection clause, using column expressions as arguments. Here the column names indicate the column data types, and the **DBDATE** setting is MDY4/:

```
SELECT a_serial, b_date, ADD_MONTHS(b_date, a_serial),
       c_datetime, ADD_MONTHS(c_datetime, a_serial)
FROM mytab WHERE a_serial = 7;
```

In this example **ADD_MONTHS** returns DATE and DATETIME values:

```
a_serial      7
b_date        07/06/2007
(expression)  02/06/2008
c_datetime    2007-10-06 16:47:49.00000
(expression)  2008-05-06 16:47:49.00000
```

If you use a host variable to store the argument to **ADD_MONTHS**, but the data type of the argument is not known at prepare time, Dynamic Server assumes that the data type is DATETIME YEAR TO FRACTION(5). If at runtime, after the statement has been prepared, the user supplies a DATE value for the host variable, error -9750 is issued. To prevent this error, specify the data type of the host variable by using a cast, as in this program fragment:

```
sprintf(query, "  
    \"select add_months(?::date, 6) from mytab\"");  
EXEC SQL prepare selectq from :query;  
EXEC SQL declare select_cursor cursor for selectq;  
EXEC SQL open select_cursor using :hostvar_date_input;  
  
EXEC SQL fetch select_cursor into :var_date_output;
```

DATE Function:

The **DATE** function converts its argument to a DATE value. Its non-DATE argument can be any expression that can be converted to a DATE value, usually a CHAR, DATETIME, or INTEGER value. The following WHERE clause specifies a quoted string as its CHAR argument:

```
WHERE order_date < DATE('12/31/07')
```

When the **DATE** function interprets a CHAR non-DATE expression, it expects this expression to conform to any DATE format that the **DBDATE** environment variable specifies. For example, suppose **DBDATE** is set to Y2MD/ when you execute the following query:

```
SELECT DISTINCT DATE('02/01/2008') FROM ship_info;
```

This SELECT statement generates an error, because the **DATE** function cannot convert this string expression. The **DATE** function interprets the first part of the date string (02) as the year and the second part (01) as the month.

For the third part (2008), the **DATE** function encounters four digits when it expects a two-digit day (valid day values must be between 01 and 31). It therefore cannot convert the value. For the SELECT statement to execute successfully with the Y2MD/ value for **DBDATE**, the argument would need to be '08/02/01'. For information on the format of **DBDATE**, see the *IBM Informix Guide to SQL: Reference*.

When you specify a positive INTEGER value for the non-DATE expression, the **DATE** function interprets this as the number of days after December 31, 1899.

If the integer value is negative, the **DATE** function interprets the value as the number of days before December 31, 1899. The following WHERE clause specifies an INTEGER value for the non-DATE expression:

```
WHERE order_date < DATE(365)
```

The database server searches for rows with an **order_date** value less than December 31, 1900 (which is 12/31/1899 plus 365 days).

DAY Function:

The **DAY** function takes a DATE or DATETIME argument and returns the day of the month as an integer in the range from 1 to the number of days in the current month. The following statement fragment calls the **DAY** function with the **CURRENT** function as its argument to compare **order_date** column values to the current day of the month:

```
WHERE DAY(order_date) > DAY(CURRENT)
```

MONTH Function:

The **MONTH** function returns an integer corresponding to the *month* portion of its DATE or DATETIME argument. The following example returns a number that can range from 1 through 12 to indicate the month when the order was placed:

```
SELECT order_num, MONTH(order_date) FROM orders;
```

WEEKDAY Function:

The **WEEKDAY** function takes a DATE or DATETIME argument and returns an integer in the range from 0 to 6 that represents the day of the week. Zero (0) represents Sunday, one (1) represents Monday, and so on. The following example lists all the orders that were paid on the same day of the week as the current day:

```
SELECT * FROM orders
WHERE WEEKDAY(paid_date) = WEEKDAY(CURRENT);
```

YEAR Function:

The **YEAR** function takes a DATE or DATETIME argument and returns a four-digit integer that represents the year.

The following example lists orders in which the **ship_date** is earlier than the beginning of the current year:

```
SELECT order_num, customer_num FROM orders
WHERE year(ship_date) < YEAR(TODAY);
```

Similarly, because a DATE value is a simple calendar date, you cannot add or subtract a DATE value with an INTERVAL value whose *last* qualifier is smaller than DAY. In this case, convert the DATE value to a DATETIME value.

MONTHS_BETWEEN Function:

The **MONTHS_BETWEEN** function requires two arguments, each of which can be a DATE or DATETIME expression.

The value returned is a DECIMAL data type, representing the difference between the two arguments, expressed as a DECIMAL value in units based on 31-day months. If the first argument is a point in time later than the second argument, the sign of the returned value is positive. If the first argument is earlier than the second argument, the sign of the returned value is negative

If the dates of the arguments are both the same days of a month or are both the last days of a months, the result is a whole number. Otherwise, the fractional portion of the result is calculated, based on a month of 31 days . This fractional part can also include the difference in *hour*, *minute*, and *second* time units, unless both arguments are DATE expressions.

The following query calls the **MONTHS_BETWEEN** function in the Projection clause, using two DATE values returned by **TO_DATE** expressions as arguments.

```
SELECT MONTHS_BETWEEN(TO_DATE('2-2-2005', '%m-%d-%Y'),
                     TO_DATE('1-1-2005', '%m-%d-%Y'))
AS months FROM systables WHERE tabid = 1;
```

The value returned by the query expresses the 32-day difference between the two DATE arguments as a positive number of 31-day months:

```
months
1.03225806451613
```

The next example returns the DATETIME column expression arguments to MONTHS_BETWEEN expressions, and their differences in months for two rows of a table:

```
SELECT d_datetime, e_datetime,
       MONTHS_BETWEEN(d_datetime, e_datetime) AS months_between
FROM mytab1;
```

```
d_datetime      2007-11-01 09:00:00.00000
e_datetime      2007-12-07 14:30:12.12345
months_between  -1.2009453405018
```

```
d_datetime      2007-12-13 09:40:30.00000
e_datetime      2007-11-13 08:40:30.00000
months_between  1.000000000000000
```

Here the first MONTHS_BETWEEN result includes differences in time units smaller than days. The second result has no fractional part, because the *day* time units of the arguments had the same value.

The MONTHS_BETWEEN expressions in the next example compares DATE and DATETIME values:

```
SELECT col_datetime, col_date,
       MONTHS_BETWEEN(col_datetime, col_date) AS months_between
FROM mytab2;
```

```
col_datetime    2008-12-13 08:40:30.00000
col_date        11/13/2007
months_between  13.000000000000000
```

Because both arguments specify the same day of the month, the result has no fractional part.

LAST_DAY Function:

The LAST_DAY function requires a DATE or DATETIME expression as its only argument. It returns the date of the last day of the month that its argument specifies. The data type of this returned value is the same data type as the argument. The difference between the returned value and the argument is the number of days remaining in that month.

The following query returns the DATE representation of the current date, the date of the last day in the current month, and the integer number of days (calculated by subtracting the first DATE value from second) before the last day in the current month:

```
SELECT TODAY AS today, LAST_DAY(TODAY) AS last,
       LAST_DAY(TODAY) - TODAY AS days_left
FROM systables WHERE tabid = 1;
```

If the query were issued on 12 April 2008, with MDY4/ as the DBDATE setting for the default locale, it would return the following information:

```

today      last      days_left
03/12/2008 03/31/2008      19

```

In the SELECT statement of this example, there is no name conflict in the Projection clause between the TODAY operator and the identifier **today**, because the AS keyword indicates to Dynamic Server that **today** is a display label.

If you use a host variable to store the argument to **LAST_DAY**, but the data type of the argument is not known at prepare time, Dynamic Server assumes that the data type is DATETIME YEAR TO FRACTION(5). If at runtime, after the statement has been prepared, the user supplies a DATE value for the host variable, error -9750 is issued. To prevent this error, specify the data type of the host variable by using a cast, as in this program fragment:

```

sprintf(query, ",
  \"select last_day(?::date) from mytab\");
EXEC SQL prepare selectq from :query;
EXEC SQL declare select_cursor cursor for selectq;
EXEC SQL open select_cursor using :hostvar_date_input;

EXEC SQL fetch select_cursor into :var_date_output;

```

NEXT_DAY Function:

The **NEXT_DAY** function requires a DATE or DATETIME expression as its first argument, and requires a second *weekday* argument that is a quoted string representing the abbreviation of the English name for a day of the week. Successful execution of this function returns the earliest calendar date that satisfies each of two conditions:

- The date is later than the date specified by the first argument.
- The date falls on the day of the week specified by the second argument.

NEXT_DAY accepts the following abbreviation strings for days of the week:

Table 4-7. Weekday abbreviations valid as arguments to NEXT_DAY function

Day of Week	Abbreviation	Day of Week	Abbreviation
Sunday	'SUN'	Wednesday	'WED'
Monday	'MON'	Thursday	'THU'
Tuesday	'TUE'	Friday	'FRI'
		Saturday	'SAT'

Any characters that follow the 3rd character of these abbreviation strings are ignored. For example, both 'MONDAY' and 'MONTAG' are valid specification for the 2nd argument, each specifying the next Monday after the date in the first argument. Dynamic Server issues an error, however, if the second argument is a string such as 'MODNAY' whose first three characters do not match one of the *weekday* abbreviations in Table 4-7.

The following query, for example, includes a valid **NEXT_DAY** expression:

```

SELECT ship_date, NEXT_DAY(ship_date, 'SAT') AS next_saturday,
       NEXT_DAY(ship_date, 'SAT') - ship_date AS num_days FROM orders;

```

The result set of this query might include the following data from the **orders** table:

ship_date	next_saturday	num_days
06/01/2006	06/03/2006	2
02/12/2007	02/17/2007	5
05/31/2007	06/02/2007	2
05/23/2007	05/26/2007	3

The value returned by **NEXT_DAY** has the same data type as the first argument. If this argument is a DATE type, **NEXT_DAY** returns a DATE value. If the first argument is a DATETIME type, **NEXT_DAY** returns a DATETIME YEAR TO FRACTION(5) value.

Because **ship_date** in the preceding example is a DATE column, the returned dates are formatted as DATE values, rather than in DATETIME format.

If you use a host variable to store the argument to **NEXT_DAY**, but the data type of the argument is not known at prepare time, Dynamic Server assumes that the data type is DATETIME YEAR TO FRACTION(5). If at runtime, after the statement has been prepared, the user supplies a DATE value for the host variable, error -9750 is issued. To prevent this error, specify the data type of the host variable by using a cast, as in this program fragment:

```

sprintf(query, "
  "select next_day(?:::date, 'SUN') from mytab");
EXEC SQL prepare selectq from :query;
EXEC SQL declare select_cursor cursor for selectq;
EXEC SQL open select_cursor using :hostvar_date_input;

EXEC SQL fetch select_cursor into :var_date_output;

```

EXTEND Function:

The **EXTEND** function adjusts the precision of a DATETIME or DATE value. The expression cannot be a quoted string representation of a DATE value.

If you do not specify *first* and *last* qualifiers, the default qualifiers are YEAR TO FRACTION(3).

If the expression contains fields that are not specified by the qualifiers, those fields are discarded.

If the *first* qualifier specifies a larger (that is, more significant) time unit than what exists in the expression, the new fields are filled in with values returned by the **CURRENT** function. If the *last* qualifier specifies a smaller time unit (that is, less significant) than what exists in the expression, the new fields are filled in with constant values. A missing MONTH or DAY field is filled in with 1, and the missing HOUR to FRACTION fields are filled in with 0.

In the following expression, the **EXTEND** call returns the **call_dtime** column value with YEAR TO SECOND precision:

```
EXTEND (call_dtime, YEAR TO SECOND)
```

You can use the **EXTEND** function to perform addition or subtraction with a DATETIME value and an INTERVAL value that do not have the same time unit qualifiers. The next expression expands a literal DATETIME YEAR TO DAY value to a precision of YEAR TO MINUTE so that an interval YEAR TO MINUTE value can be subtracted from it:

```

EXTEND (DATETIME (2009-8-1) YEAR TO DAY, YEAR TO MINUTE)
  - INTERVAL (720) MINUTE (3) TO MINUTE

```

You can use the **EXTEND** function to selectively update a subset of the time units in a DATETIME value. The UPDATE statement in the next example updates only the *hour* and *minute* time unit values in a DATETIME YEAR TO MINUTE column.

```
UPDATE cust_calls SET call_dtime = call_dtime -  
    (EXTEND(call_dtime, HOUR TO MINUTE) - DATETIME (11:00)  
    HOUR TO MINUTE) WHERE customer_num = 106;
```

Subtracting 11:00 from the DATETIME HOUR TO MINUTE value returned by **EXTEND** yields a positive or negative INTERVAL HOUR TO MINUTE value. Subtracting this difference from the original value in the **call_dtime** column forces the updated *hour* and *minute* time unit values to 11:00 in the **cust_calls.call_dtime** column.

MDY Function:

The **MDY** function takes as its arguments three integer expressions that represent the *month*, *day*, and *year*, and returns a type DATE value.

- The first argument represents the number of the month (1 to 12).
- The second argument represents the number of the day of the month (1 to 28, 29, 30, or 31, as appropriate for the month)
- The third expression represents the 4-digit year. You cannot use a 2-digit abbreviation.

The following example finds the **paid_date** for the order number 8052 and sets its value to the first day of the present month:

```
UPDATE orders SET paid_date = MDY(MONTH(TODAY), 1, YEAR(TODAY))  
    WHERE po_num = '8052';
```

TO_CHAR Function (IDS):

The **TO_CHAR** function converts an expression that evaluates to a DATE, DATETIME or numeric value to a character string. The returned character string represents the data value that was specified in the first argument, using a formatting mask that the second argument, a *format_string* argument specifies.

The first argument to this function must be of a DATE, DATETIME, or built-in numeric data type, or a character string that can be converted to one of these data types. If the value of the initial DATE, DATETIME or numeric argument is NULL, the function returns a NULL value.

The second argument to this function is a character string that specifies a formatting mask. What characters that are appropriate for the formatting mask primarily depend on whether the first argument to the **TO_CHAR** function represents a point in time or a number.

- Formatting DATE and DATETIME expressions

The *format_string* argument does not need to imply the same time units as the value in the first argument to the **TO_CHAR** function. When the precision implied in the *format_string* is different from the DATETIME qualifier in first argument, the **TO_CHAR** function extends the DATETIME value as if it had called the **EXTEND** function.

In the following example, the user wants to convert the **begin_date** column of the **tab1** table to a character string. The **begin_date** column is defined as a DATETIME

YEAR TO SECOND data type. The user uses a SELECT statement with the **TO_CHAR** function to perform this conversion:

```
SELECT TO_CHAR(begin_date, '%A %B %d, %Y %R') FROM tab1;
```

The symbols in the *format_string* parameter in this example have the following meanings.

Symbol	Meaning
%A	Full weekday name as defined in the locale
%B	Full month name as defined in the locale
%d	Day of the month as a decimal number
%Y	Year as a 4-digit decimal number
%R	Time in 24-hour notation

Note that the comma (,) that follows the %d specification in this example is a literal character, rather than a separator of arguments to the **TO_CHAR** function. For a complete list of DATE and DATETIME format symbols and their meanings, see the **GL_DATE** and **GL_DATETIME** environment variables in the *IBM Informix GLS User's Guide*.

Applying this *format_string* to the **begin_date** column returns this result:

```
Wednesday July 25, 2007 18:45
```

The query in the next example calls **TO_CHAR** to apply the same format string to an **ADD_MONTHS** expression, and shows the results of the query:

```
SELECT ship_date, TO_CHAR(ADD_MONTHS(ship_date, 1), '%A %B %d, %Y')
AS survey_date FROM orders;
```

```
ship_date    03/12/2007
survey_date  Thursday April 12, 2007
```

If you omit the *format_string* argument when a DATETIME or DATE expression is the first argument, the **TO_CHAR** function uses as a default the setting of the **DBTIME** or **DBDATE** environment variables to format the value represented in the first argument. In non-default locales, the default format for DATETIME and DATE values is specified by environment variables such as **GL_DATETIME** and **GL_DATE**.

- Formatting numeric and MONEY expressions

The *format_string* argument to the **TO_CHAR** function supports the same numeric formatting masks that are used for ESQL functions like **rfmtdec()**, **rfmtdouble()**, and **rfmtlong()**. A detailed description of the IDS numeric-formatting masks for numeric values (when formatting numeric expressions as strings) is in the *IBM Informix ESQL/C Programmer's Manual*. Below is a short summary description of the formatting masks.

A numeric-formatting mask specifies a format to apply to some numeric value when formatting a numeric expression as a string. This mask is a combination of the following formatting characters:

Symbol	Meaning
--------	---------

- * This character fills with asterisks any positions in the display field that would otherwise be blank
- & This character fills with zeros any positions in the display field that would otherwise be blank
- # This character changes leading zeros to blanks. Use this character to specify the maximum leftward extent of a field.
- < This character left-justifies the numbers in the display field. It changes leading zeros to a null string.
- ,
- This character indicates the symbol that separates groups of three digits (counting leftward from the units position) in the whole-number part of the value. By default, this symbol is a comma. You can set the symbol with the **DBMONEY** environment variable. In a formatted number, this symbol appears only if the whole-number part of the value has four or more digits.
- .
- This character indicates the symbol that separates the whole-number part of a money value from the fractional part. By default, this symbol is a period. You can set the symbol with the **DBMONEY** environment variable. You can have only one period in a format string.
-
- This character is a literal. It appears as a minus sign when `expr1` is less than zero. When you group several minus signs in a row, a single minus sign floats to the rightmost position that it can occupy; it does not interfere with the number and its currency symbol.
- +
- This character is a literal. It appears as a plus sign when `expr1` is greater than or equal to zero and as a minus sign when `expr1` is less than zero. When you group several plus signs in a row, a single plus or minus sign floats to the rightmost position that it can occupy; it does not interfere with the number and its currency symbol
- (
- This character is a literal. It appears as a left parenthesis to the left of a negative number. It is one of the pair of accounting parentheses that replace a minus sign for a negative number. When you group several in a row, a single left parenthesis floats to the rightmost position that it can occupy; it does not interfere with the number and its currency symbol.
-)
- This is one of the pair of accounting parentheses that replace a minus sign for a negative value.
- \$
- This character displays the currency symbol that appears at the front of the numeric value. By default, the currency symbol is the dollar sign (\$). You can set the currency symbol with the **DBMONEY** environment variable. When you group several dollar signs in a row, a single currency symbol floats to the rightmost position that it can occupy; it does not interfere with the number.

Any other characters in the formatting mask are reproduced literally in the formatted value that the **TO_CHAR** function returns.

In the next three examples, the value of the **d_int** column expression argument to the **TO_CHAR** function is -12344455.

This query specifies no formatting mask in a call to **TO_CHAR**:

```
SELECT TO_CHAR(d_int) FROM tab_numbers;
```

The following table shows the output of this SELECT statement.

(expression)
-12344455

The following query specifies a monetary format mask:

```
SELECT TO_CHAR(d_int, "$*****.**") FROM tab_numbers;
```

The following table shows the output of this SELECT statement.

(expression)
\$12344455.00

```
SELECT TO_CHAR(d_int, "-$*****.**") FROM tab_numbers;
```

The query returns - \$12344455.00.

```
SELECT TO_CHAR(12344455, "-$*****.**") FROM tab_numbers;
```

The following table shows the output of this SELECT statement.

(constant)
\$12344455.00

The currency (\$) symbol from the formatting mask argument is applied, but the minus (-) symbol has no effect, because the value of the first argument is greater than zero.

Note that the **TO_CHAR** function is a time expression only when its first argument is a DATE or DATETIME expression, or is a character string that can be formatted as a DATE or DATETIME expression. When a numeric or monetary value is its first argument, **TO_CHAR** returns a representation of the value of that argument as a character string, but it does not return a time expression.

TO_DATE Function (IDS):

The **TO_DATE** function converts a character string to a DATETIME value. The function evaluates the *char_expression* parameter as a date according to the date format you specify in the *format_string* parameter and returns the equivalent date. If *char_expression* is NULL, then a NULL value is returned.

Any argument to the **TO_DATE** function must be of a built-in data type.

If you omit the *format_string* parameter, the **TO_DATE** function applies the default DATETIME format to the DATETIME value. The default DATETIME format is specified by the **GL_DATETIME** environment variable.

In the following example, the user wants to convert a character string to a DATETIME value in order to update the **begin_date** column of the **tab1** table with the converted value. The **begin_date** column is defined as a DATETIME YEAR TO SECOND data type. The user uses an UPDATE statement that contains a **TO_DATE** function to accomplish this result:

```
UPDATE tab1
  SET begin_date = TO_DATE('Wednesday July 25, 2007 18:45',
    '%A %B %d, %Y');
```

The *format_string* parameter in this example tells the **TO_DATE** function how to format the converted character string in the **begin_date** column. For a table that shows the meaning of each format symbol in this format string, see “**TO_CHAR Function (IDS)**” on page 4-135.

TO_NUMBER Function

The **TO_NUMBER** function can convert a number or a character expression representing a number value to a DECIMAL data type.

TO_NUMBER Function:

```
TO_NUMBER(char_expr | num_expr)
```

Element	Description	Restrictions	Syntax
<i>char_expression</i>	Expression to be converted to a DECIMAL value	Must be a literal, host variable, expression, or column of a character data type	“Expression” on page 4-40
<i>num_expression</i>	Expression that evaluates to a real number	Must return a numeric data type	“Expression” on page 4-40

The **TO_NUMBER** function converts its argument to a DECIMAL data type. The argument can be the character string representation of a number or a numeric expression.

The following example retrieves a DECIMAL value that the **TO_NUMBER** function returns from the literal representation of a MONEY value:

```
SELECT TO_NUMBER('$100.00') from mytab;
```

The following table shows the output of this SELECT statement.

(expression)
100.000000000000

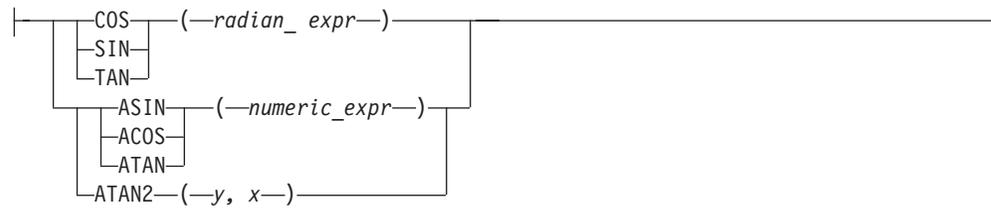
In this example, the currency symbol is discarded from the '\$100.00' string.

The **TO_NUMBER** function is not required in most contexts, because by default, Dynamic Server converts numbers that include a decimal point (and quoted strings in the format of a literal number that has a decimal point) to a DECIMAL data type. This function can be useful, however, when you are migrating SQL applications that were originally written for other database servers, if the application makes calls to a function of this name that returns a DECIMAL value.

Trigonometric Functions

The built-in trigonometric functions have the following syntax.

Trigonometric Functions:



Element	Description	Restrictions	Syntax
<i>numeric_expr</i>	Expression that serves as an argument to the ASIN, ACOS, or ATAN functions	Must return a value between -1 and 1, inclusive	"Expression" on page 4-40
<i>radian_expr</i>	Expression that represents the number of radians	Must return a numeric value	"Expression" on page 4-40
<i>x</i>	Expression that represents the x coordinate of the rectangular coordinate pair (<i>x</i> , <i>y</i>)	Must return a numeric value	"Expression" on page 4-40
<i>y</i>	Expression that represents the y coordinate of the rectangular coordinate pair (<i>x</i> , <i>y</i>)	Must return a numeric value	"Expression" on page 4-40

Formulas for Radian Expressions: The **COS**, **SIN**, and **TAN** functions take the number of radians (*radian_expr*) as an argument. If you are using degrees and want to convert degrees to radians, use the following formula:

degrees * p/180 = # radians

To convert radians to degrees, use the following formula:

radians * 180/p = # degrees

COS Function:

The **COS** function returns the cosine of a radian expression. The following example returns the cosine of the values of the degrees column in the **anglestbl** table. The expression passed to the **COS** function in this example converts degrees to radians.

```
SELECT COS(degrees*180/3.1416) FROM anglestbl;
```

SIN Function:

The **SIN** function returns the sine of a radian expression. This example returns the sine of the values in the **radians** column of the **anglestbl** table:

```
SELECT SIN(radians) FROM anglestbl;
```

TAN Function:

The **TAN** function returns the tangent of a radian expression. This example returns the tangent of the values in the **radians** column of the **anglestbl** table:

```
SELECT TAN(radians) FROM anglestbl;
```

ACOS Function:

The **ACOS** function returns the arc cosine of a numeric expression. The following example returns the arc cosine of the value (-0.73) in radians:

```
SELECT ACOS(-0.73) FROM anglestbl;
```

ASIN Function:

The **ASIN** function returns the arc sine of a numeric expression. The following example returns the arc sine of the value (-0.73) in radians:

```
SELECT ASIN(-0.73) FROM anglestbl;
```

ATAN Function:

The **ATAN** function returns the arc tangent of a numeric expression. The following example returns the arc tangent of the value (-0.73) in radians:

```
SELECT ATAN(-0.73) FROM anglestbl;
```

ATAN2 Function:

The **ATAN2** function computes the angular component of the polar coordinates (r , q) associated with (x, y) . The following example compares *angles* to q for the rectangular coordinates (4, 5):

```
WHERE angles > ATAN2(4,5)    --determines q for (4,5) and  
                             --compares to angles
```

You can determine the length of the radial coordinate r using the expression that the following example shows:

```
SQRT(POW(x,2) + POW(y,2))    --determines r for (x,y)
```

You can determine the length of the radial coordinate r for the rectangular coordinates (4,5) using the expression that the following example shows:

```
SQRT(POW(4,2) + POW(5,2))    --determines r for (4,5)
```

String-Manipulation Functions

String-manipulation functions perform various operations on strings of characters. The syntax for string-manipulation functions is as follows.

String-Manipulation Functions:

	(1)	
CONCAT Function	(2)	
ASCII Function	(3)	
TRIM Function	(4)	(5)
LTRIM Function	(4)	(6)
RTRIM Function	(4)	(7)
SUBSTRING Function	(4)	(8)
SUBSTR Function	(4)	(9)
REPLACE Function	(4)	(10)
LPAD Function	(4)	(11)
RPAD Function	(4)	(12)
Case-Conversion Functions		

Notes:

- 1 See "CONCAT Function"
- 2 See "ASCII Function" on page 4-146
- 3 See " TRIM Function" on page 4-147
- 4 Informix extension
- 5 See "LTRIM Function" on page 4-149
- 6 See "RTRIM Function" on page 4-149
- 7 See "SUBSTRING Function" on page 4-150
- 8 See " SUBSTR Function" on page 4-152
- 9 See "REPLACE Function" on page 4-153
- 10 See " LPAD Function" on page 4-154
- 11 See " RPAD Function" on page 4-155
- 12 See " Case-Conversion Functions" on page 4-156

Sections that follow describe each of the built-in string manipulation functions.

CONCAT Function:

The **CONCAT** function accepts two expressions as arguments, and returns a single character that appends the string representation of the value returned by its second argument to the string representation of the value returned by its first argument.

CONCAT Function:

|—CONCAT—(—*expr_1*—,—*expr_2*—)———|

Element	Description	Restrictions	Syntax
<i>expr_1</i> , <i>expr_2</i>	Expressions whose string representations of their values are to be concatenated	Cannot return a complex, user-defined, or large object type. If a host variable, it must be long enough to store the resulting combined strings.	"Expression" on page 4-40

Each arguments to the **CONCAT** function can evaluate to a character, number, or time data type. If one of its arguments returns a NULL value, the function returns the string representation of its other argument.

Unlike other built-in string manipulation functions of Dynamic Server, the **CONCAT** function cannot be overloaded.

CONCAT is the operator function of the concatenation (||) operator, For a given pair of expression arguments, **CONCAT** returns the same string as that operator returns from the same expressions as operands. See "Concatenation Operator" on page 4-56 for additional information about concatenation operations, and for restrictions on the SQL and Dynamic SQL statements in which you can invoke the **CONCAT** function.

Return Types from CONCAT and String Functions:

The data type of the return value from a successful call to the **CONCAT** function (or to other built-in string-manipulation functions that follow the same rules as **CONCAT** for determining their return type) depends on the data types of the arguments and on the length of the resulting string. The order of the two arguments is not significant in determining the return type.

Dynamic Server applies the following rules for the return type from operations that concatenate values that arguments of more than one data type specify

- If one of the types is National Language Support (namely **NVARCHAR**, **NCHAR**):
 - the return type is **NVARCHAR** if the resulting length is less than 255 bytes
 - the return type is **NCHAR** otherwise.
- If one of the arguments is **VARCHAR** or a number type,
 - the return type is **NVARCHAR** if the resulting length is less than 255 bytes
 - the return type is **LVARCHAR** otherwise.
- An exception to these rules, however, can occur in certain cross-server operations in which a remote routine is executed locally, and a concatenation expression is evaluated locally before its return value is sent to a remote database server. For remote servers that do not support the **LVARCHAR** data type in distributed transactions, the concatenated result is sent as a **CHAR** data type if sending the **LVARCHAR** type returns an error. Dynamic Server instances earlier than Version 11.10 require a **CHAR** return value in this scenario.

(See also “Return String Types in Distributed Transactions” on page 4-145 for the data types that can be returned from concatenation expressions that are evaluated by remote Dynamic Server instances earlier than Version 11.50.xC2.)

In the following table, the rows list the valid data types of the first argument to the **CONCAT** function, and the columns list the type of the second argument. The cell at the intersection of each row and column shows the possible returned type or types. The row and the column labelled as **Other** represent arguments that evaluate to non-character data types, such as numeric or time types like **DECIMAL** or **DATE**.

Table 4-8. Return Types from Operations on Two Arguments (in Version 11.50.xC2)

	NCHAR	NVARCHAR	CHAR	VARCHAR	LVARCHAR	Other
NCHAR	nchar	nvarchar or nchar	nchar	nvarchar or nchar	nvarchar or nchar	nvarchar or nchar
NVARCHAR	nvarchar or nchar	nvarchar or nchar	nvarchar or nchar	nvarchar or nchar	nvarchar or nchar	nvarchar or nchar
CHAR	nchar	nvarchar or nchar	char	varchar or lvarchar	lvarchar	varchar or lvarchar
VARCHAR	nvarchar or nchar	nvarchar or nchar	varchar or lvarchar	varchar or lvarchar	lvarchar	varchar or lvarchar
LVARCHAR	nvarchar or nchar	nvarchar or nchar	lvarchar	lvarchar	lvarchar	lvarchar
Other	nvarchar or nchar	nvarchar or nchar	varchar or lvarchar	varchar or lvarchar	lvarchar	varchar or lvarchar

For string manipulation functions other than **CONCAT**, arguments of **DATE**, **DATETIME**, or **MONEY** data types always return an **NVARCHAR** or **NCHAR** value, depending on the length of the resulting string.

This table is symmetrical, because the order of arguments has no effect on the return data type. User-defined data types, large-object types, complex types, and other extended data types are not valid as arguments to the built-in string-manipulation functions or operators.

This table also describes the return data types of expressions that use the concatenation (||) operator.

Not shown here is the result of concatenation operations in which the sum of the argument lengths exceeds the approximately 32Kb limit for CHAR, NCHAR, and LVARCHAR data types. This returns error -881, rather than a concatenated data value. Because the maximum LVARCHAR size is 32,739 bytes, and the CHAR and NCHAR limits are both 32,767 bytes, error -881 is usually associated with VARCHAR and NVARCHAR objects, whose limit is 255 bytes, but automatic return type promotion can reduce the incidence of this error.

The following string-manipulation functions support the same rules as **CONCAT** for return type promotion:

- **LPAD**
- **RPAD**
- **REPLACE**
- **SUBSTR**
- **SUBSTRING**
- **TRIM**
- **LTRIM**
- **RTRIM**

The following table summarizes how Dynamic Server determines the return type from these string manipulating functions, based on the argument types:

Table 4-9. String Manipulation Functions that Support Return Type Promotion

Function	How the Return Type of the Function is Determined
CONCAT, 	Return type is based on both arguments. Refer to Table 4-8 on page 4-143..
SUBSTR, SUBSTRING	Return type is the same as the <i>source string</i> type. If <i>source string</i> is a host variable, the return type is NVARCHAR or NCHAR, depending on the length of the result.
TRIM, LTRIM, RTRIM	Return type depends on the source type and the returned length: <ul style="list-style-type: none"> • NVARCHAR <i>returns</i> NVARCHAR • VARCHAR <i>returns</i> VARCHAR • CHAR <i>returns</i> VARCHAR (if length <= 255 bytes) • CHAR <i>returns</i> LVARCHAR (if length > 255 bytes) • NCHAR <i>returns</i> NVARCHAR (if length <= 255 bytes) • NCHAR <i>returns</i> LVARCHAR (if length > 255 bytes) • LVARCHAR <i>returns</i> LVARCHAR
LPAD, RPAD	Return type is based on the <i>source_string</i> and <i>pad_string</i> arguments. If <i>pad_string</i> is not specified, the return type is based on the data type of <i>source_string</i> .
REPLACE	Return type is based on the <i>source_string</i> and <i>old_string</i> arguments (and on the <i>new_string</i> argument, if that is specified). If any argument is a host variable, the return type is NCHAR.

Table 4-9. String Manipulation Functions that Support Return Type Promotion (continued)

Function	How the Return Type of the Function is Determined
ENCRYPT_AES, ENCRYPT_TDES, DECRYPT_BINARY, DECRYPT_CHAR,	For arguments that are not BLOB or CLOB variables, the return type is based on the data types of the <i>data</i> and <i>encrypted_data</i> arguments. Refer to Table 4-8 on page 4-143.

Return String Types in Distributed Transactions:

In cross-database distributed queries that access tables in different databases of the same Dynamic Server instance, the same types are returned by **CONCAT** (and by other built-in string manipulation functions that follow the same rules for return type promotion) that the section Return Types from the **CONCAT** Function describes.

The same types are also returned in cross-server distributed queries, if every participating Dynamic Server instance that evaluates these string-manipulation function expressions is no earlier than Version 11.50.xC2.

For cross-server distributed operations in which the return value is evaluated on a remote Dynamic Server instance earlier than Version 11.50.xC2, the following table (in the same format as the table for Version 11.50.xC2) lists the possible return data types (or -881 overflow error) for the specified data types of the arguments to the string-manipulation function (or for the operands of the concatenation (||) operator:

Table 4-10. Return Types from Distributed Operations (in Version 11.50.xC1 and earlier)

	NCHAR	NVARCHAR	CHAR	VARCHAR	LVARCHAR	Other
NCHAR	nchar	nvarchar or EM -881	nchar	nchar	nchar	nchar
NVARCHAR	nchar or EM -881	nvarchar or EM -881				
CHAR	nchar	nvarchar or EM -881	char	varchar or EM -881	char	char
VARCHAR	nchar	nvarchar or EM -881	varchar or EM -881	varchar or EM -881	varchar or EM -881	varchar or EM -881
LVARCHAR	nchar	nvarchar or EM -881	char	varchar or EM -881	char	char
Other	nchar	nvarchar or EM -881	char	varchar or EM -881	char	char

The following are among the differences between the return values in this release and what Dynamic Server versions earlier than 11.50.xC2 return:

- Earlier releases accept **LVARCHAR** arguments, but cannot return **LVARCHAR** values.
- If the result is longer than the maximum size of the argument of the longest data type, earlier releases do not support data-type promotion, but issue error -881. (This is typically with **VARCHAR** or **NVARCHAR** arguments, if the length of the resulting string would be greater than 255 bytes.)

For all versions of Dynamic Server, error -881 is issued if the length of a returned string exceeds 32Kb.

ASCII Function:

The **ASCII** function returns the decimal representation of the first character in a character string, based on its codepoint in the ASCII character set.

ASCII Function:

|—ASCII—(—*char_expr*—)|

Element	Description	Restrictions	Syntax
<i>char_expr</i>	Expression that evaluates to a character data type	Must be of type CHAR, LVARCHAR, NCHAR, NVARCHAR, or VARCHAR	“Identifier ” on page 5-21

The **ASCII** function takes a single argument of any character data type. It returns an integer value, based on the first character of the argument, corresponding to the decimal representation of the codepoint of that character within the ASCII character set.

If the argument is NULL, or if the argument is an empty string, the **ASCII** function returns a NULL value.

The following query returns the ASCII value of uppercase H:

```
SELECT ASCII("HELLO") FROM systables WHERE tabid = 1;
```

The following table shows the output of this SELECT statement.

(constant)
72

The following query returns the ASCII value of lowercase h:

```
SELECT ASCII("hello") FROM systables WHERE tabid = 1;
```

The following table shows the output of this SELECT statement.

(constant)
104

The following query returns the **ASCII** output from an empty string argument:

```
SELECT ASCII("") FROM systables WHERE tabid = 1;
```

The following table shows the NULL output of this SELECT statement.

(constant)

The following query returns the **ASCII** output from a NULL argument:

```
SELECT ASCII(NULL) FROM systables WHERE tabid = 1;
```

The following table shows the NULL output of this SELECT statement.

(constant)

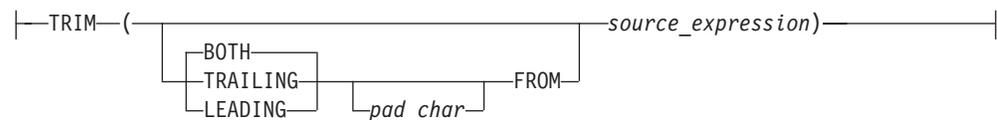
The **ASCII** function interprets this argument as a NULL expression, rather than as a value that begins with uppercase N.

For a table of the numeric values of the codepoints in the ASCII character set, see “Collating Order for U.S. English Data” on page 4-195.

TRIM Function:

The **TRIM** function removes specified leading or trailing pad characters from a string. (See also the descriptions of the **LTRIM** and **RTRIM** functions, which provide similar functionality, but support a different syntax.)

TRIM Function:



Element	Description	Restrictions	Syntax
<i>pad_char</i>	Expression that evaluates to a single character or NULL. The default is a blank space (= ASCII 32).	Must be a character expression	“Expression” on page 4-40
<i>source_expression</i>	Character expression, including a character column name, or a call to another TRIM function	Cannot be a DISTINCT data type	“Expression” on page 4-40

The **TRIM** function returns a character string identical to its *source_expression* argument, except that any leading or trailing pad characters, as specified by the **LEADING**, **TRAILING**, or **BOTH** keywords, are deleted. If no trim qualifier (**LEADING**, **TRAILING**, or **BOTH**) is specified, **BOTH** is the default. If no *pad_char* is specified, a single blank space is the default, and leading or trailing blank spaces, as specified by the qualifying keyword, are deleted from the returned value.

If either the *pad_char* or the *source_expression* evaluates to NULL, the result of the **TRIM** function is NULL.

The data type of the returned value depends on the *source_expression* argument:

- If the argument is longer than 255 bytes, the returned value is of type **LVARCHAR**.
- If the argument has 255 bytes or fewer, the data type of the returned value depends on the data type of the argument:
 - If the argument is of type **CHAR** or **VARCHAR**, a **VARCHAR** value is returned.
 - If the argument is of type **NCHAR** or **NVARCHAR**, an **NVARCHAR** value is returned.
 - If the argument is of type **LVARCHAR**, an **LVARCHAR** value is returned.

The length of the returned value is 255 bytes or fewer for VARCHAR or NVARCHAR *source_expression* arguments, and no more than 32,739 bytes for CHAR, NCHAR, or LVARCHAR arguments.

The following example shows some generic uses for the **TRIM** function:

```
SELECT TRIM (c1) FROM tab;
SELECT TRIM (TRAILING '#' FROM c1) FROM tab;
SELECT TRIM (LEADING FROM c1) FROM tab;
UPDATE c1='xyz' FROM tab WHERE LENGTH(TRIM(c1))=5;
SELECT c1, TRIM(LEADING '#' FROM TRIM(TRAILING '%' FROM
    '###abc%%')) FROM tab;
```

In Dynamic SQL, when you use the DESCRIBE statement with a SELECT statement that calls the **TRIM** function in the projection list, the data type of the trimmed column that DESCRIBE returns depends on the data type of the *source_expression*, for SQL data type constants defined in the **sqltypes.h** header file of the Informix ESQL/C source file. For further information on the GLS aspects of the **TRIM** function in Informix ESQL/C, see the *IBM Informix GLS User's Guide*.

Fixed Character Columns:

The **TRIM** function can be specified on fixed-length character columns. If the length of the string is not completely filled, the unused characters are padded with blank space. Figure 4-3 shows this concept for the column entry '##A2T##', where the column is defined as CHAR(10).

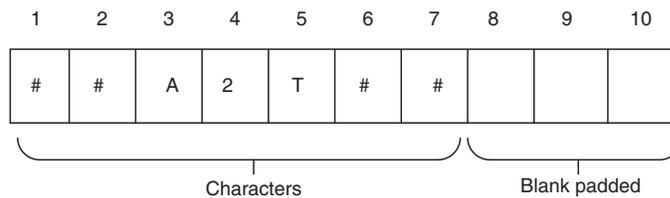


Figure 4-3. Column Entry in a Fixed-Length Character Column

If you want to trim the sharp sign (#) *pad_char* from the column, you need to consider the blank padded spaces as well as the actual characters.

For example, if you specify the keyword **BOTH**, the result from the trim operation is A2T##, because the **TRIM** function does not match the blank padded space that follows the string. In this case, the only sharp signs (#) trimmed are those that precede the other characters. The SELECT statement is shown, followed by Figure 4-4, which presents the result.

```
SELECT TRIM(LEADING '#' FROM col1) FROM taba;
```

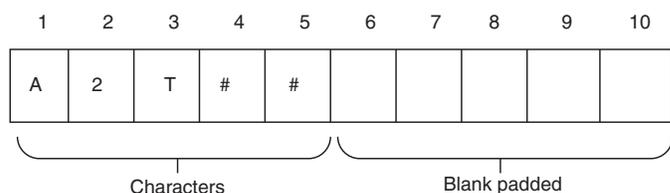


Figure 4-4. Result of TRIM Operation

This SELECT statement removes all occurrences of the sharp (#) sign:

```
SELECT TRIM(BOTH '#' FROM TRIM(TRAILING ' ' FROM col1)) FROM taba;
```

LTRIM Function:

The **LTRIM** function removes specified leading pad characters from a string.

LTRIM Function:

```
|—LTRIM—(—source_string—, —pad_string—)
```

Element	Description	Restrictions	Syntax
<i>pad_string</i>	Expression that specifies one or more characters to delete from <i>source_string</i>	Must be a character expression	“Expression” on page 4-40
<i>source_string</i>	Expression that specifies a character string from which characters in <i>pad_string</i> are deleted	Pad characters to the right of any character not in <i>pad_string</i> are not deleted	“Expression” on page 4-40

The first argument to the **LTRIM** function must be a character expression from which to delete leading pad characters. The optional second argument is a character expression that evaluates to a string of pad characters. If no second argument is provided, only blank characters are regarded as pad characters.

The return data type of the **LTRIM** function is based on its *source_string* argument, using the return type promotion rules that the section Return Types from the CONCAT Function describes.

The value returned contains a substring of *source_string*, but from which any leading pad characters to the left of the first non-pad character have been removed. If a host variable is used, an LVARCHAR data type is returned.

The **LTRIM** function scans a copy of the *source_string* from the left, deleting any leading characters that appear in the *pad_string*. If no *pad_string* argument is specified, only leading blanks are deleted from the returned value. When the first non-pad character is encountered, the function returns its result string and terminates.

In the following example, the *pad_string* is 'Hello':

```
SELECT LTRIM('Hellohello world!', 'Hello') FROM mytab;
```

The following table shows the output of this SELECT statement.

(constant)
hello world!

Here the first five characters of the *source_string* were dropped because they matched characters in the *pad_string*, but the function terminated after it encountered the lowercase h character, which preserved the trailing 'ello' pad characters to its right.

RTRIM Function:

The **RTRIM** function removes specified trailing pad characters from a string.

RTRIM Function:

|—RTRIM—(—*source_string*— [,—*pad_string*—])—|

Element	Description	Restrictions	Syntax
<i>pad_string</i>	Expression that specifies one or more characters to delete from <i>source_string</i>	Must be a character expression	“Expression” on page 4-40
<i>source_string</i>	Expression that specifies a character string from which characters in <i>pad_string</i> are deleted	Pad characters to the left of any character not in <i>pad_string</i> are not deleted	“Expression” on page 4-40

The first argument to the **RTRIM** function must be a character expression from which to delete trailing pad characters. The optional second argument is a character expression that evaluates to a string of pad characters. If no second argument is provided, only blank characters are regarded as pad characters.

The return data type of the **LTRIM** function is based on its *source_string* argument, using the return type promotion rules that the section Return Types from the CONCAT Function describes.

The value returned contains a substring of *source_string*, but from which any trailing pad characters to the right of the first non-pad character have been removed. If a host variable is used, an LVARCHAR data type is returned.

The **RTRIM** function scans a copy of the *source_string* from the right, deleting any trailing characters that appear in the *pad_string*. If no *pad_string* argument is specified, only trailing blanks are deleted from the returned value. When the first non-pad character is encountered, the function returns its result string and terminates.

In the following example, the *pad_string* is ' theend!*#?':

```
SELECT RTRIM('good night... *!#?theend ', ' theend!*#?') AS closing FROM mytab;
```

The following table shows the output of this SELECT statement.

(constant)

good night...

Here the last fifteen characters of the *source_string* were dropped because they matched characters in the *pad_string*, but the function terminated after it encountered the period (.) characters, which preserved the leading 'thn' pad characters to the left.

SUBSTRING Function:

The **SUBSTRING** function returns a subset of a character string.

SUBSTRING Function:

|—SUBSTRING—(—*source_string*—FROM—*start_position*— [FOR—*length*—])—|

Element	Description	Restrictions	Syntax
<i>length</i>	Number of characters to return from <i>source_string</i>	Must be an expression, constant, column, or host variable that returns an integer	“Literal Number ” on page 4-185
<i>source_string</i>	String argument to the SUBSTRING function	Must be an expression, constant, column, or host variable whose value can be converted to a character data type	“Expression” on page 4-40
<i>start_position</i>	Position in <i>source_string</i> of first returned character	Must be an expression, constant, column, or host variable that returns an integer	“Literal Number ” on page 4-185

Any argument to the **SUBSTRING** function must be of a built-in data type.

The return data type is that of the *source_string* argument. If a host variable is the source, the return value is either NVARCHAR or NCHAR, according to the length of the returned string, using the return type promotion rules that the section Return Types from the CONCAT Function describes.

The subset begins at the column position that *start_position* specifies. The following table shows how the database server determines the starting position of the returned subset based on the input value of the *start_position*.

Value of Start_Position	How the Database Server Determines the Starting Position of the Return Subset
Positive	Counts forward from the first character in <i>source_string</i> For example, if <i>start_position</i> = 1, the first character in the <i>source_string</i> is the first character in the returned subset.
Zero (0)	Counts from one position before (that is, to the left of) the first character in <i>source_string</i> For example, if <i>start_position</i> = 0 and <i>length</i> = 1, the database server returns NULL, whereas if <i>length</i> = 2, the database server returns the first character in <i>source_string</i> .
Negative	Counts backward from one position after (that is, to the right of) the last character in <i>source_string</i> For example, if <i>start_position</i> = -1, the starting position of the returned subset is the last character in <i>source_string</i> .

In locales for languages with a right-to-left writing direction, such as Arabic, Farsi, or Hebrew, *right* should replace *left* in the preceding table.

The size of the subset is specified by *length*. The *length* parameter refers to the number of logical characters, rather than to the number of bytes. If you omit the *length* parameter, or if you specify a *length* that is greater than the number of characters from *start_position* to the end of *source_string*, the **SUBSTRING** function returns the entire portion of *source_string* that begins at *start_position*. The following example specifies that the subset of the source string that begins in column position 3 and is two characters long should be returned:

```
SELECT SUBSTRING('ABCDEFGH' FROM 3 FOR 2) FROM mytable;
```

The following table shows the output of this SELECT statement.

(constant)
CD

In the following example, the user specifies a negative *start_position* for the return subset:

```
SELECT SUBSTRING('ABCDEFG' FROM -3 FOR 7)
      FROM mytable;
```

The database server starts at the -3 position (four positions before the first character) and counts forward for 7 characters. The following table shows the output of this SELECT statement.

(constant)
ABC

SUBSTR Function:

The **SUBSTR** function has the same purpose as the **SUBSTRING** function (to return a subset of a source string), but it uses different syntax.

SUBSTR Function:

|—SUBSTR—(—*source_string*—,—*start_position*—,—*length*—)—|

Element	Description	Restrictions	Syntax
<i>length</i>	Number of characters to be returned from <i>source_string</i>	Must be an expression, literal, column, or host variable that returns an integer	“Expression” on page 4-40
<i>source_string</i>	String that serves as input to the SUBSTR function	Must be an expression, literal, column, or host variable of a data type that can be converted to a character data type	“Expression” on page 4-40
<i>start_position</i>	Column position in <i>source_string</i> where the SUBSTR function starts to return characters	Must be an integer expression, literal, column, or host variable. Can have a plus sign (+), a minus sign (-), or no sign.	“Literal Number ” on page 4-185

Any argument to the **SUBSTR** function must be of a built-in data type.

The **SUBSTR** function returns a subset of *source_string*. The subset begins at the column position that *start_position* specifies. The following table shows how the database server determines the starting position of the returned subset based on the input value of the *start_position*.

Value of Start_Position	How the Database Server Determines the Starting Position of the Returned Subset
Positive	Counts forward from the first character in <i>source_string</i>
Zero (0)	Counts forward from the first character in <i>source_string</i> (that is, treats a <i>start_position</i> of 0 as equivalent to 1)

Value of Start_Position	How the Database Server Determines the Starting Position of the Returned Subset
Negative	Counts backward from an origin that immediately follows the last character in <i>source_string</i> . A value of -1 returns the last character in <i>source_string</i> .

The *length* parameter specifies the number of logical characters (not the number of bytes) in the subset. If you omit the *length* parameter, the **SUBSTR** function returns the entire portion of *source_string* that begins at *start_position*.

If you specify a negative *start_position* whose absolute value is greater than the number of characters in *source_string*, or if *length* is greater than the number of characters from *start_position* to the end of *source_string*, **SUBSTR** returns NULL. (In this case, the behavior of **SUBSTR** is different from that of the **SUBSTRING** function, which returns all the characters from *start_position* to the last character of *source_string*, rather than returning NULL.)

The return data type is that of the *source_string* argument. If a host variable is the source, the return value is either NVARCHAR or NCHAR, according to the length of the returned string, using the return type promotion rules that the section Return Types from the CONCAT Function describes.

The following example specifies that the string of characters to be returned begins at a starting position 3 characters before the end of a 7-character *source_string*. This implies that the starting position is the fifth character of *source_string*. Because the user does not specify a value for *length*, the database server returns a string that includes all characters from character-position 5 to the end of *source_string*.

```
SELECT SUBSTR('ABCDEFG', -3)
FROM mytable;
```

The following table shows the output of this SELECT statement.

(constant)
EFG

REPLACE Function:

The **REPLACE** function replaces specified characters within a source string with different characters.

REPLACE Function:

```
|—REPLACE—(—source_string—,—old_string— [,—new_string—] )—|
```

Element	Description	Restrictions	Syntax
<i>new_string</i>	Character or characters that replace <i>old_string</i> in the return string	Must be an expression, constant, column, or host variable of a data type that can be converted to a character data type	“Expression” on page 4-40
<i>old_string</i>	Character or characters in <i>source_string</i> that are to be replaced by <i>new_string</i>	Must be an expression, constant, column, or host variable of a data type that can be converted to a character data type	“Expression” on page 4-40

Element	Description	Restrictions	Syntax
<i>source_string</i>	String of characters argument to the REPLACE function	Must be an expression, constant, column, or host variable of a data type that can be converted to a character data type	"Expression" on page 4-40

Any argument to the **REPLACE** function must be of a built-in data type.

The **REPLACE** function returns a copy of *source_string* in which every occurrence of *old_string* is replaced by *new_string*. If you omit the *new_string* option, every occurrence of *old_string* is omitted from the return string.

The return data type is its *source_string* argument. If a host variable is the source, the return value is either NVARCHAR or NCHAR, according to the length of the returned string, using the return type promotion rules that the section Return Types from the CONCAT Function describes.

In the following example, the **REPLACE** function replaces every occurrence of xz in the source string with t:

```
SELECT REPLACE('Mighxzy xzime', 'xz', 't')
FROM mytable;
```

The following table shows the output of this SELECT statement.

(constant)

Mighty time

LPAD Function:

The **LPAD** function returns a copy of *source_string* that is left-padded to the total number of characters specified by *length*.

LPAD Function:

|—LPAD—(—*source_string*—,—*length*—,—*pad_string*—)—

Element	Description	Restrictions	Syntax
<i>length</i>	Integer value that specifies total number of characters in the returned string	Must be an expression, constant, column, or host variable of a data type that can be converted to an integer data type	"Literal Number " on page 4-185
<i>pad_string</i>	String that specifies the pad character or characters	Must be an expression, constant, column, or host variable of a data type that can be converted to a character data type	"Expression" on page 4-40
<i>source_string</i>	String that serves as input to the LPAD function	Must be an expression, constant, column, or host variable of a data type that can be converted to a character data type	"Expression" on page 4-40

Any argument to the **LPAD** function must be of a built-in data type.

The *pad_string* parameter specifies the character or characters to be used for padding the source string. The sequence of pad characters occurs as many times as necessary to make the return string the length specified by *length*.

The series of pad characters in *pad_string* is truncated if it is too long to fit into *length*. If you specify no *pad_string*, the default value is a single blank.

|
|
|

The return data type is based on the three arguments, using the return type promotion rules that the section Return Types from the CONCAT Function describes.

In the following example, the user specifies that the source string is to be left-padded to a total length of 16 characters. The user also specifies that the pad characters are a series consisting of a hyphen and an underscore (-_).

```
SELECT LPAD('Here we are', 16, '-_') FROM mytable;
```

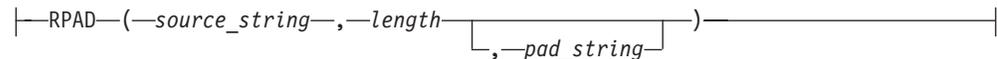
The following table shows the output of this SELECT statement.

(constant)
--Here we are

RPAD Function:

The **RPAD** function returns a copy of *source_string* that is right-padded to the total number of characters that *length* specifies.

RPAD Function:



Element	Description	Restrictions	Syntax
<i>length</i>	The number of characters in the returned string	Must be an expression, constant, column, or host variable that returns an integer	“Literal Number ” on page 4-185
<i>pad_string</i>	String that specifies the pad character or characters	Must be an expression, column, constant, or host variable of a data type that can be converted to a character data type	“Expression” on page 4-40
<i>source_string</i>	String that serves as input to the RPAD function	Same as for <i>pad_string</i>	“Expression” on page 4-40

Any argument to the **RPAD** function must be of a built-in data type.

The *pad_string* parameter specifies the pad character or characters to be used to pad the source string.

The series of pad characters occurs as many times as necessary to make the return string reach the length that *length* specifies. The series of pad characters in *pad_string* is truncated if it is too long to fit into *length*. If you omit the *pad_string* parameter, the default value is a single blank space.

|
|
|
|
|

The return data type is based on the *source_string* and *pad_string* arguments, if both are specified. If a host variable is the source, the return value is either NVARCHAR or NCHAR, according to the length of the returned string, using the return type promotion rules that the section Return Types from the CONCAT Function describes.

In the following example, the user specifies that the source string is to be right-padded to a total length of 18 characters. The user also specifies that the pad characters to be used are a sequence consisting of a question mark and an exclamation point (?!)

```
SELECT RPAD('Where are you', 18, '?!')
      FROM mytable;
```

The following table shows the output of this SELECT statement.

(constant)
Where are you?!?!?

Case-Conversion Functions

The case-conversion functions perform lettercase conversion on alphabetic characters. In the default locale, only the ASCII characters A - Z and a - z can be modified by these functions, which enable you to perform case-insensitive searches in your queries and to specify the format of the output.

The case-conversion functions are **UPPER**, **LOWER**, and **INITCAP**. The following diagram shows the syntax of these case-conversion functions.

Case-Conversion Functions:



Element	Description	Restrictions	Syntax
<i>expression</i>	Expression returning a character string	Must be a character type. If a host variable, its length must be long enough to store the converted string.	"Expression" on page 4-40

The *expression* must return a character data type. When the column is described, the data type returned by the database server is that of *expression*. For example, if the input type is CHAR, the output type is also CHAR.

Argument to these functions must be of the built-in data types.

In all locales, the byte length returned from the description of a column with a case-conversion function is the input byte length of the source string. If you use a case-conversion function with a multibyte *expression* argument, the conversion might increase or decrease the length of the string. If the byte length of the result string exceeds the byte length *expression*, the database server truncates the result string to fit into the byte length of *expression*.

Only characters designated as ALPHA class in the locale file are converted, and this occurs only if the locale recognizes the construct of lettercase.

If *expression* is NULL, the result of a case-conversion function is also NULL.

The database server treats a case-conversion function as an SPL routine in the following instances:

- If it has no argument
- If it has one argument, and that argument is a named argument

- If it has more than one argument
- If it appears in a projection list with a host variable as an argument

If none of the conditions in the preceding list are met, the database server treats a case-conversion function as a system function.

The following example uses all the case-conversion functions in the same query to specify multiple output formats for the same value:

Input value:

SAN Jose

Query:

```
SELECT City, LOWER(City), LOWER("City"),
       UPPER (City), INITCAP(City)
FROM Weather;
```

Query output:

SAN Jose san jose city SAN JOSE San Jose

UPPER Function:

The **UPPER** function accepts an *expression* argument and returns a character string in which every lowercase alphabetical character in the *expression* is replaced by a corresponding uppercase alphabetic character.

The following example uses the **UPPER** function to perform a case-insensitive search on the **lname** column for all employees with the last name of **Curran**:

```
SELECT title, INITCAP(fname), INITCAP(lname) FROM employees
WHERE UPPER (lname) = "CURRAN"
```

Because the **INITCAP** function is specified in the projection list, the database server returns the results in a mixed-case format. For example, the output of one matching row might read: accountant James Curran.

LOWER Function:

The **LOWER** function accepts an *expression* argument and returns a character string in which every uppercase alphabetic character in the *expression* is replaced by a corresponding lowercase alphabetic character.

The following example shows how to use the **LOWER** function to perform a case-insensitive search on the **City** column. This statement directs the database server to replace all instances (that is, any variation) of the words san jose, with the mixed-case format, San Jose.

```
UPDATE Weather SET City = "San Jose"
WHERE LOWER (City) = "san jose";
```

INITCAP Function:

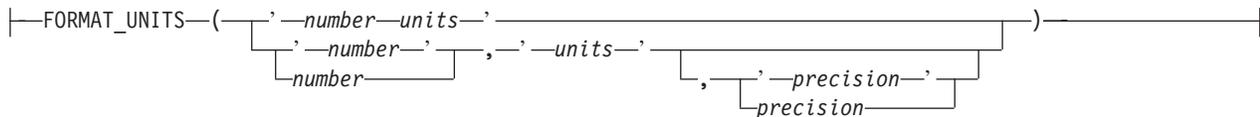
The **INITCAP** function returns a copy of the *expression* in which every word in the *expression* begins with an uppercase letter. With this function, a *word* begins after any character other than a letter. Thus, in addition to a blank space, symbols such as commas, periods, colons, and so on, introduce a new word.

For an example of the **INITCAP** function, see "UPPER Function."

FORMAT_UNITS Function

The **FORMAT_UNITS** function can interpret strings that specify a number and the names of abbreviated units of memory or of storage. This built-in function can accept one, two, or three quoted string arguments, and can be called internally in the **sysmaster** database by the SQL Administration API functions **ADMIN** and **TASK**, but you can invoke **FORMAT_UNITS** in SQL statements that process size specifications expressed with standard abbreviations for bytes or for larger units (such as kilobytes, megabytes, gigabytes, and so forth) of memory or of mass storage.

FORMAT_UNITS Function:



Element	Description	Restrictions	Syntax
<i>number</i>	Expression that evaluates to the number of storage or memory <i>units</i>	Must be a literal number or a quoted string specifying a number that can be converted to FLOAT	"Expression" on page 4-40
<i>precision</i>	Integer number of significant digits to return from <i>number</i>	Must be a literal number or a quoted string specifying an integer	"Expression" on page 4-40
<i>units</i>	Abbreviation of a unit of storage or memory; the default is 'B' (for bytes)	Must begin with 'B', 'K', 'M', 'G', 'T', 'P', or 'PB' (or the lowercase form of these letters). Any trailing characters are ignored.	"Quoted String" on page 4-189

This built-in function can accept one, two, or three arguments. The returned value is a character string that shows the specified *number* and an appropriate format label that shows the storage units. If you specify a *precision* as the last argument, the *number* is returned with that precision. Otherwise, the *number* is formatted to precision 3 (%3.31f) by default.

The same notation also applies to arguments to all commands of the **ADMIN** or **TASK** function of the SQL Administration API (except for the Enterprise Replication cdr utility commands) as described in "Mnemonic Initial Commands" on page 6-6, that specify sizes of memory, of disk storage, or of offsets:

Notation

Corresponding Units

'B' or 'b'

Bytes (= 2 to the power 0)

'K' or 'k'

Kilobytes (= 2 to the power 10)

'M' or 'm'

Megabytes (= 2 to the power 20)

'G' or 'g'

Gigabytes (= 2 to the power 30)

'T' or 't'

Terabytes (= 2 to the power 40)

'PB'

Petabytes (= 2 to the power 50)

'P' Pages (= 2 kilobytes or 4 kilobytes, depending on the base page size of the system)

The initial letter in the *unit* specification ('B', 'K', 'M', 'G' or 'T') determines the units of measure, and any trailing characters are ignored. An exception, however, is if the initial letter 'P' (or 'p') is immediately followed by 'B' or 'b' in the string, because in this case the string is interpreted as petabytes. Any other string starting with "P" (such as "PA", "pc", "PhD", "papyrus" and so forth) is interpreted as specifying *pages*, rather than *petabytes*.

If one argument provides both the *number* and *units* specifications, Dynamic Server ignores any whitespace that separates the *number* specification from the *units* specification within the same argument to the **FORMAT_UNITS**, **ADMIN**, or **TASK** functions. For example, the specifications '128M' and '128 m' are both interpreted as 128 megabytes.

The following examples invoke the **FORMAT_UNITS** function with a single argument:

```
EXECUTE FUNCTION FORMAT_UNITS('1024 M');
```

The following character string value is returned.

(expression)

1.00 GB

```
SELECT FORMAT_UNITS('1024 k') FROM systables WHERE tabid=1;
```

The following character string value is returned.

(expression)

1.00 MB

```
SELECT FORMAT_UNITS(tabid || 'M') FROM systables WHERE tabid=100;
```

The following character string value is returned.

(expression)

100 MB

The following examples show calls to the **FORMAT_UNITS** function with two arguments:

```
EXECUTE FUNCTION FORMAT_UNITS(1024, 'k');
```

The following character string value is returned.

(expression)

1.00 MB

```
SELECT FORMAT_UNITS( SUM(chksize), 'P') SIZE,  
       FORMAT_UNITS( SUM(nfree), 'p') FREE FROM syschunks;
```

```
size 117 MB  
free 8.05 MB
```

This query returns the string values size 117 MB and free 8.05 MB.

The following examples show calls to the **FORMAT_UNITS** function with three arguments:

```
EXECUTE FUNCTION FORMAT_UNITS(1024, 'k', 4);
```

The following character string value is returned.

(expression)

1.000 MB

```
SELECT FORMAT_UNITS( SUM(chksize), 'P', 4), SIZE,  
       FORMAT_UNITS( SUM(nfree), 'p', 4) FREE FROM syschunks;
```

```
size 117.2 MB  
free 8.049 MB
```

This query returns the string values `size 117.2 MB` and `free 8.047 MB`. These results differ from the previous example of a query only in their non-default precision, which the last argument to **FORMAT_UNITS** specifies.

IFX_ALLOW_NEWLINE Function

The **IFX_ALLOW_NEWLINE** function sets a newline mode that allows newline characters in quoted strings or disallows newline characters in quoted strings within the current session.

IFX_ALLOW_NEWLINE Function:

```
|—IFX_ALLOW_NEWLINE—(— [ ' t ' ] — ) —————|  
                    | [ ' f ' ] |
```

If you enter 't' as the argument of this function, you enable newline characters in quoted strings in the session. If you enter 'f' as the argument, you disallow newline characters in quoted strings in the session.

You can set the newline mode for all sessions by setting the **ALLOW_NEWLINE** parameter in the **ONCONFIG** file to a value of 0 (newline characters not allowed) or to a value of 1 (newline characters allowed). If you do not set this configuration parameter, the default value is 0. Each time you start a session, the new session inherits the newline mode set in the **ONCONFIG** file. To change the newline mode for the session, execute the **IFX_ALLOW_NEWLINE** function. Once you have set the newline mode for a session, the mode remains in effect until the end of the session or until you execute the **IFX_ALLOW_NEWLINE** function again within the session.

In the following example, assume that you did not specify any value for the **ALLOW_NEWLINE** parameter in the **ONCONFIG** file, so, by default, newline characters are not allowed in quoted strings in any session. After you start a new session, you can enable newline characters in quoted strings in that session by executing the **IFX_ALLOW_NEWLINE** function:

```
EXECUTE PROCEDURE IFX_ALLOW_NEWLINE('t');
```

In **ESQL/C**, the newline mode that is set by the **ALLOW_NEWLINE** parameter in the **ONCONFIG** file or by the execution of the **IFX_ALLOW_NEWLINE** function in a session applies only to quoted-string literals in SQL statements. The newline mode does not apply to quoted strings contained in host variables in SQL

statements. Host variables can contain newline characters within string data regardless of the newline mode currently in effect.

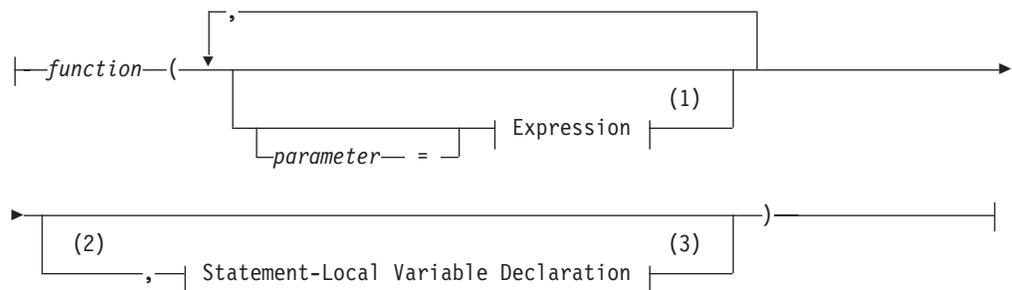
For example, you can use a host variable to insert data that contains newline characters into a column even if the `ALLOW_NEWLINE` parameter in the `ONCONFIG` file is set to 0.

For further information on how the `IFX_ALLOW_NEWLINE` function affects quoted strings, see “Quoted String” on page 4-189. For further information on the `ALLOW_NEWLINE` parameter in the `ONCONFIG` file, see the *IBM Informix Administrator’s Reference*.

User-Defined Functions

A user-defined function (UDF) is a function that you write in SPL or in a language external to the database, such as C or Java. A UDF as an expression has the following syntax:

User-Defined Functions:



Notes:

- 1 See “Expression” on page 4-40
- 2 Dynamic Server only
- 3 See “Statement-Local Variable Declaration (IDS)” on page 4-162

Element	Description	Restrictions	Syntax
<i>function</i>	Name of the function	Function must exist	“Database Object Name ” on page 5-16
<i>parameter</i>	Name of an argument that was declared in a CREATE FUNCTION statement	If you use the <i>parameter</i> = option for any argument in the called function, you must use it for all arguments	“Identifier ” on page 5-21

You can call user-defined functions within SQL statements. Unlike built-in functions, user-defined functions can only be invoked by the creator of the function, by the DBA, and by users who have been granted the Execute privilege on the function. For more information, see “GRANT ” on page 2-395.

The following examples show some user-defined function expressions. The first example omits the *parameter* option when it lists the function argument:

```
read_address('Miller')
```


Notes:

- 1 See “Built-In Data Types” on page 4-21
- 2 See “Complex Data Type (IDS)” on page 4-34

Element	Description	Restrictions	Syntax
<i>distinct_data_type</i>	Name of a distinct data type	The distinct data type must already exist in the database	“Identifier ” on page 5-21
<i>opaque_data_type</i>	Name of an opaque data type	The opaque data type must already exist in the database	“Identifier ” on page 5-21
<i>slv_name</i>	Name of a statement local variable you are defining	The <i>slv_name</i> is valid only for the life of the statement, and must be unique within the statement	“Identifier ” on page 5-21

You can declare an SLV in a call to a user-defined function if both of the following conditions are true:

- The UDF has one or more OUT or INOUT parameters
- The SLV is declared when the UDF is invoked in the WHERE clause of a query.

The SLV declaration in the WHERE clause assigns the value of an OUT or INOUT parameter to the SLV, with the sharp (#) symbol between the identifier of the SLV and its declared data type. The UDF can be written in the SPL, C, or Java language. For example, if you register a function with the following CREATE FUNCTION statement, you can assign the value of its *y* parameter, which is an OUT parameter, to an SLV in a WHERE clause:

```
CREATE FUNCTION find_location(a FLOAT, b FLOAT, OUT y INTEGER)
  RETURNING VARCHAR(20)
  EXTERNAL NAME "/usr/lib/local/find.so"
LANGUAGE C;
```

In this example, **find_location()** accepts two FLOAT values that represent a latitude and a longitude and return the name of the nearest city with an extra value of type INTEGER that represents the population rank of the city.

You can now call **find_location()** in a WHERE clause:

```
SELECT zip_code_t FROM address
  WHERE address.city = find_location(32.1, 35.7, rank # INT)
  AND rank < 101;
```

The function expression passes two FLOAT values to **find_location()** and declares an SLV named **rank** of type INT. In this case, **find_location()** will return the name of the city nearest latitude 32.1 and longitude 35.7 (which might be a heavily populated area) whose population rank is between 1 and 100. The statement then returns the zip code that corresponds to that city.

The SLV can be declared only in a call to the UDF in the WHERE clause of the SELECT statement. The scope of reference of the SLV includes other parts of the same SELECT statement. The following SELECT statement, however, is *invalid* because the SLV declaration is in the projection list of the Projection clause, rather than in the WHERE clause:

```
-- invalid SELECT statement
SELECT title, contains(body, 'dog and cat', rank # INT), rank
  FROM documents;
```

The data type that you specify when you declare the SLV must be the same data type as the corresponding OUT or INOUT parameter in the CREATE FUNCTION statement. If you use different but compatible data types, such as INTEGER and FLOAT, the database server automatically performs the cast between the data types.

SLVs share the name space with UDR variables and the column names of the table involved in the SQL statement. Therefore, the database uses the following descending order of precedence to resolve name conflicts among the following objects:

- UDR variables
- Column names
- SLVs

After the call to the UDF assigns the value of an OUT or INOUT parameter to the SLV, you can reference the SLV in other parts of the same query. For more information, see “Statement-Local Variable Expressions (IDS).”

Statement-Local Variable Expressions (IDS)

The Statement-Local Variable Expression specifies a statement-local variable (SLV) that you can use elsewhere in the same SELECT statement.

Statement-Local Variable Expressions:

|—*SLV_variable*—|

Element	Description	Restrictions	Syntax
<i>SLV_variable</i>	Statement-local variable (SLV) assigned in a call to a user-defined function in the same query	The <i>SLV_variable</i> exists only while the query is executing. Its name must be unique within the query	“Identifier ” on page 5-21

You define an SLV in the call to a user-defined function in the WHERE clause of the SELECT statement. This user-defined function must be defined with one or more OUT or INOUT parameters. The call to the user-defined function assigns the value of the OUT or INOUT parameters to the SLVs. For more information, see “Statement-Local Variable Declaration (IDS)” on page 4-162.

Once the user-defined function assigns its OUT or INOUT parameters to the SLVs, you can use these values in other parts of the same SELECT statement, subject to the following scope-of-reference rules:

- The SLV is *read-only* throughout the query (or subquery) in which it is defined.
- The scope of an SLV extends from the query in which the SLV is defined down into all nested subqueries.
- In nested queries, the scope of an SLV does not extend upwards.

In other words, if a query contains one or more subqueries, an SLV that is defined in the query is also visible to all the subqueries of that query. But if the SLV is defined in the subquery, it is not visible to the parent query.

- In queries that include the UNION operator, the SLV is only visible in the query in which it is defined.

The SLV is not visible to any other queries specified in the UNION.

- For INSERT, DELETE, and UPDATE statements, an SLV is not visible outside the SELECT portion of the statement.

Within this SELECT portion of a DML statement, all the above scoping rules apply.

Important: A statement-local variable is in scope only for the duration of a single SQL statement.

The following SELECT statement calls the `find_location()` function in a WHERE clause and defines the `rank` SLV. Here `find_location()` accepts two values that represent a latitude and a longitude and return the name of the nearest city with an extra value of type INTEGER that represents the population rank of the city.

```
SELECT zip_code_t FROM address
  WHERE address.city = find_location(32.1, 35.7, rank # INT)
  AND rank < 101;
```

When execution of the `find_location()` function completes successfully, the function has initialized the `rank` SLV. The SELECT then uses this `rank` value in a second WHERE clause condition. In this example, the Statement-Local Variable Expression is the variable `rank` in the second WHERE clause condition:

```
rank < 101
```

The number of OUT and INOUT parameters and SLVs that a UDF can have is not restricted. (Releases of Dynamic Server earlier than Version 9.4 restricted user-defined functions to a single OUT parameter and no INOUT parameters, thereby restricting the number of SLVs to no more than one.)

If the user-defined function that initializes the SLVs is *not* executed in an iteration of the statement, the SLVs each have a value of NULL. Values of SLVs do not persist across iterations of the statement. At the start of each iteration, the database server sets the SLV values to NULL.

The following partial statement calls two user-defined functions with OUT parameters, whose values are referenced with the SLV names `out1` and `out2`:

```
SELECT...
  WHERE func_2(x, out1 # INTEGER) < 100
  AND (out1 = 12 OR out1 = 13)
  AND func_3(a, out2 # FLOAT) = "SAN FRANCISCO"
  AND out2 = 3.1416;
```

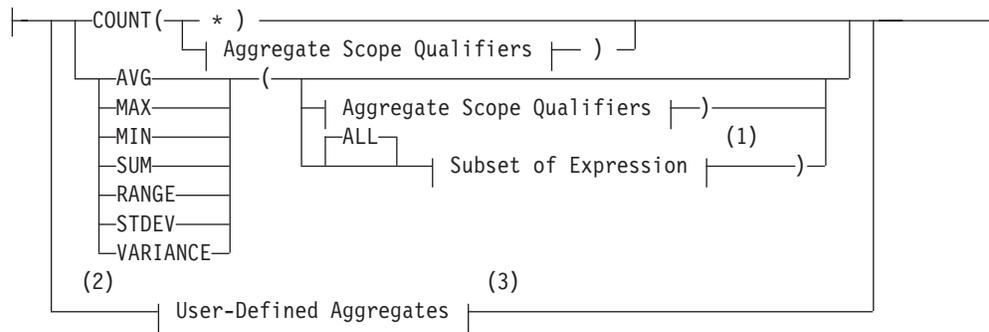
If a function assigns one or more OUT or INOUT parameter values from another database of the local database server to SLVs, the values must be of built-in data types, or DISTINCT data types whose base types are built-in data types (and that you cast explicitly to built-in data types), or must be opaque UDTs that you cast explicitly to built-in data types. All the opaque UDTs, DISTINCT types, type hierarchies, and casts must be defined exactly the same way in all of the participating databases.

For more information on how to write a user-defined function with OUT or INOUT parameters, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*.

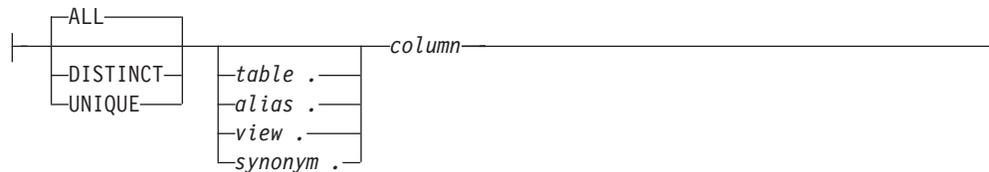
Aggregate Expressions

An aggregate expression uses an aggregate function to summarize selected database data. The built-in aggregate functions have the following syntax.

Aggregate Expressions:



Aggregate Scope Qualifiers:



Notes:

- 1 See "Subset of Expressions Valid in an Aggregate Expression" on page 4-167
- 2 Dynamic Server only
- 3 See "User-Defined Aggregates" on page 4-167

Element	Description	Restrictions	Syntax
<i>column</i>	Column to which aggregate function is applied	See headings for individual keywords on pages that follow	"Identifier" on page 5-21
<i>alias, synonym, table, view</i>	Synonym, table, view, or alias that contains <i>column</i>	<i>Synonym</i> and the <i>table</i> or <i>view</i> to which it points must exist	"Identifier" on page 5-21

You cannot use an aggregate expression in a condition that is part of a WHERE clause unless you use the aggregate expression within a subquery. You cannot apply an aggregate function to a BYTE or TEXT column. For other general restrictions, see "Subset of Expressions Valid in an Aggregate Expression" on page 4-167.

An aggregate function returns one value for a set of queried rows. The following examples show aggregate functions in SELECT statements:

```
SELECT SUM(total_price) FROM items WHERE order_num = 1013;
```

```
SELECT COUNT(*) FROM orders WHERE order_num = 1001;
```

```
SELECT MAX(LENGTH(fname) + LENGTH(lname)) FROM customer;
```

If you use an aggregate function and one or more columns in the projection list of the Projection clause, you must include all the column names that are not used as part of an aggregate or time expression in the GROUP BY clause.

Types of Aggregate Expressions

SQL statements can include *built-in* aggregates and *user-defined* aggregates. The built-in aggregates include all the aggregates shown in the syntax diagram in “Aggregate Expressions” on page 4-165 except for the “User-Defined Aggregates” category. User-defined aggregates are any new aggregates that the user creates with the CREATE AGGREGATE statement.

Built-in Aggregates: Built-in aggregates are aggregate functions that are defined by the database server, such as AVG, SUM, and COUNT. These aggregates work only with built-in data types, such as INTEGER and FLOAT. You can extend these built-in aggregates to work with extended data types. To extend built-in aggregates, you must create UDRs that overload several binary operators.

After you overload the binary operators for a built-in aggregate, you can use that aggregate with an extended data type in an SQL statement. For example, if you have overloaded the **plus** operator for the SUM aggregate to work with a specified row type and assigned this row type to the **complex** column of the **complex_tab** table, you can apply the SUM aggregate to the **complex** column:

```
SELECT SUM(complex) FROM complex_tab;
```

For more information on how to extend built-in aggregates, see *IBM Informix User-Defined Routines and Data Types Developer's Guide*. For information on how to invoke built-in aggregates, see the descriptions of individual built-in aggregates in the following pages.

User-Defined Aggregates: A user-defined aggregate is an aggregate that you define to perform an aggregate computation that the database server does not provide. For example, you can create a user-defined aggregate named SUMSQ that returns the sum of the squared values of a specified column. User-defined aggregates can work with built-in data types or extended data types or both, depending on how you define the support functions for the user-defined aggregate.

To create a user-defined aggregate, use the CREATE AGGREGATE statement. In this statement you name the new aggregate and specify the support functions for the aggregate. Once you create the new aggregate and its support functions, you can use the aggregate in SQL statements. For example, if you created the SUMSQ aggregate and specified that it works with the FLOAT data type, you can apply the SUMSQ aggregate to a FLOAT column named **digits** in the **test** table:

```
SELECT SUMSQ(digits) FROM test;
```

For more information on how to create user-defined aggregates, see “CREATE AGGREGATE ” on page 2-92 and the discussion of user-defined aggregates in *IBM Informix User-Defined Routines and Data Types Developer's Guide*. For information on how to invoke user-defined aggregates, see “ User-Defined Aggregates (IDS)” on page 4-175.

Subset of Expressions Valid in an Aggregate Expression

As indicated in the diagrams for “Aggregate Expressions” on page 4-165 and “ User-Defined Aggregates (IDS)” on page 4-175, not all expressions are available when you use an aggregate expression. The argument of an aggregate function, for example, cannot itself contain an aggregate function. You cannot use aggregate functions in the following contexts:

- In a WHERE clause, unless it is contained in a subquery, or unless the aggregate is on a correlated column from a parent query and the WHERE clause is in a subquery within a HAVING clause
- As an argument to an aggregate function
The following nested aggregate expression is invalid:
`MAX (AVG (order_num))`
- On a BYTE or TEXT column

You cannot use a column that is a collection data type as an argument to the following aggregate functions:

- **AVG**
- **SUM**
- **MIN**
- **MAX**

Expression or *column* arguments to built-in aggregates (except for **COUNT**, **MAX**, **MIN**, and **RANGE**) must return numeric or INTERVAL data types, but **RANGE** also accepts DATE and DATETIME arguments.

For **SUM** and **AVG**, you cannot use the difference between two DATE values directly as the argument to an aggregate, but you can use DATE differences as operands within arithmetic expression arguments. For example:

```
SELECT . . . AVG(ship_date - order_date);
```

returns error -1201, but the following equivalent expression is valid:

```
SELECT . . . AVG((ship_date - order_date)*1);
```

Including or Excluding Duplicates in the Row Set

The **DISTINCT** keyword restricts the argument to unique values from the specified column. The **UNIQUE** and **DISTINCT** keywords are synonyms.

The **ALL** keyword specifies that all values selected from the column or expression, including any duplicate values, are used in the calculation.

AVG Function

The **AVG** function returns the average of all values in the specified column or expression. You can apply the **AVG** function only to number columns. If you use the **DISTINCT** keyword, the average (meaning the *mean*) is calculated from only the distinct values in the specified column or expression. The query in the following example finds the average price of a helmet:

```
SELECT AVG(unit_price) FROM stock WHERE stock_num = 110;
```

NULLs are ignored unless every value in the column is NULL. If every column value is NULL, the **AVG** function returns a NULL for that column.

Overview of COUNT Functions

The **COUNT** function is actually a set of functions that enable you to count column values in different ways, according to arguments after the **COUNT** keyword. Each form of the **COUNT** function is explained in the following subsections. For a comparison of the different forms of the **COUNT** function, see “Comparison of the Different COUNT Functions” on page 4-170.

COUNT(*) Function

The **COUNT (*)** function returns the number of rows that satisfy the WHERE clause of a SELECT statement. The following example finds how many rows in the **stock** table have the value HRO in the **manu_code** column:

```
SELECT COUNT(*) FROM stock WHERE manu_code = 'HRO';
```

The following example queries one of the System Management Interface (SMI) tables to find the number of extents in the **customer** table:

```
SELECT COUNT(*) FROM sysextents WHERE dbs_name = 'stores' AND tablename = customer";
```

You can use **COUNT(*)** as the Projection clause in queries of this general format to obtain information from the SMI tables. For information about **sysextents** and other SMI tables, see the *IBM Informix Administrator's Reference* chapter that describes the **sysmaster** database.

If the SELECT statement does not have a WHERE clause, the **COUNT (*)** function returns the total number of rows in the table. The following example finds how many rows are in the **stock** table:

```
SELECT COUNT(*) FROM stock;
```

If the SELECT statement contains a GROUP BY clause, the **COUNT (*)** function reflects the number of values in each group. The following example is grouped by the first name; the rows are selected if the database server finds more than one occurrence of the same name:

```
SELECT fname, COUNT(*) FROM customer GROUP BY fname
HAVING COUNT(*) > 1;
```

If the value of one or more rows is NULL, the **COUNT (*)** function includes the NULL columns in the count unless the WHERE clause explicitly omits them.

COUNT DISTINCT and COUNT UNIQUE Functions

The **COUNT DISTINCT** function returns the number of unique values in the column or expression, as the following example shows. If the **COUNT DISTINCT** function encounters NULL values, it ignores them:

```
SELECT COUNT (DISTINCT item_num) FROM items;
```

NULLs are ignored unless every value in the specified column is NULL. If every column value is NULL, the **COUNT DISTINCT** function returns zero (0).

The **UNIQUE** keyword has the same meaning as the **DISTINCT** keyword in **COUNT** functions. The **UNIQUE** keyword instructs the database server to return the number of unique non-NULL values in the column or expression. The following example calls the **COUNT UNIQUE** function, but it is equivalent to the preceding example that calls the **COUNT DISTINCT** function:

```
SELECT COUNT (UNIQUE item_num) FROM items;
```

COUNT column Function

The **COUNT column** function returns the total number of non-NULL values in the column or expression, as the following example shows:

```
SELECT COUNT (item_num) FROM items;
```

The ALL keyword can precede the specified column name for clarity, but the query result is the same whether you include the ALL keyword or omit it.

The following example shows how to include the ALL keyword in the COUNT column function:

```
SELECT COUNT (ALL item_num) FROM items;
```

Comparison of the Different COUNT Functions

You can use the different forms of the COUNT function to retrieve different types of information about a table. The following table summarizes the meaning of each form of the COUNT function.

COUNT Function	Description
COUNT (*)	Returns the number of rows that satisfy the query. If you do not specify a WHERE clause, this function returns the total number of rows in the table.
COUNT (DISTINCT) or COUNT (UNIQUE)	Returns the number of unique non-NULL values in the specified column.
COUNT (column) or COUNT (ALL column)	Returns the total number of non-NULL values in the specified column.

Some examples can help to show the differences among the different forms of the COUNT function. Most of the following examples query against the ship_instruct column of the orders table in the demonstration database. For information on the structure of the orders table and the data in the ship_instruct column, see the description of the demonstration database in the *IBM Informix Guide to SQL: Reference*.

Examples of the COUNT(*) Function:

In the following example, the user wants to know the total number of rows in the orders table. So the user calls the COUNT(*) function in a SELECT statement without a WHERE clause:

```
SELECT COUNT(*) AS total_rows FROM orders;
```

The following table shows the result of this query.

total_rows
23

In the following example, the user wants to know how many rows in the orders table have a NULL value in the ship_instruct column. The user calls the COUNT(*) function in a SELECT statement with a WHERE clause, and specifies the IS NULL condition in the WHERE clause:

```
SELECT COUNT (*) AS no_ship_instruct FROM orders  
WHERE ship_instruct IS NULL;
```

The following table shows the result of this query.

no_ship_instruct
2

In the following example, the user wants to know how many rows in the **orders** table have the value **express** in the **ship_instruct** column. So the user calls the **COUNT(*)** function in the projection list and specifies the equals (=) relational operator in the WHERE clause.

```
SELECT COUNT (*) AS ship_express FROM ORDERS
WHERE ship_instruct = 'express';
```

The following table shows the result of this query.

ship_express
6

Examples of the COUNT DISTINCT Function:

In the next example, the user wants to know how many unique non-NULL values are in the **ship_instruct** column of the **orders** table. The user calls the **COUNT DISTINCT** function in the projection list of the SELECT statement:

```
SELECT COUNT(DISTINCT ship_instruct) AS unique_notnulls
FROM orders;
```

The following table shows the result of this query.

unique_notnulls
16

Examples of the COUNT column Function:

In the following example the user wants to know how many non-NULL values are in the **ship_instruct** column of the **orders** table. The user invokes the **COUNT(column)** function in the Projection list of the SELECT statement:

```
SELECT COUNT(ship_instruct) AS total_notnulls FROM orders;
```

The following table shows the result of this query.

total_notnulls
21

A similar query for non-NULL values in the **ship_instruct** column can include the **ALL** keyword in the parentheses that follow the **COUNT** keyword:

```
SELECT COUNT(ALL ship_instruct) AS all_notnulls FROM orders;
```

The following table shows that the query result is the same whether you include or omit the **ALL** keyword (because **ALL** is the default).

all_notnulls
21

MAX Function

The **MAX** function returns the largest value in the specified column or expression. Using the **DISTINCT** keyword does not change the results. The query in the following example finds the most expensive item that is in stock but has not been ordered:

```
SELECT MAX(unit_price) FROM stock
WHERE NOT EXISTS (SELECT * FROM items
WHERE stock.stock_num = items.stock_num AND
stock.manu_code = items.manu_code);
```

NULLs are ignored unless every value in the column is NULL. If every column value is NULL, the **MAX** function returns a NULL for that column.

MIN Function

The **MIN** function returns the lowest value in the column or expression. Using the **DISTINCT** keyword does not change the results. The following example finds the least expensive item in the **stock** table:

```
SELECT MIN(unit_price) FROM stock;
```

NULL values are ignored unless every value in the column is NULL. If every column value is NULL, the **MIN** function returns a NULL for that column.

SUM Function

The **SUM** function returns the sum of all the values in the specified column or expression, as the following example shows. If you use the **DISTINCT** keyword, the sum is for only distinct values in the column or expression:

```
SELECT SUM(total_price) FROM items WHERE order_num = 1013;
```

NULL values are ignored unless every value in the column is NULL. If every column value is NULL, the **SUM** function returns a NULL for that column. You cannot use the **SUM** function with a non-numeric column.

RANGE Function

The **RANGE** function computes the range of returned values. It calculates the difference between the maximum and the minimum values, as follows:

```
range(expr) = max(expr) - min(expr);
```

You can apply the **RANGE** function only to numeric columns. The following query finds the range of ages for a population:

```
SELECT RANGE(age) FROM u_pop;
```

As with other aggregates, the **RANGE** function applies to the rows of a group when the query includes a **GROUP BY** clause, as the next example shows:

```
SELECT RANGE(age) FROM u_pop GROUP BY birth;
```

Because **DATE** values are stored internally as integers, you can use the **RANGE** function on **DATE** columns. With a **DATE** column, the return value is the number of days between the earliest and latest dates in the column.

NULL values are ignored unless every value in the column is NULL. If every column value is NULL, the **RANGE** function returns a NULL for that column.

Important: All computations for the **RANGE** function are performed in 32-digit precision, which should be sufficient for many sets of input data. The computation, however, loses precision or returns incorrect results when all of the input data values have 16 or more digits of precision.

STDEV Function

The **STDEV** function computes the standard deviation of a data set, which is the square root of the **VARIANCE** function. You can apply the **STDEV** function only to numeric columns. The next query finds the standard deviation:

```
SELECT STDEV(age) FROM u_pop WHERE u_pop.age > 0;
```

As with the other aggregates, the **STDEV** function applies to the rows of a group when the query includes a **GROUP BY** clause, as this example shows:

```
SELECT STDEV(age) FROM u_pop GROUP BY birth WHERE STDEV(age) > 0;
```

NULL values are ignored unless every value in the specified column is NULL. If every column value is NULL, **STDEV** returns a NULL for that column.

Important: All computations for the **STDEV** function are performed in 32-digit precision, which should be sufficient for many sets of input data. The computation, however, loses precision or returns incorrect results when all of the input data values have 16 or more digits of precision.

You cannot use this function on columns of type **DATE**.

Within a **SELECT** Statement with **GROUP BY** clause, **STDEV** returns a zero variance for a count of 1. You can omit this special case through appropriate query construction (for example, "HAVING COUNT(*) > 1"). Otherwise, a data set that has only a few cases might block the rest of the query result.

VARIANCE Function

The **VARIANCE** function returns an estimate of the population variance, as the standard deviation squared. **VARIANCE** calculates the following value:

$$(\text{SUM}(X_i^2) - (\text{SUM}(X_i)^2)/N)/(N - 1)$$

In this formula, X_i is each value in the column and N is the total number of non-NULL values in the column (unless all values are NULL, in which case the variance is logically undefined, and the **VARIANCE** function returns NULL).

You can apply the **VARIANCE** function only to numeric columns.

The following query estimates the variance of **age** values for a population:

```
SELECT VARIANCE(age) FROM u_pop WHERE u_pop.age > 0;
```

As with the other aggregates, the **VARIANCE** function applies to the rows of a group when the query includes a **GROUP BY** clause, as in this example:

```
SELECT VARIANCE(age) FROM u_pop GROUP BY birth
WHERE VARIANCE(age) > 0;
```

As previously noted, **VARIANCE** ignores NULL values unless every qualified row is NULL for a specified column. If every value is NULL, then **VARIANCE** returns a NULL result for that column. (This typically indicates missing data, and is not necessarily a good estimate of underlying population variance.)

If *N*, the total number of qualified non-NULL column values, equals one, then the **VARIANCE** function returns zero (another implausible estimate of the true population variance). To omit this special case, you can modify the query. For example, you might include a **HAVING COUNT(*) > 1** clause.

Important: All calculations for the **VARIANCE** function are performed in 32-digit precision, which should be sufficient for many sets of input data. The computation, however, loses precision or returns incorrect results when all of the input data values have 16 or more digits of precision.

Although **DATE** values are stored internally as an integer, you cannot use the **VARIANCE** function on columns of data type **DATE**.

Error Checking in ESQL/C

Aggregate functions always return exactly one row. If no rows are selected, the function returns a **NULL**. You can use the **COUNT (*)** function to determine whether any rows were selected, and you can use an indicator variable to determine whether any selected rows were empty. Fetching a row with a cursor that is associated with an aggregate function always returns one row; hence, 100 for end of data is never returned into the **sqlcode** variable for a first **FETCH** attempt.

You can also use the **GET DIAGNOSTICS** statement for error checking.

Summary of Aggregate Function Behavior

An example can help to summarize the behavior of the aggregate functions. Assume that the **testtable** table has a single **INTEGER** column that is named **num**. The contents of this table are as follows.

num
2
2
2
3
3
4
(NULL)

You can use aggregate functions to obtain information about the **num** column and the **testtable** table. The following query uses the **AVG** function to obtain the average of all the non-NULL values in the **num** column:

```
SELECT AVG(num) AS average_number FROM testtable;
```

The following table shows the result of this query.

average_number
2.66666666666667

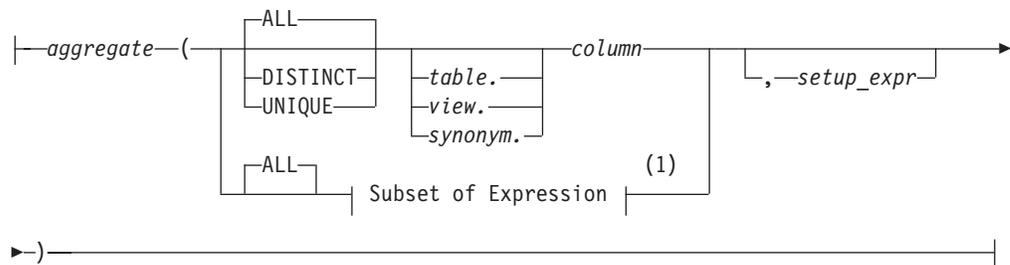
You can use the other aggregate functions in **SELECT** statements that are similar to the preceding example. If you enter a series of **SELECT** statements that have different aggregate functions in the projection list and do not include a **WHERE** clause, you receive the results that the following table shows.

Function	Results	Function	Results
COUNT (*)	7	MAX	4
COUNT (DISTINCT)	3	MAX(DISTINCT)	4
COUNT (ALL num)	6	MIN	2
COUNT (num)	6	MIN(DISTINCT)	2
AVG	2.66666666666667	RANGE	2
AVG (DISTINCT)	3.00000000000000	SUM	16
STDEV	0.74535599249993	SUM(DISTINCT)	9
VARIANCE	0.55555555555556		

User-Defined Aggregates (IDS)

You can create your own aggregate expressions with the CREATE AGGREGATE statement and then invoke these aggregates wherever you can invoke the built-in aggregates. The following diagram shows the syntax for invoking a user-defined aggregate.

User-Defined Aggregates:



Notes:

- 1 See "Subset of Expressions Valid in an Aggregate Expression" on page 4-167

Element	Description	Restrictions	Syntax
<i>aggregate</i>	Name of the user-defined aggregate to invoke	The <i>aggregate</i> and the support functions defined for <i>aggregate</i> must exist	"Identifier " on page 5-21
<i>column</i>	Name of a column within <i>table</i>	Must exist and have a numeric data type	"Quoted String" on page 4-189
<i>setup_expr</i>	Set-up expression that customizes <i>aggregate</i> for a specific invocation	Cannot be a lone host variable. Any columns referenced in <i>setup_expr</i> must be in the GROUP BY clause of the query	"Expression" on page 4-40
<i>synonym, table, view</i>	Synonym, table, or view in which <i>column</i> occurs	The <i>synonym</i> and the table or view to which it points must exist	"Identifier " on page 5-21

Use the DISTINCT or UNIQUE keywords to specify that the user-defined aggregate is to be applied only to unique values in the named column or expression. Use the ALL keyword to specify that the aggregate is to be applied to all values in the named column or expression.

If you omit the DISTINCT, UNIQUE, and ALL keywords, ALL is the default. For further information on the DISTINCT, UNIQUE, and ALL keywords, see "Including or Excluding Duplicates in the Row Set" on page 4-168.

When you specify a setup expression, this value is passed to the **INIT** support function that was defined for the user-defined aggregate in the **CREATE AGGREGATE** statement.

In the following example, you apply the user-defined aggregate named **my_avg** to all values of the **quantity** column in the **items** table:

```
SELECT my_avg(quantity) FROM items
```

In the following example, you apply the user-defined aggregate named **my_sum** to unique values of the **quantity** column in the **items** table. You also supply the value 5 as a setup expression. This value might specify that the initial value of the sum that **my_avg** will compute is 5.

```
SELECT my_sum(DISTINCT quantity, 5) FROM items
```

In the following example, you apply the user-defined aggregate named **my_max** to all values of the **quantity** column in the remote **items** table:

```
SELECT my_max(remote.quantity) FROM rdb@rserv:items remote
```

If the **my_max** aggregate is defined as **EXECUTEANYWHERE**, then the distributed query can be pushed to the remote database server, **rserv**, for execution. If the **my_max** aggregate is not defined as **EXECUTEANYWHERE**, then the distributed query scans the remote **items** table and computes the **my_max** aggregate on the local database server.

You cannot qualify a user-defined aggregate with the name of a remote database server, as the following example shows. In this case, the database server returns an error:

```
SELECT rdb@rserv:my_max(remote.quantity)
FROM rdb@rserv:items remote
```

For further information on user-defined aggregates, see “**CREATE AGGREGATE**” on page 2-92 and the discussion of user-defined aggregates in *IBM Informix User-Defined Routines and Data Types Developer’s Guide*.

Related Information

For a discussion of expressions in the context of the **SELECT** statement, see the *IBM Informix Guide to SQL: Tutorial*.

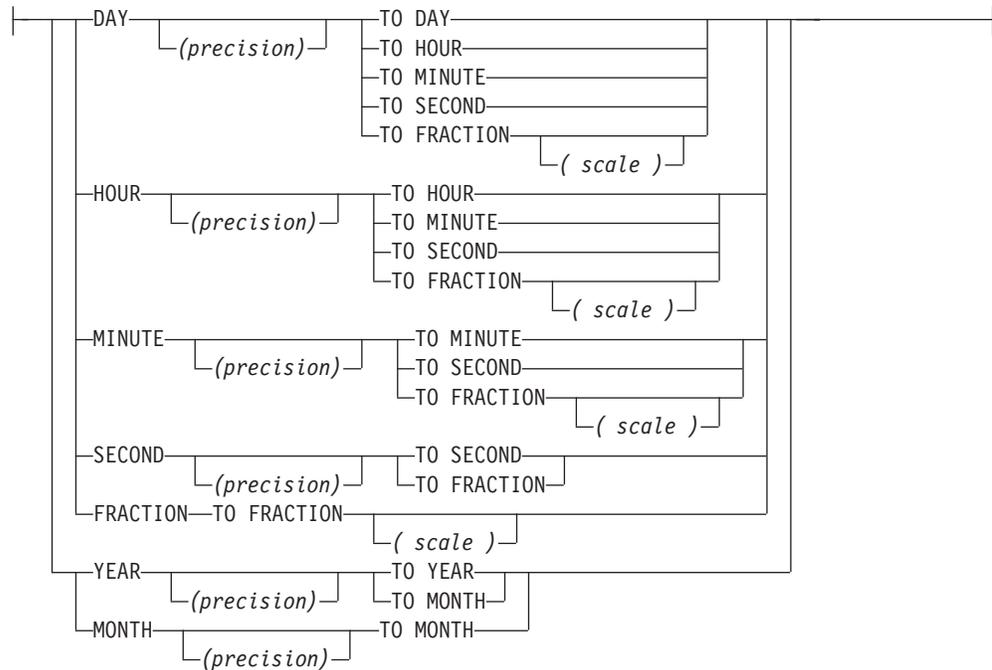
For discussions of column expressions, length functions, and the **TRIM** function, see the *IBM Informix GLS User’s Guide*.

INTERVAL Field Qualifier

The **INTERVAL** field qualifier specifies the precision, in time units, for an **INTERVAL** value. Use the **INTERVAL** Field Qualifier segment whenever you see a reference to an **INTERVAL** field qualifier in a syntax diagram.

Syntax

INTERVAL Field Qualifier:



Element	Description	Restrictions	Syntax
<i>scale</i>	Integer number of digits in FRACTION field. Default is 3.	Must be in the range from 1 to 5	"Literal Number " on page 4-185
<i>precision</i>	Integer number of digits in the largest time unit that the INTERVAL includes. For YEAR, the default is 4. For all other time units except FRACTION, the default is 2.	Must be in the range from 1 to 9	"Literal Number " on page 4-185

Usage

This segment specifies the precision and scale of an INTERVAL data type.

A keyword specifying the *largest* time unit must be the first keyword, and a keyword specifying the *smallest* time unit must follow the TO keyword. These can be the same keyword. This segment resembles the syntax of a " DATETIME Field Qualifier" on page 4-38, but with these exceptions:

- If the largest time unit keyword is YEAR or MONTH, the smallest time unit keyword cannot specify a time unit smaller than MONTH.
- You can specify up to 9-digit *precision* after the first time unit, unless FRACTION is the first time unit (in which case no *precision* is valid after the first FRACTION keyword, but you can specify up to 5 digits of *scale* after the second FRACTION keyword).

Because *year* and *month* are not fixed-length units of time, the database server treats INTERVAL data types that include the YEAR or MONTH keywords in their qualifiers as incompatible with INTERVAL data types whose qualifiers are time units smaller than MONTH. The database server supports no implicit casts between these two categories of INTERVAL data types.

The next two examples show YEAR TO MONTH qualifiers of INTERVAL data types. The first example can hold an interval of up to 999 years and 11 months,

because it gives 3 as the precision of the YEAR field. The second example uses the default precision on the YEAR field, so it can hold an interval of up to 9,999 years and 11 months.

YEAR (3) TO MONTH

YEAR TO MONTH

When you want a value to specify only one kind of time unit, the first and last qualifiers are the same. For example, an interval of whole years is qualified as YEAR TO YEAR or YEAR (5) TO YEAR, for an interval of up to 99,999 years.

The following examples show several forms of INTERVAL field qualifiers:

YEAR(5) TO MONTH

DAY (5) TO FRACTION(2)

DAY TO DAY

FRACTION TO FRACTION (4)

Related Information

For information about how to specify INTERVAL field qualifiers and how to use INTERVAL data in arithmetic and relational operations, see the discussion of the INTERVAL data type in the *IBM Informix Guide to SQL: Reference*.

Literal Collection

Use the Literal Collection segment to specify values for a collection data type. For the syntax of expressions that return values of individual elements within a collection, see “Collection Constructors” on page 4-83.

Syntax

Literal Collection:



Literal Data:



Notes:

- 1 See “Element Literal Value” on page 4-179
- 2 See “Nested Quotation Marks” on page 4-180

Usage

You can specify literal collection values for SET, MULTISET, or LIST data types.

To specify a single literal-collection value, specify the collection type and the literal values. The following SQL statement inserts four integer values into a column called `set_col` that was declared as `SET(INT NOT NULL)`:

```
INSERT INTO table1 (set_col) VALUES (SET{6, 9, 9, 4});
```

Specify an empty collection with an empty pair of braces (`{ }`) symbols. This example inserts an empty list into a column `list_col` that was declared as `LIST(INT NOT NULL)`:

```
INSERT INTO table2 (list_col) VALUES ('LIST{}');
```

A pair of single (`'`) or double (`"`) quotation marks can delimit the collection. However, double quotation marks are not valid in databases where delimited identifiers are enabled, except to delimit SQL identifiers.

If you are passing a literal collection as an argument to an SPL routine, make sure that there is a blank space between the parentheses that surround the arguments and quotation marks that indicate the beginning and end of the literal collection.

Element Literal Value

The diagram for “Literal Collection” on page 4-178 refers to this section. Elements of a collection can be literal values for the following data types.

For a Collection of Type	Literal Value Syntax
BOOLEAN	t or f, representing TRUE or FALSE as a quoted string
CHAR, VARCHAR, NCHAR, NVARCHAR, CHARACTER VARYING, LVARCHAR, DATE	“Quoted String” on page 4-189
DATETIME	“Literal DATETIME ” on page 4-181
DECIMAL, MONEY, FLOAT, INTEGER, INT8, SMALLFLOAT, SMALLINT	“Literal Number ” on page 4-185
INTERVAL	“Literal INTERVAL” on page 4-183
Opaque data types	“Quoted String” on page 4-189. The string literal must be recognized by the input support function for the associated opaque type.
Row Type	“Literal Row” on page 4-186. When the collection element type is a named ROW type, you do not need to cast the inserted values to the named ROW type.

Important: You cannot specify the simple-large-object data types (BYTE and TEXT) as the element type for a collection.

Quoted strings must be specified with a different type of quotation mark than the quotation marks that encompass the collection, so that the database server can parse the quoted strings. Thus, if you use double (`"`) quotation marks to specify the collection, use single (`'`) quotation marks to specify individual, quoted-string elements. (In databases where delimited identifiers are enabled, however, double quotation marks are not valid, except to delimit SQL identifiers.)

Nested Quotation Marks

The diagram for “Literal Collection” on page 4-178 refers to this section.

A *nested collection* is a collection that is the element type for another collection.

Whenever you nest collection literals, use nested quotation marks. In these cases, you must follow the rule for nesting quotation marks. Otherwise, the database server cannot correctly parse the strings.

The general rule is that you must double the number of quotation marks for each new level of nesting. For example, if you use double (") quotation marks for the first level, you must use two double quotation marks for the second level, four double quotation marks for the third level, eight for the fourth level, sixteen for the fifth level, and so on.

Likewise, if you use single (') quotation marks for the first level, you must use two single quotation marks for the second level and four single quotation marks for the third level. There is no limit to the number of levels you can nest, as long as you follow this rule.

The following examples illustrate the case for two levels of nested collection literals, using double (") quotation marks. Here table **tab5** is a single-column table whose only column, **set_col**, is a nested collection type.

The following statement creates the **tab5** table:

```
CREATE TABLE tab5 (set_col SET(SET(INT NOT NULL) NOT NULL));
```

The following statement inserts values into the table **tab5**:

```
INSERT INTO tab5 VALUES ( "SET{"SET{34, 56, 23, 33}"}) );
```

For each literal value, the opening quotation mark and the closing quotation mark must match. Thus, if you open a literal with two double quotation marks, you must close that literal with two double quotation marks ("a literal value").

To specify nested quotation marks within an SQL statement in an Informix ESQL/C program, use the C escape character for every double quotation mark inside a string that is delimited by single quotation marks. Otherwise, the Informix ESQL/C preprocessor cannot correctly interpret the literal collection value. For example, the preceding INSERT statement on the **tab5** table would appear in an Informix ESQL/C program as follows:

```
EXEC SQL insert into tab5
  values ('set{\set{34, 56, 23, 33}\}');
```

For more information, see the chapter on complex data types in the *IBM Informix ESQL/C Programmer's Manual*.

If the collection is a nested collection, you must include the collection-constructor syntax for each level of collection type. Suppose you define the following column:

```
nest_col SET(MULTISET (INT NOT NULL) NOT NULL);
```

The following statement inserts three elements into the **nest_col** column:

```
INSERT INTO tabx (nest_col)
  VALUES ("SET{'MULTISET{1, 2, 3}'});
```


Element	Description	Restrictions	Syntax
<i>ffff</i>	Fraction of a second, in digits	$0 \leq ffff \leq 9999$	"Literal Number " on page 4-185
<i>hh</i>	Hour of day, expressed in digits	$0 \leq hh \leq 23$	"Literal Number " on page 4-185
<i>mi</i>	Minute of hour, expressed in digits	$0 \leq mi \leq 59$	"Literal Number " on page 4-185
<i>mo</i>	Month of year, expressed in digits	$1 \leq mo \leq 11$	"Literal Number " on page 4-185
<i>space</i>	Blank space (ASCII 32)	Exactly 1 blank character	Literal blank space
<i>ss</i>	Second of minute, in digits	$0 \leq ss \leq 59$	"Literal Number " on page 4-185
<i>yyyy</i>	Year, expressed in digits	No more than 4 digits	"Literal Number " on page 4-185

Usage

You must specify both a numeric date and a DATETIME field qualifier for this date in the Literal DATETIME segment. The DATETIME field qualifier must correspond to the numeric date you specify. For example, if you specify a numeric date that includes a year as the largest unit and a minute as the smallest unit, you must also specify YEAR TO MINUTE as the DATETIME field qualifier.

If you specify two digits for the year, the database server uses the setting of the **DBCENTURY** environment variable to expand the abbreviated year value to four digits. If the **DBCENTURY** is not set, the first two digits of the current year are used to expand the abbreviated year value.

The following examples show literal DATETIME values:

```
DATETIME (07-3-6) YEAR TO DAY
```

```
DATETIME (09:55:30.825) HOUR TO FRACTION
```

```
DATETIME (07-5) YEAR TO MONTH
```

The following example shows a literal DATETIME value used with the EXTEND function:

```
EXTEND (DATETIME (2007-8-1) YEAR TO DAY, YEAR TO MINUTE)
- INTERVAL (720) MINUTE (3) TO MINUTE
```

Casting Numeric Date and Time Strings to DATE Data Types

The database server provides a built-in cast to convert DATETIME values to DATE values, as in the following SPL program fragment:

```
DEFINE my_date DATE;
DEFINE my_dt DATETIME YEAR TO SECOND;
. . .
LET my_date = CURRENT;
```

Here the DATETIME value that CURRENT returns is implicitly cast to DATE. You can also cast DATETIME to DATE explicitly:

```
LET my_date = CURRENT::DATE;
```

Both of these LET statements assign the *year*, *month*, and *day* information from the DATETIME value to the local SPL variable **my_date** of type DATE.

Similarly, you can cast explicitly a string that has the format of the Numeric Date and Time segment, as defined in the "Literal DATETIME " on page 4-181 syntax diagram, to a DATETIME data type, as in the following example:

```
LET my_dt =  
    ('2008-02-22 05:58:44.000')::DATETIME YEAR TO SECOND;
```

There is neither an implicit nor an explicit built-in cast, however, for directly converting a character string that has the Numeric Date and Time format to a DATE value. Both of the following statements, for example, fail with error -1218:

```
LET my_date = ('2008-02-22 05:58:44.000');  
LET my_date = ('2008-02-22 05:58:44.000')::DATE;
```

To convert a character string that specifies a valid numeric date and time value to a DATE data type, you must first cast the string to DATETIME, and then cast the resulting DATETIME value to DATE, as in this example:

```
LET my_date =  
    ('2008-02-22 05:58:44.000')::DATETIME YEAR TO SECOND::DATE;
```

A direct string-to-DATE cast can succeed only if the string specifies a valid DATE value.

Related Information

For discussions of the DATETIME data type and the **DBCENTURY** environment variable, see the *IBM Informix Guide to SQL: Reference*.

For a discussion of how to use the **GL_DATETIME** environment variable to customize the display format of DATETIME values in non-default locales, see the *IBM Informix GLS User's Guide*.

Literal INTERVAL

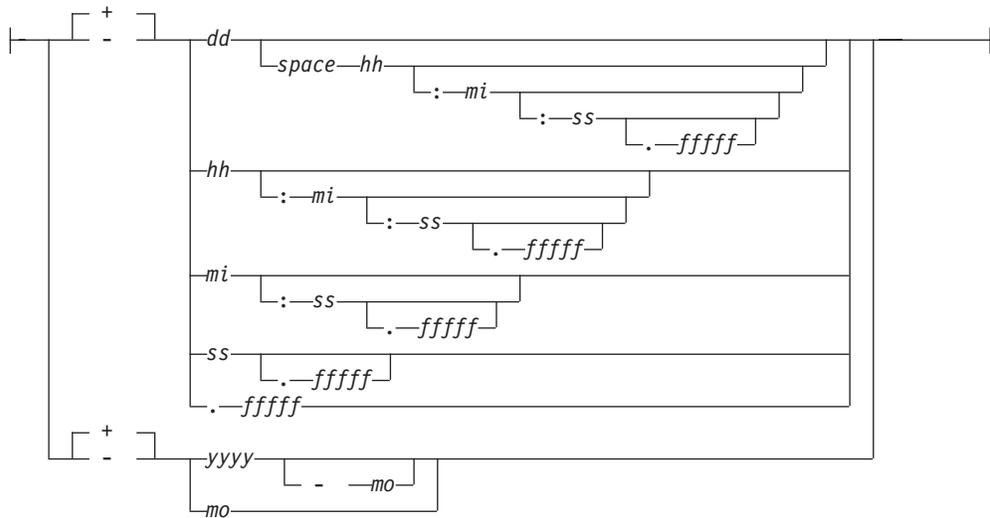
The Literal INTERVAL segment specifies a literal INTERVAL value. Use this whenever you see a reference to a literal INTERVAL in a syntax diagram.

Syntax

Literal INTERVAL:

|—INTERVAL—(—| Numeric Time Span |—)| INTERVAL Field Qualifier |—————|
(1)

Numeric Time Span:



Notes:

- 1 See "INTERVAL Field Qualifier " on page 4-176

Element	Description	Restrictions	Syntax
<i>dd</i>	Number of days	$-10^{**}10 < dd < 10^{**}10$	"Literal Number " on page 4-185
<i>fffff</i>	Fractions of a second	$0 \leq fffff \leq 9999$	"Literal Number " on page 4-185
<i>hh</i>	Number of hours	If not first, $0 \leq hh \leq 23$	"Literal Number " on page 4-185
<i>mi</i>	Number of minutes	If not first, $0 \leq mi \leq 59$	"Literal Number " on page 4-185
<i>mo</i>	Number of months	If not first, $0 \leq mo \leq 11$	"Literal Number " on page 4-185
<i>space</i>	Blank space (ASCII 32)	Exactly 1 blank character is required	Literal blank space
<i>ss</i>	Number of seconds	If not first, $0 \leq ss \leq 59$	"Literal Number " on page 4-185
<i>yyyy</i>	Number of years	$-10^{**}10 < yyyy < 10^{**}10$	"Literal Number " on page 4-185

Usage

Unlike DATETIME literals, INTERVAL literals can include the unary plus (+) or unary minus (-) sign. If you specify no sign, the default is plus.

The precision of the first time unit can be specified by the INTERVAL qualifier. Except for FRACTION, which can have no more than 5 digits of precision, the first time unit can have up to 9 digits of precision, if you specified a nondefault precision in the declaration of the INTERVAL column or variable.

The following examples show literal INTERVAL values:

```
INTERVAL (3-6) YEAR TO MONTH
INTERVAL (09:55:30.825) HOUR TO FRACTION
INTERVAL (40 5) DAY TO HOUR
INTERVAL (299995.2567) SECOND(6) TO FRACTION(4)
```

Only the last of these examples has nondefault precision. For the syntax of declaring the precision of INTERVAL data types and the default values for each time unit, refer to “INTERVAL Field Qualifier ” on page 4-176.

Related Information

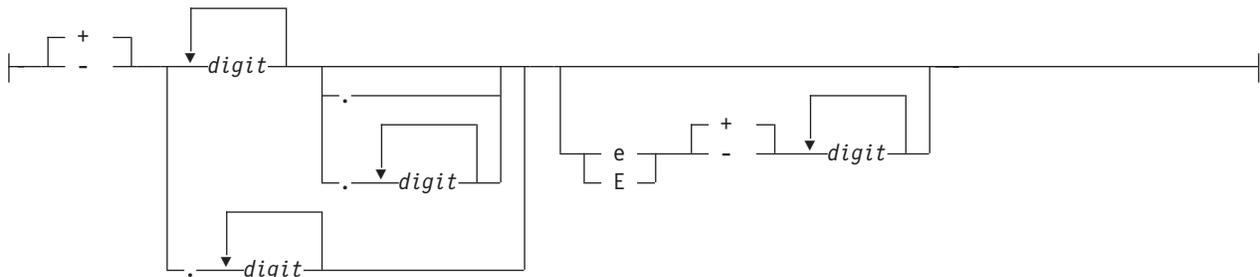
For information on how to use INTERVAL data in arithmetic and relational operations, see the discussion of the INTERVAL data type in the *IBM Informix Guide to SQL: Reference*.

Literal Number

A *literal number* is the base-10 representation of a real number as an integer, as a fixed-point decimal number, or in exponential notation. Use the Literal Number segment whenever you see a reference to a literal number in a syntax diagram.

Syntax

Literal Number:



Element	Description	Restrictions	Syntax
<i>digit</i>	Integer in range 0 through 9	Must be an ASCII digit	Literal entered from the keyboard.

Usage

You cannot include comma (,) or blank (ASCII 32) character. The unary plus (+) or minus (-) sign can precede a literal number, mantissa, or exponent.

You cannot include non-ASCII digits in literal numbers, such as the Hindi numbers that some nondefault locales support.

Integer Literals

An integer has no fractional part and cannot include a decimal point. Built-in data types of SQL that can be exactly represented as literal integers include BIGINT, BIGSERIAL, DECIMAL(*p*, 0), INT, INT8, SERIAL, SERIAL8, and SMALLINT.

If you use the representation of a number in a base other than 10 (such as a binary, octal, or hexadecimal) in any context where a literal integer is valid, the database server will attempt to interpret the value as a base-10 literal integer. For most data values, the result will be incorrect.

The following examples show some valid literal integers:

```
10          -27          +25567
```

Thousands separators (such as comma symbols) are not valid in literal integers, nor in any other literal number.

Fixed-Point Decimal Literals

Fixed-point decimal literals can exactly represent `DECIMAL(p,s)` and `MONEY` values. These can include a decimal point:

-123.456 00123456 +123456.0

The digits to the right of the decimal point in these examples are the fractional portions of the numbers.

Floating-Point Decimal Literals

Floating-point literals can exactly represent `FLOAT`, `SMALLFLOAT`, and `DECIMAL(p)` values, using a decimal point or exponential notation, or both. They can approximately represent real numbers in exponential notation. The next examples show floating point numbers:

-123.45E6 1.23456E2 123456.0E-3

The E in the previous examples is the symbol for exponential notation. The digit that follows E is the value of the exponent. For example, the number 3E5 (or 3E+5) means 3 multiplied by 10 to the fifth power, and the number 3E-5 means 3 multiplied by the reciprocal of 10 to the fifth power.

Literal Numbers and the MONEY Data Type

When you use a literal number as a `MONEY` value, do not include a currency symbol or include commas. The `DBMONEY` environment variable or the locale file can format how `MONEY` values are displayed in output.

Related Information

For discussions of numeric data types, such as `DECIMAL`, `FLOAT`, `INTEGER`, and `MONEY`, see the *IBM Informix Guide to SQL: Reference*.

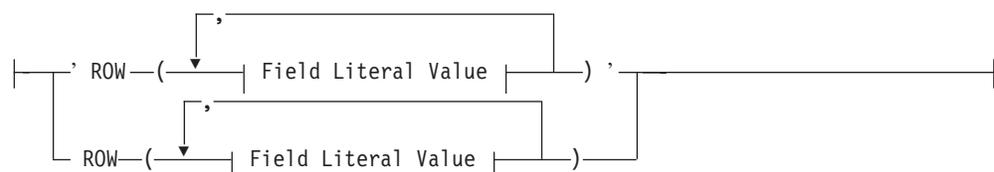
Literal Row

The Literal Row segment specifies the syntax for literal values of named and unnamed `ROW` data types.

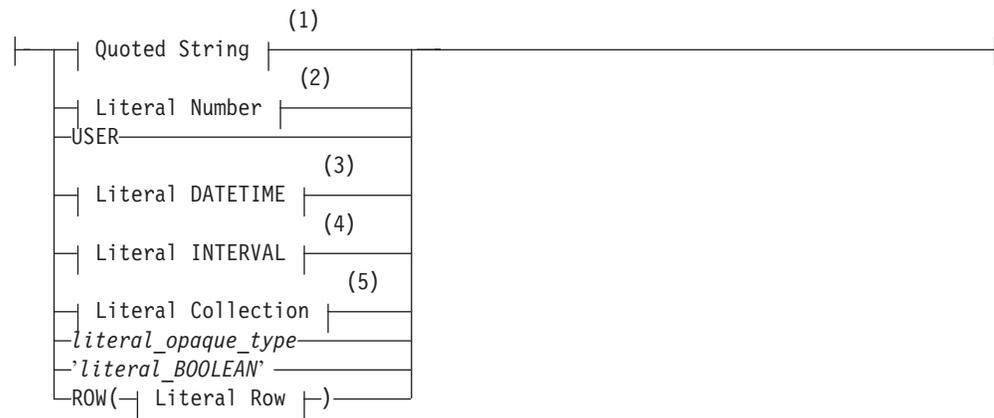
Only Dynamic Server supports this syntax. For expressions that evaluate to field values within a `ROW` data type, see "ROW Constructors" on page 4-82.

Syntax

Literal Row:



Field Literal Value:



Notes:

- 1 See “Quoted String” on page 4-189
- 2 See “Literal Number ” on page 4-185
- 3 See “Literal DATETIME ” on page 4-181
- 4 See “Literal INTERVAL” on page 4-183
- 5 See “Literal Collection” on page 4-178

Element	Description	Restrictions	Syntax
<i>literal_opaque_type</i>	Literal representation for an opaque data type	Must be a literal that is recognized by the input support function for the associated opaque data type	Defined by the developer of the opaque data type
<i>literal_BOOLEAN</i>	Literal representation of a BOOLEAN value	Must be either 't' (= TRUE) or 'f' (= FALSE) specified as a quoted string	“Quoted String” on page 4-189

Usage

You can specify literal values for named ROW and unnamed ROW data types. A ROW constructor introduces a literal ROW value, which can optionally be enclosed between quotation marks.

The format of the value for each field of the ROW type must be compatible with the data type of the corresponding field.

Important: You cannot specify simple-large-object data types (BYTE or TEXT) as the field type for a row.

Fields of a row can be literal values for the data types in the following table.

For a Field of Type	Literal Value Syntax
BOOLEAN	t or f, representing TRUE or FALSE
CHAR, VARCHAR, LVARCHAR, NCHAR, NVARCHAR, CHARACTER VARYING, DATE	“Quoted String” on page 4-189
DATETIME	“Literal DATETIME ” on page 4-181

For a Field of Type	Literal Value Syntax
DECIMAL, MONEY, FLOAT, INTEGER, INT8, SMALLFLOAT, SMALLINT	"Literal Number " on page 4-185
INTERVAL	"Literal INTERVAL" on page 4-183
Opaque data types	"Quoted String" on page 4-189 The string must be a literal that is recognized by the input support function for the associated opaque type.
Collection type (SET, MULTISSET, LIST)	"Literal Collection" on page 4-178 For information on literal collection values as variable or column values, see "Nested Quotation Marks" on page 4-180. For information on literal collection values for a ROW type, see " Literals for Nested Rows" on page 4-189.
Another ROW type (named or unnamed)	For information on ROW type values, see " Literals for Nested Rows" on page 4-189.

Literals of an Unnamed Row Type

To specify a literal value for an unnamed ROW type, introduce the literal row with the ROW constructor; you must enclose the values between parentheses. For example, suppose that you define the **rectangles** table as follows:

```
CREATE TABLE rectangles
(
  area FLOAT,
  rect ROW(x INTEGER, y INTEGER, length FLOAT, width FLOAT),
)
```

The following INSERT statement inserts values into the **rect** column of the **rectangles** table:

```
INSERT INTO rectangles (rect)
VALUES ("ROW(7, 3, 6.0, 2.0)")
```

Literals of a Named Row Type

To specify a literal value for a named ROW type, introduce the literal row with the ROW type constructor and enclose the literal values for each field in parentheses. In addition, you can cast the row literal to the appropriate named ROW type to ensure that the row value is generated as a named ROW type. The following statements create the named ROW type **address_t** and the **employee** table:

```
CREATE ROW TYPE address_t
(
  street CHAR(20),
  city CHAR(15),
  state CHAR(2),
  zipcode CHAR(9)
);

CREATE TABLE employee
(
  name CHAR(30),
  address address_t
);
```

The following INSERT statement inserts values into the **address** column of the **employee** table:

```
INSERT INTO employee (address)
VALUES (
"ROW('103 Baker St', 'Tracy','CA', 94060)::address_t)
```

Literals for Nested Rows

If the literal value is for a nested row, specify the ROW type constructor for each row level. If you include quotation marks as delimiters, they should enclose the outermost row. For example, suppose that you create the **emp_tab** table:

```
CREATE TABLE emp_tab
(
  emp_name CHAR(10),
  emp_info ROW( stats ROW(x INT, y INT, z FLOAT))
);
```

The following INSERT statement adds a row to the **emp_tab** table:

```
INSERT INTO emp_tab VALUES ('joe boyd', "ROW(ROW(8,1,12.0))" );
```

Similarly, if the row-string literal contains a nested collection, only the outermost literal row can be enclosed between quotation marks. Do not put quotation marks around an inner, nested collection type.

Related Information

Related statements: CREATE ROW TYPE, INSERT, UPDATE, and SELECT

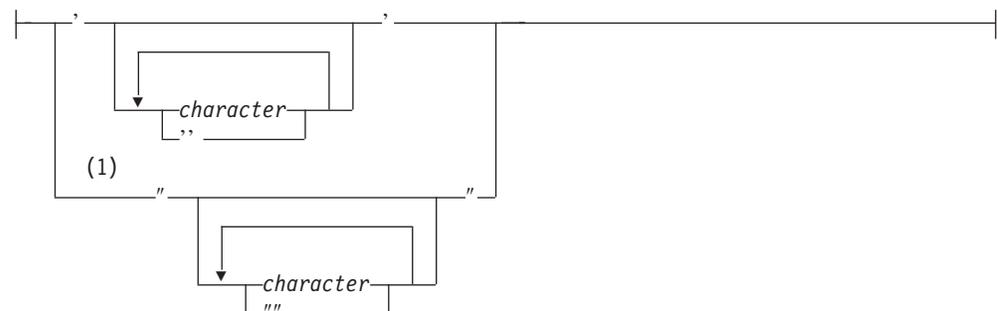
For information on ROW constructors, see the Expression segment. See also the Collection Literal segment.

Quoted String

A quoted string is a string literal between quotation marks. Use this segment whenever you see a reference to a quoted string in a syntax diagram.

Syntax

Quoted String:



Notes:

- 1 Informix extension

Element	Description	Restrictions	Syntax
<i>character</i>	Code set element within quoted string	Cannot enclose between double quotation marks if the DELIMIDENT environment variable is set	Literal value from the keyboard

Usage

Use quoted strings to specify string literals in data-manipulation statements and other SQL statements. For example, you can use a quoted string in an INSERT statement to insert a value into a column of a character data type.

Restrictions on Specifying Characters in Quoted Strings

You must observe the following restrictions on *character* in quoted strings:

- If you are using the ASCII code set, you can specify any printable ASCII character, including a single (') quotation mark or double (") quotation mark. For restrictions that apply to using quotation marks in quoted strings, see “Using Quotation Marks in Strings” on page 4-192.
- In some locales, you can specify non-ASCII characters, including multibyte characters, that the locale supports. See the discussion of quoted strings in the *IBM Informix GLS User's Guide*.
- If you enable newline characters for quoted strings, you can embed newline characters in quoted strings. For further information, see “Newline Characters in Quoted Strings” on page 4-191.
- You can enter DATETIME and INTERVAL data values as quoted strings. For the restrictions that apply to entering DATETIME and INTERVAL data in quoted-string format, see “DATETIME and INTERVAL Values as Strings” on page 4-192.
- Quoted strings that are used with the LIKE or MATCHES keyword in a search condition can include wildcard characters that have a special meaning in the search condition. For further information, see “LIKE and MATCHES in a Condition” on page 4-192.
- When you insert a value that is a quoted string, you must observe a number of restrictions. For further information, see “Inserting Values as Quoted Strings” on page 4-192.

The DELIMIDENT Environment Variable

If the **DELIMIDENT** environment variable is set on the database server, you cannot use double quotation marks (") to delimit literal strings. If **DELIMIDENT** is set, the database server interprets strings enclosed in double quotation marks as SQL identifiers, not as literal strings. If **DELIMIDENT** is not set, a string between double quotation marks is interpreted as a literal string, not an identifier. For further information, see “Using Quotation Marks in Strings” on page 4-192, and the description of **DELIMIDENT** in *IBM Informix Guide to SQL: Reference*.

DELIMIDENT is also supported on client systems, where it can be set to *y*, to *n*, or to no setting.

- *y* specifies that client applications must use single quotation mark (') symbols to delimit literal strings, and must use double quotation mark (") symbols only around delimited SQL identifiers. Delimited identifiers can support a larger character set than is valid for undelimited identifiers. Letters within delimited strings or delimited identifiers are case-sensitive.

- `n` specifies that client applications can use double quotation mark (`"`) or single quotation mark (`'`) symbols to delimit character strings, but not to delimit SQL identifiers. If the database server encounters a string delimited by double or single quotation mark symbols in a context where an SQL identifier is required, it issues an error. An owner name, however, that qualifies an SQL identifier can be delimited by single quotation mark (`'`) symbols. You must use a pair of the same quotation mark symbols to delimit a character string.
- Specifying **DELIMIDENT** with no value on the client system requires client applications to use the **DELIMIDENT** setting that is the default for their application programming interface (API).

Client APIs of Dynamic Server use the following default **DELIMIDENT** settings:

- For OLE DB and .NET, the default **DELIMIDENT** setting is `y`
- For ESQL/C, JDBC, and ODBC, the default **DELIMIDENT** setting is `n`
- APIs that have ESQL/C as an underlying layer, such as Informix 4GL, the DataBlade API (LIBDMI), and the C++ API, behave as ESQL/C, and use `'n'` as the default if no value for **DELIMIDENT** is specified on the client system.

Even if **DELIMIDENT** is set, you can use single quotation mark (`'`) symbols to delimit *authorization identifiers* as the owner name component of a database object name, as in the following example:

```
RENAME COLUMN 'owner'.table2.column3 TO column3;
```

The general rule, however, is that when **DELIMIDENT** is set, the SQL parser interprets strings delimited by single quotation marks as string literals, and interprets character strings delimited by double quotation marks (`"`) as SQL identifiers.

Newline Characters in Quoted Strings

By default, the string constant must be written on a single line. That is, you cannot use embedded newline characters in a quoted string. You can, however, override this default behavior in one of two ways:

- To enable newline characters in quoted strings in all sessions, set the **ALLOW_NEWLINE** parameter to 1 in the **ONCONFIG** file.
- To enable newline characters in quoted strings for the current session, execute the built-in function **IFX_ALLOW_NEWLINE**.

This enables newline characters in quoted strings for the current session:

```
EXECUTE PROCEDURE IFX_ALLOW_NEWLINE('T');
```

If newline characters in quoted strings are not enabled for a session, the following statement is invalid and returns an error:

```
SELECT 'The quick brown fox
      jumped over the old gray fence'
FROM customer
WHERE customer_num = 101;
```

If you enable newline characters in quoted strings for the session, however, the statement in the preceding example is valid and executes successfully.

For more information on the **IFX_ALLOW_NEWLINE** function, see “**IFX_ALLOW_NEWLINE** Function” on page 4-160. For more information on the **ALLOW_NEWLINE** parameter in the **ONCONFIG** file, see your *IBM Informix Administrator's Reference*.

Using Quotation Marks in Strings

The single quotation mark (') has no special significance in string literals delimited by double quotation marks. Conversely, double quotation mark (") has no special significance in strings delimited by single quotation marks. For example, these strings are valid:

```
"Nancy's puppy jumped the fence"  
'Billy told his kitten, "No!" '
```

A string delimited by double quotation marks can include a double quotation mark character by preceding it with another double quotation mark, as the following string shows:

```
"Enter ""y"" to select this row"
```

When the **DELIMITED** environment variable is set, double quotation marks can only delimit SQL identifiers, not strings. For more information on delimited identifiers, see “Delimited Identifiers” on page 5-23.

DATETIME and INTERVAL Values as Strings

You can enter DATETIME and INTERVAL data in the literal forms described in the “Literal DATETIME ” on page 4-181 and “Literal INTERVAL” on page 4-183, or you can enter them as quoted strings.

Valid literals that are entered as character strings are converted automatically into DATETIME or INTERVAL values.

These statements enter INTERVAL and DATETIME values as quoted strings:

```
INSERT INTO cust_calls(call_dtime) VALUES ('2007-5-4 10:12:11');  
INSERT INTO manufact(lead_time) VALUES ('14');
```

The format of the value in the quoted string must exactly match the format specified by the INTERVAL or DATETIME qualifiers of the column. For the first INSERT in the preceding example, the **call_dtime** column must be defined with the qualifiers YEAR TO SECOND for the INSERT statement to be valid.

LIKE and MATCHES in a Condition

Quoted strings with the LIKE or MATCHES keyword in a condition can include wildcard characters. For a complete description of how to use wildcard characters, see “Condition ” on page 4-5.

Inserting Values as Quoted Strings

In the default locale, if you are inserting a value that is a quoted string, you must adhere to the following restrictions:

- Enclose CHAR, VARCHAR, NCHAR, NVARCHAR, DATE, DATETIME, INTERVAL, and (for Dynamic Server) LVARCHAR values in quotation marks.
- Specify DATE values in the *mm/dd/yyyy* format (or in the format that the **DBDATE** or **GL_DATE** environment variable specifies, if set).
- In Extended Parallel Server, you cannot insert strings longer than 256 bytes.
- In Dynamic Server, you cannot insert strings longer than 32 kilobytes.
- Numbers with decimal values must include a decimal separator. Comma (,) is not valid as a decimal separator in the default locale.
- MONEY values cannot include a dollar sign (\$) or commas.

- You can enter NULL in a column only if it accepts null values.

Numeric Operations on Character Columns

Avoid comparing number literals to character columns. It requires that all of the strings compared be converted to numbers, which takes much longer than comparing two strings. For example, suppose that you wish to find all customers within the 356 telephone exchange code:

```
SELECT lname FROM customer WHERE phone [5,7] = '356';
```

Notice that the operand whose value is 356 is enclosed in quotes. The quotes indicate that the database server must handle the filter as a character string. By contrast, when the operand is not in quotes, the server treats each retrieved value as a number, and must implicitly cast each value retrieved from the table to a numeric data type.

The following example causes implicit data type conversion of the substrings:

```
SELECT lname FROM customer WHERE phone [5,7] = 356;
```

When the phone column has a distribution that an UPDATE STATISTICS MEDIUM or by an UPDATE STATISTICS HIGH statement, the query fails with a -1213 conversion error. To avoid this error, you can take any one of the following actions:

- Be sure to enclose the operand between quotes. (This is equivalent to the rule "Do not apply numeric filters to character columns.")
- Store only characters in the range ASCII 0x30 through 0x39, and decimal point (ASCII 0x2e). This range is also known as "semi-numeric."
- When developing your schema and UPDATE STATISTICS strategy, ensure that character columns contain either semi-numerics, or characters, but not both, if you expect that numeric filters will be useful.

Related Information

For a discussion of the **DELIMITED** environment variable, see the *IBM Informix Guide to SQL: Reference*.

For a discussion of the GLS aspects of quoted strings, see the *IBM Informix GLS User's Guide*.

Relational Operator

A relational operator compares two expressions quantitatively. Use the Relational Operator segment whenever you see a reference to a relational operator in a syntax diagram.

Syntax

Relational Operator:

<	_____
<=	_____
>	_____
=	_____
==	_____
>=	_____
<>	_____
(1)	_____
!=	_____

Notes:

- 1 Informix extension

Usage

The relational operators of SQL have the following meanings.

**Relational Operator
Meaning**

- < Less than
- <= Less than or equal to
- > Greater than
- = or == Equal to
- >= Greater than or equal to
- <> or != Not equal to

Usage

For number expressions, *greater than* means to the right on the real line.

For DATE and DATETIME expressions, *greater than* means later in time.

For INTERVAL expressions, *greater than* means a longer span of time.

For CHAR, VARCHAR, and LVARCHAR expressions, *greater than* means *after* in code-set order. (For NCHAR and NVARCHAR expressions, *greater than* means *after* in the localized collation order, if one exists; otherwise, it means in code-set order.)

Locale-based collation order, if defined for the locale, is used for NCHAR and NVARCHAR expressions. So for NCHAR and NVARCHAR expressions, *greater than* means *after* in the locale-based collation order. For more information on locale-based collation order and the NCHAR and NVARCHAR data types, see the *IBM Informix GLS User's Guide*.

Using Operator Functions in Place of Relational Operators

Each relational operator is bound to a particular operator function, as the table shows. The operator function accepts two values and returns a boolean value of true, false, or unknown.

**Relational Operator
Associated Operator Function**

```

<    lessthan( )
<=  lessthanorequal( )
>    greaterthan( )
>=  greaterthanorequal( )
= or ==
    equal( )
<> or !=
    notequal( )

```

Connecting two expressions with a relational operator is equivalent to invoking the operator function on the expressions. For example, the next two statements both select orders with a shipping charge of \$18.00 or more.

The `>=` operator in the first statement implicitly invokes the `greaterthanorequal()` operator function:

```

SELECT order_num FROM orders
   WHERE ship_charge >= 18.00;

SELECT order_num FROM orders
   WHERE greaterthanorequal(ship_charge, 18.00);

```

The database server provides the operator functions associated with the relational operators for all built-in data types. When you develop a user-defined data type, you must define the operator functions for that type for users to be able to use the relational operator on the type.

If you define `lessthan()`, `greaterthan()`, and the other operator functions for a user-defined type, then you should also define `compare()`. Similarly, if you define `compare()`, then you should also define `lessthan()`, `greaterthan()`, and the other operator functions. All of these functions must be defined in a consistent manner, to avoid the possibility of incorrect query results when UDT values are compared in the WHERE clause of a SELECT.

Collating Order for U.S. English Data

If you are using the default locale (U.S. English), the database server uses the code-set order of the default code set when it compares the character expressions that precede and follow the relational operator.

On UNIX, the default code set is the ISO8859-1 code set, which consists of the following sets of characters:

- The ASCII characters have code points in the range of 0 to 127.
This range contains control characters, punctuation symbols, English-language characters, and numerals.
- The 8-bit characters have code points in the range 128 to 255.
This range includes many non-English-language characters (such as é, â, ö, and ñ) and symbols (such as £, ©, and ¸).

In Windows, the default code set is Microsoft 1252. This code set includes both the ASCII code set and a set of 8-bit characters.

This table lists the ASCII code set. The **Num** columns show ASCII code point numbers, and the **Char** columns display corresponding ASCII characters. In the

default locale, ASCII characters are sorted according to their code-set order. Thus, lowercase letters follow uppercase letters, and both follow digits. In this table, ASCII 32 is the blank character, and the caret symbol (^) stands for the CTRL key. For example, ^X means CONTROL-X.

Num	Char												
0	^@	20	^T	40	(60	<	80	P	100	d	120	x
1	^A	21	^U	41)	61	=	81	Q	101	e	121	y
2	^B	22	^V	42	*	62	>	82	R	102	f	122	z
3	^C	23	^W	43	+	63	?	83	S	103	g	123	{
4	^D	24	^X	44	,	64	@	84	T	104	h	124	
5	^E	25	^Y	45	-	65	A	85	U	105	i	125	}
6	^F	26	^Z	46	.	66	B	86	V	106	j	126	~
7	^G	27	esc	47	/	67	C	87	W	107	k	127	del
8	^H	28	^\	48	0	68	D	88	X	108	l		
9	^I	29	^]	49	1	69	E	89	Y	109	m		
10	^J	30	^^	50	2	70	F	90	Z	110	n		
11	^K	31	^_	51	3	71	G	91	[111	o		
12	^L	32		52	4	72	H	92	\	112	p		
13	^M	33	!	53	5	73	I	93]	113	q		
14	^N	34	"	54	6	74	J	94	^	114	r		
15	^O	35	#	55	7	75	K	95	_	115	s		
16	^P	36	\$	56	8	76	L	96	`	116	t		
17	^Q	37	%	57	9	77	M	97	a	117	u		
18	^R	38	&	58	:	78	N	98	b	118	v		
19	^S	39	'	59	;	79	O	99	c	119	w		

Support for ASCII Characters in Nondefault Code Sets (GLS)

Most code sets for nondefault locales (called *nondefault code sets*) support the ASCII characters. In a nondefault locale, the database server uses ASCII code-set order for ASCII data in CHAR and VARCHAR expressions, if the code set supports these ASCII characters. If the current collation (as specified by **DB_LOCALE** or by SET COLLATION) supports a localized collating order, however, that localized order is used when the database server sorts NCHAR or NVARCHAR values.

Literal Numbers as Operands

You might obtain unexpected results if a literal number that you specify as an operand is not in a format that can exactly represent the data type of another value with which it is compared by a relational operator. Because of rounding errors, for example, a relational operator like = or the **equals()** operator function generally cannot return TRUE if one operand returns a FLOAT value and the other an INTEGER. For information about which of the built-in data types store values that can be exactly represented as literal numbers, see the section "Literal Number " on page 4-185.

Related Information

For a discussion of relational operators in the SELECT statement, see the *IBM Informix Guide to SQL: Tutorial*.

For a discussion of the GLS aspects of relational operators, see the *IBM Informix GLS User's Guide*.

Chapter 5. Other Syntax Segments

In This Chapter

Syntax segments are language elements, such as database object names or optimizer directives, that appear as a subdiagram reference in the syntax diagrams of some SQL or SPL statements. Most segments that can occur in only one statement are described in Chapter 2, “SQL Statements,” on page 2-1 or Chapter 3, “SPL Statements,” on page 3-1 within the description of the statement. For the sake of clarity, ease of use, and comprehensive treatment, however, most segments that can occur in various SQL or SPL statements, and that are not data types nor expressions, are discussed separately here.

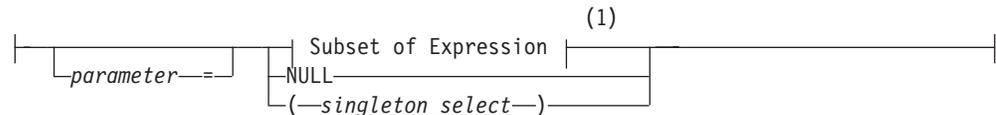
The previous chapter described the syntax segments that specify data types and expressions. This chapter describes additional syntax segments that are neither data types, expressions, nor complete SQL statements or SPL statements. These segments are referenced in various syntax diagrams that appear in Chapter 2, “SQL Statements,” on page 2-1 and in other chapters of this document.

Arguments

Use the Argument segment to pass a specific value as input to a routine. Use this segment wherever you see a reference to an *argument* in a syntax diagram.

Syntax

Argument:



Notes:

- 1 See “Subset of Expressions Valid as an Argument” on page 5-3

Element	Description	Restrictions	Syntax
<i>parameter</i>	A parameter whose value you specify	Must match a name that CREATE FUNCTION or CREATE PROCEDURE statement declared	“Identifier” on page 5-21
<i>singleton_select</i>	Embedded query that returns a single value	Must return exactly one value of a data type and length compatible with <i>parameter</i>	“SELECT” on page 2-528

Usage

The CREATE PROCEDURE or CREATE FUNCTION statement can define a parameter list for a UDR. If the parameter list is not empty, you must enter arguments when you invoke the UDR. An *argument* is a specific value whose data type is compatible with that of the corresponding UDR parameter.

When you execute a UDR, you can enter arguments in either of two ways:

- With a parameter name (in the form *parameter name = expression*), even if the arguments are not in the same order as the parameters
- By position, with no *parameter* name, where each *expression* is in the same order as the parameter to which the argument corresponds. (This is sometimes called *ordinal* format.)

You cannot mix these two ways of specifying arguments within a single invocation of a routine. If you specify a *parameter* name for one argument, for example, you must use parameter names for all the arguments.

In the following example, both statements are valid for a user-defined procedure that expects three character arguments, **t**, **d**, and **n**:

```
EXECUTE PROCEDURE add_col (t ='customer', d ='integer',
                          n ='newint');

EXECUTE PROCEDURE add_col ('customer','newint','integer') ;
```

Comparing Arguments to the Parameter List

When you create or register a UDR with CREATE PROCEDURE or CREATE FUNCTION, you declare a *parameter list* with the names and data types of the parameters that the UDR expects. (Parameter names are optional for external routines written in the C or Java languages.) See “Routine Parameter List” on page 5-67 for details of declaring parameters.

User-defined routines can be *overloaded*, if different routines have the same identifier, but have different numbers of declared parameters. For more information about overloading, see “Routine Overloading and Routine Signatures (IDS)” on page 5-19.

If you attempt to execute a UDR with more arguments than the UDR expects, you receive an error.

If you invoke a UDR with fewer arguments than the UDR expects, the omitted arguments are said to be *missing*. The database server initializes missing arguments to their corresponding default values. This initialization occurs before the first executable statement in the body of the UDR.

If missing arguments have no default values, Dynamic Server issues an error. Extended Parallel Server can invoke a routine with missing arguments that have no default values, but the routine might fail if a missing argument causes another error (for example, an undefined value for a variable).

Named parameters cannot be used to invoke UDRs that overload data types in their routine signatures. Named parameters are valid in resolving non-unique routine names only if the signatures have different numbers of parameters:

```
func( x::integer, y );    -- VALID if only these 2 routines
func( x::integer, y, z ); -- have the same 'func' identifier

func( x::integer, y );    -- NOT VALID if both routines have
func( x::float, y );      -- same identifier and 2 parameters
```

For both ordinal and named parameters, the routine with the fewest parameters is executed if two or more UDR signatures have multiple numbers of defaults:

```
func( x, y default 1 )
func( x, y default 1, z default 2 )
```

If two registered UDRs that are both called **func** have the signatures shown above, then the statement `EXECUTE func(100)` invokes **func(100,1)**.

You cannot supply a subset of default values using named parameters unless they are in the positional order of the routine signature. That is, you cannot skip a few arguments and rely on the database server to supply their default values.

For example, given the signature:

```
func( x, y default 1, z default 2 )
```

you can execute:

```
func( x=1, y=3 )
```

but you cannot execute:

```
func( x=1, z=3 )
```

Subset of Expressions Valid as an Argument

The diagram for “Arguments ” on page 5-1 refers to this section.

You can use any expression as an argument, except an aggregate function. If you use a subquery or function call as an argument, the subquery or function must return a single value of the appropriate data type and size. For the syntax and usage of SQL expressions, see “Expression” on page 4-40.

Arguments to UDRs in Remote Databases

UDRs are valid in cross-database and in cross-server distributed operations in most contexts where a UDR is valid in the local database, but every participating database must have the same logging mode.

Excluding `BIGSERIAL`, `BYTE`, `SERIAL`, `SERIAL8`, and `TEXT`, the data types that are valid as arguments to cross-server UDRs include the built-in SQL data types that are not opaque, as listed in “Data Types in Distributed Queries” on page 2-535, and these additional built-in opaque and `DISTINCT` data types:

- `BOOLEAN`
- `LVARCHAR`
- `DISTINCT` of built-in types that are not opaque
- `DISTINCT` of `BOOLEAN`
- `DISTINCT` of `LVARCHAR`
- `DISTINCT` of the `DISTINCT` types listed above.

These data types can be arguments to SPL, C, or Java language UDRs, if the UDRs are defined in all the participating databases. Any implicit or explicit casts defined over these data types must be duplicated across all the participating Dynamic Server instances. The `DISTINCT` data types must have exactly the same data type hierarchy defined in all databases that participate in the distributed query.

The same data types are valid as arguments in calls to UDRs in other databases of the same Dynamic Server instance, as well as arguments of the following additional types:

- `BLOB`
- `CLOB`
- UDTs that you cast explicitly to built-in types

All the UDRs, UDTs, DISTINCT data types, DISTINCT type hierarchies, casts, and cast functions must be registered in all of the participating databases. For more information on DISTINCT types in distributed operations, see “DISTINCT Types in Distributed Operations” on page 4-33.

Related Information

Related Statements: “ALTER FUNCTION ” on page 2-28, “ALTER PROCEDURE” on page 2-32, “ALTER ROUTINE” on page 2-33, “Receiving Input from the Called UDR” on page 3-4 “CREATE FUNCTION ” on page 2-113, “CREATE FUNCTION FROM ” on page 2-122, “CREATE PROCEDURE FROM ” on page 2-163, “EXECUTE FUNCTION ” on page 2-356, “EXECUTE PROCEDURE ” on page 2-364.

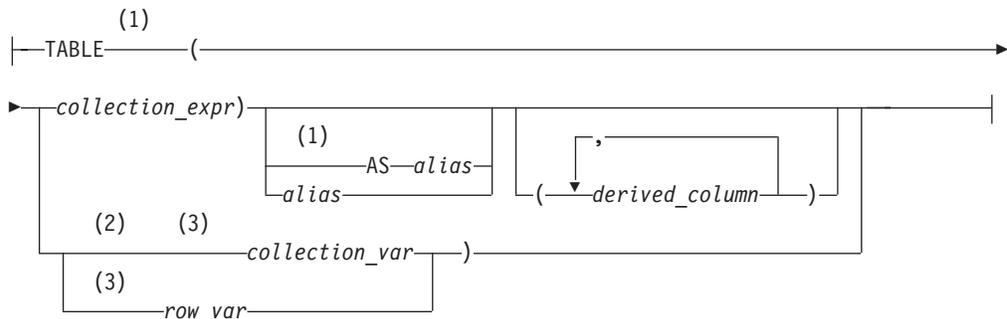
Related Segments: “Routine Parameter List” on page 5-67

Collection-Derived Table

A *collection-derived table* is a virtual table in which the values in the rows of the table are equivalent to elements of a collection. Use this segment where you see a reference to Collection-Derived Table in a syntax diagram. Only Dynamic Server supports this syntax, which is an extension to the ANSI/ISO standard for SQL.

Syntax

Collection-Derived Table:



Notes:

- 1 Informix extension
- 2 Stored Procedure Language
- 3 ESQL/C

Element	Description	Restrictions	Syntax
<i>alias</i>	Temporary name for a collection-derived table whose scope is a SELECT statement. The default is implementation dependent.	If potentially ambiguous, you must precede <i>alias</i> with the AS keyword. See “The AS Keyword” on page 2-545.	“Identifier ” on page 5-21
<i>collection_expr</i>	Any expression that evaluates to the elements of a single collection	See “Restrictions with the Collection-Expression Format” on page 5-6.	“Expression” on page 4-40

Element	Description	Restrictions	Syntax
<i>collection_var</i> , <i>row_var</i>	Name of a typed or untyped collection variable, or an Informix ESQL/C row variable that holds the collection-derived table	Must have been declared in an Informix ESQL/C program or (for <i>collection_var</i>) in an SPL routine	See the <i>IBM Informix ESQL/C Programmer's Manual</i> or "DEFINE " on page 3-8.
<i>derived_column</i>	Temporary name for a derived column in a table	If the underlying collection is not of a ROW data type, you can specify only one derived-column name	"Identifier " on page 5-21

Usage

A collection-derived table can appear where a *table* name is valid in the UPDATE statement, in the FROM clause of the SELECT or DELETE statement, or in the INTO clause of an INSERT statement.

Use the collection-derived-table segment to accomplish these tasks:

- Access the elements of a collection as you would the rows of a table.
- Specify a collection variable to access, instead of a table name.
- Specify an ESQL/C **row** variable to access, instead of a table name.

The TABLE keyword converts a collection into a virtual table. You can use the collection expression format to query a collection column, or you can use the **collection** variable or **row** variable format to manipulate the data in a collection column.

Accessing a Collection Through a Virtual Table

When you use the collection expression format of the collection-derived table segment to access the elements of a collection, you can select elements of the collection directly through a virtual table. You can use this format in the FROM clause of a SELECT statement. The FROM clause can be in either a query or a subquery.

With this format you can use joins, aggregates, the WHERE clause, expressions, the ORDER BY clause, and other operations that are not available when you use the collection-variable format. This format reduces the need for multiple cursors and temporary tables.

Examples of possible collection expressions include column references, scalar subquery, dotted expression, functions, operators (through overloading), collection subqueries, literal collections, collection constructors, cast functions, and so on.

The following example uses a SELECT statement in the FROM clause whose result set defines a virtual table consisting of the fifty-first through seventieth qualifying rows, ordered by the **employee_id** column value.

```
SELECT * FROM TABLE(MULTISET(SELECT SKIP 50 FIRST 20 * FROM employees
ORDER BY employee_id)) vt(x,y), tab2 WHERE tab2.id = vt.x;
```

The following example uses a join query to create a virtual table of no more than twenty rows (beginning with the 41st row), ordered by value in the **salary** column of the collection-derived table:

```

SELECT emp_id, emp_name, emp_salary
FROM TABLE(MULTISET(SELECT SKIP 40 LIMIT 20 id, name, salary
                     FROM e1, e2
                     WHERE e1.id = e2.id ORDER BY salary ))
AS etab(emp_id, emp_name, emp_salary);

```

Table Expressions in the FROM Clause

Dynamic Server supports ANSI/ISO standard syntax for table expressions in the FROM clause of SELECT queries and subqueries as a substitute for the Informix-extension collection-derived table syntax. The keywords TABLE and MULTISET were required in version 10.00 and in earlier releases. These extensions to the ANSI/ISO standard for SQL are supported but no longer required for collection-derived table specifications in the FROM clause of SELECT statements.

The following two queries return the same result set, but only the second query complies with the ANSI/ISO standard:

```

SELECT * FROM TABLE(MULTISET(SELECT col1 FROM tab1 WHERE col1 = 100))
AS vtab(c1),
(SELECT col1 FROM tab1 WHERE col1 = 10) AS vtab1(vc1) ORDER BY c1;

SELECT * FROM (SELECT col1 FROM tab1 WHERE col1 = 100) AS vtab(c1),
(SELECT col1 FROM tab1 WHERE col1 = 10) AS vtab1(vc1)
ORDER BY c1;

```

The same SELECT statement can combine instances of both the Informix-extension and ANSI/ISO syntax for derived tables:

```

SELECT * FROM (SELECT col1 FROM tab1 WHERE col1 = 100) AS vtab(c1),
TABLE(MULTISET(SELECT col1 FROM tab1 WHERE col1 = 10)) AS vtab1(vc1)
ORDER BY c1;

```

The subquery must be delimited by parentheses in both formats, but the outer set of parentheses (()) that immediately follows the TABLE keyword and encloses the MULTISET collection subquery specification is an extension to the ANSI/ISO syntax. This ANSI/ISO syntax is valid only in the FROM clause of the SELECT statement. You cannot omit these keywords and parentheses from a collection subquery specification in any other context.

Restrictions with the Collection-Expression Format

When you use the collection-expression format, certain restrictions apply:

- A collection-derived table is read-only.
 - It cannot be the target of INSERT, UPDATE, or DELETE statements.
 - To perform insert, update, and delete operations, you must use the collection-variable format.
 - It cannot be the underlying table of an updatable cursor or view.
- In the FROM clause of the SELECT statement, the CALL keyword of SPL cannot precede the TABLE keyword of a table expression.
- If the collection is a LIST data type, the resulting collection-derived table does not preserve the order of the elements in the LIST.
- The underlying collection expression cannot evaluate to NULL.
- The collection expression cannot contain a reference to a collection on a remote database server.
- The collection expression cannot contain column references to tables that appear in the same FROM clause. That is, the collection-derived table must be independent of other tables in the FROM clause.

For example, the following statement returns an error because the collection-derived table, TABLE (parents.children), refers to the **parents** table, which is also referenced in the FROM clause:

```
SELECT COUNT(*)
  FROM parents, TABLE(parents.children) c_table
 WHERE parents.id = 1001;
```

To counter this restriction, you might write a query that contains a subquery in the Projection clause:

```
SELECT (SELECT COUNT(*)
        FROM TABLE(parents.children) c_table)
  FROM parents WHERE parents.id = 1001;
```

Additional Restrictions That Apply to ESQL/C

In addition to the previously described restrictions, the following restrictions also apply when you use the collection-expression format with Informix ESQL/C:

- You cannot specify an untyped COLLECTION as the host-variable data type.
- You cannot use the format TABLE(?).

The data type of the underlying collection variable must be determined statically. To counter this restriction, you can explicitly cast the variable to a typed collection data type (SET, MULTISSET, or LIST) that the database server recognizes. For example,

```
TABLE(CAST(? AS type))
```

- You cannot use the format TABLE(:hostvar).

To counter this restriction, you must explicitly cast the variable to a typed collection data type (SET, MULTISSET, or LIST) that the database server recognizes. For example,

```
TABLE(CAST(:hostvar AS type))
```

Row Type of the Resulting Collection-Derived Table

If you do not specify a derived-column name, the behavior of the database server depends on the data types of the elements in the underlying collection.

Although a collection-derived table appears to contain columns of individual data types, these columns are, in fact, the fields of a ROW data type. The data type of the ROW type as well as the column name depend on several factors.

If the data type of the elements of the underlying collection expression is *type*, the database server determines the ROW type of the collection-derived table by the following rules:

- If *type* is a ROW data type, and no derived-column list is specified, then the ROW type of the collection-derived table is *type*.
- If *type* is a ROW data type and a derived column list is specified, then the ROW type of the collection-derived table is an unnamed ROW type whose column data types are the same as those of *type* and whose column names are taken from the derived column list.
- If *type* is not a ROW data type, the ROW type of the collection-derived table is an unnamed ROW type that contains one column of *type* and whose name is specified in the derived column list. If no name is specified, the database server assigns an implementation-dependent name to the column.

The extended examples that the following table shows illustrate these rules. The table uses the following schema for its examples:

```

CREATE ROW TYPE person (name CHAR(255), id INT);
CREATE TABLE parents
(
  name CHAR(255),
  id INT,
  children LIST (person NOT NULL)
);
CREATE TABLE parents2
(
  name CHAR(255),
  id INT,
  children_ids LIST (INT NOT NULL)
);

```

ROW Type	Explicit Derived-Column List	Resulting ROW Type of the Collection-Derived Table	Code Example
Yes	No	<i>Type</i>	<pre>SELECT (SELECT c_table.name FROM TABLE(parents.children) c_table WHERE c_table.id = 1002) FROM parents WHERE parents.id = 1001;</pre> <p>In this example, the ROW type of c_table is parents.</p>
Yes	Yes	Unnamed ROW type of which the column type is <i>Type</i> and the column name is the name in the derived-column list	<pre>SELECT (SELECT c_table.c_name FROM TABLE(parents.children) c_table(c_name, c_id) WHERE c_table.c_id = 1002) FROM parents WHERE parents.id = 1001;</pre> <p>In this example, the ROW type of c_table is ROW(c_name CHAR(255), c_id INT).</p>
No	No	Unnamed ROW that contains one column of <i>Type</i> that is assigned an implementation-dependent name	<p>In the following example, if you do not specify c_id, the database server assigns a name to the derived column. In this case, the ROW type of c_table is ROW(server_defined_name INT).</p>
No	Yes	Unnamed ROW type that contains one column of <i>Type</i> whose name is in the derived-column list	<pre>SELECT(SELECT c_table.c_id FROM TABLE(parents2.child_ids) c_table (c_id) WHERE c_table.c_id = 1002) FROM parents WHERE parents.id = 1001;</pre> <p>Here the ROW type of c_table is ROW(c_id INT).</p>

The following program fragment creates a collection-derived table using an SPL function that returns a single value:

```

CREATE TABLE wanted(person_id int);
CREATE FUNCTION
  wanted_person_count (person_set SET(person NOT NULL))
  RETURNS INT;
RETURN( SELECT COUNT (*)
  FROM TABLE (person_set) c_table, wanted
  WHERE c_table.id = wanted.person_id);
END FUNCTION;

```

The next program fragment shows the more general case of creating a collection-derived table using an SPL function that returns multiple values:

```

-- Table of categories and child categories,
-- allowing any number of levels of subcategories
CREATE TABLE CategoryChild (

```

```

        categoryId    INTEGER,
        childCategoryId SMALLINT
    );

INSERT INTO CategoryChild VALUES (1, 2);
INSERT INTO CategoryChild VALUES (1, 3);
INSERT INTO CategoryChild VALUES (1, 4);
INSERT INTO CategoryChild VALUES (2, 5);
INSERT INTO CategoryChild VALUES (2, 6);
INSERT INTO CategoryChild VALUES (5, 7);
INSERT INTO CategoryChild VALUES (7, 8);
INSERT INTO CategoryChild VALUES (7, 9);
INSERT INTO CategoryChild VALUES (4, 10);

-- "R" == ROW type
CREATE ROW TYPE categoryLevelR (
    categoryId    INTEGER,
    level        SMALLINT );

-- DROP FUNCTION categoryDescendants (
--     INTEGER, SMALLINT );
CREATE FUNCTION categoryDescendants (
    pCategoryId INTEGER,
    pLevel      SMALLINT DEFAULT 0 )
RETURNS MULTISET (categoryLevelR NOT NULL)

-- "p" == Prefix for Parameter names
-- "l" == Prefix for Local variable names
DEFINE lCategoryId LIKE CategoryChild.categoryId;
DEFINE lRetSet MULTISET (categoryLevelR NOT NULL);
DEFINE lCatRow categoryLevelR;

-- TRACE ON;
-- Must initialize collection before inserting rows
LET lRetSet = 'MULTISET{}' :: MULTISET (categoryLevelR NOT NULL);
FOREACH
SELECT childCategoryId INTO lCategoryId
FROM CategoryChild WHERE categoryId = pCategoryId;
INSERT INTO TABLE (lRetSet)
VALUES (ROW (lCategoryId, pLevel+1)::categoryLevelR);

-- INSERT INTO TABLE (lRetSet);
-- EXECUTE FUNCTION categoryDescendantsR ( lCategoryId,
-- pLevel+1 );
-- Need to iterate over results and insert into SET.
-- See the SQL Tutorial, pg. 10-52:
-- "Tip: You can only insert one value at a time
-- into a simple collection."
FOREACH
EXECUTE FUNCTION categoryDescendantsR ( lCategoryId, pLevel+1 )
INTO lCatRow;
INSERT INTO TABLE (lRetSet)
VALUES (lCatRow);
END FOREACH;
END FOREACH;

RETURN lRetSet;
END FUNCTION
;
-- "R" == recursive
-- DROP FUNCTION categoryDescendantsR (INTEGER, SMALLINT);
CREATE FUNCTION categoryDescendantsR (
    pCategoryId INTEGER,
    pLevel      SMALLINT DEFAULT 0
)
RETURNS categoryLevelR;
DEFINE lCategoryId LIKE CategoryChild.categoryId;

```

```

DEFINE lCatRow          categoryLevelR;

FOREACH
  SELECT  childCategoryId
  INTO    lCategoryId
  FROM    CategoryChild
  WHERE   categoryId = pCategoryId
  RETURN ROW (lCategoryId, pLevel+1)::categoryLevelR WITH RESUME;

      FOREACH
      EXECUTE FUNCTION categoryDescendantsR ( lCategoryId, pLevel+1 )
      INTO    lCatRow
      RETURN  lCatRow WITH RESUME;
      END FOREACH;
END FOREACH;
END FUNCTION;

-- Test the functions:
SELECT lev, col
FROM   TABLE ((
        categoryDescendants (1, 0)
      )) AS CD (col, lev);

```

Accessing a Collection Through a Collection Variable

When you use the collection-variable format of the collection-derived table segment, you use a host or program variable to access and manipulate the elements of a collection. This format allows you to modify the contents of a variable as you would a table in the database, and then update the actual table with the contents of the **collection** variable.

You can use the collection-variable format (the TABLE keyword preceding a **collection** variable) in place of the name of a table, synonym, or view in the following SQL statements (or in the FOREACH statement of SPL):

- The FROM clause of the SELECT statement to access an element of the **collection** variable
- The INTO clause of the INSERT statement to add a new element to the **collection** variable
- The DELETE statement to remove an element from the **collection** variable
- The UPDATE statement to modify an existing element in the **collection** variable
- The DECLARE statement to declare a Select or Insert cursor to access multiple elements of an Informix ESQL/C **collection** host variable
- The FETCH statement to retrieve a single element from a **collection** host variable that is associated with a Select cursor
- The PUT statement to retrieve a single element from a **collection** host variable that is associated with an Insert cursor
- The FOREACH statement to declare a cursor to access multiple elements of an SPL collection variable and to retrieve a single element from this **collection** variable

Using a Collection Variable to Manipulate Collection Elements

When you use data manipulation statements (SELECT, INSERT, UPDATE, or DELETE) of Dynamic Server in conjunction with a **collection** variable, you can modify one or more elements in a collection.

To modify elements in a collection

1. Create a **collection** variable in your SPL routine or Informix ESQL/C program. For information on how to declare a **collection** variable in Informix ESQL/C, see the *IBM Informix ESQL/C Programmer's Manual*. For information on how to define a **COLLECTION** variable in SPL, see "DEFINE " on page 3-8.
2. In Informix ESQL/C, allocate memory for the collection; see "ALLOCATE COLLECTION " on page 2-1.
3. Optionally, use a SELECT statement to select a COLLECTION column into the **collection** variable. If the variable is an untyped COLLECTION variable, you must perform a SELECT from the COLLECTION column before you use the variable in the collection-derived table segment. The SELECT statement allows the database server to obtain the collection data type.
4. Use the appropriate data manipulation statement with the collection-derived table segment to add, delete, or update elements in the collection variable. To insert more than one element or to update or delete a *specific* element of a collection, you must use a cursor for the collection variable.
 - For more information on how to use an update cursor with ESQL/C, see "DECLARE " on page 2-286.
 - For more information on how to use an update cursor with SPL, see "FOREACH " on page 3-22.
5. After the collection variable contains the correct elements, use an INSERT or UPDATE statement on the table or view that holds the actual collection column to save the changes that the collection variable holds.
 - With UPDATE, specify the collection variable in the SET clause.
 - With INSERT, specify the collection variable in the VALUES clause.

The collection variable stores the elements of the collection. It has no intrinsic connection, however, with a database column. Once the collection variable contains the correct elements, you must then save the variable into the actual collection column of the table with either an INSERT or an UPDATE statement.

Example of Deleting from a Collection in ESQL/C

Suppose that the **set_col** column of a row in the **table1** table is defined as a SET and for one row contains the values {1,8,4,5,2}. The following Informix ESQL/C code fragment uses an update cursor and a DELETE statement with a WHERE CURRENT OF clause to delete the element whose value is 4:

```
EXEC SQL BEGIN DECLARE SECTION;
    client collection set(smallint not null) a_set;
    int an_int;
EXEC SQL END DECLARE SECTION;
...
EXEC SQL allocate collection :a_set;
EXEC SQL select set_col into :a_set from table1 where int_col = 6;
EXEC SQL declare set_curs cursor for
    select * from table(:a_set) for update;

EXEC SQL open set_curs;
while (i<coll_size)
{
    EXEC SQL fetch set_curs into :an_int;
    if (an_int = 4)
    {
        EXEC SQL delete from table(:a_set) where current of set_curs;
        break;
    }
    i++;
}
```

```

EXEC SQL update table1 set set_col = :a_set
      where int_col = 6;
EXEC SQL deallocate collection :a_set;
EXEC SQL close set_curs;
EXEC SQL free set_curs;

```

After the DELETE statement executes, this collection variable contains the elements {1,8,5,2}. The UPDATE statement at the end of this code fragment saves the modified collection into the **set_col** column. Without this UPDATE statement, element 4 of the collection column is not deleted.

Example of Deleting from a Collection

Suppose that the **set_col** column of a row in the **table1** table is defined as a SET and one row contains the values {1,8,4,5,2}. The following SPL code fragment uses a FOREACH loop and a DELETE statement with a WHERE CURRENT OF clause to delete the element whose value is 4:

```

CREATE_PROCEDURE test6()

  DEFINE a SMALLINT;
  DEFINE b SET(SMALLINT NOT NULL);
  SELECT set_col INTO b FROM table1
    WHERE id = 6;
    -- Select the set in one row from the table
    -- into a collection variable
  FOREACH cursor1 FOR
    SELECT * INTO a FROM TABLE(b);
    -- Select each element one at a time from
    -- the collection derived table b into a
    IF a = 4 THEN
      DELETE FROM TABLE(b)
        WHERE CURRENT OF cursor1;
        -- Delete the element if it has the value 4
      EXIT FOREACH;
    END IF;
  END FOREACH;

  UPDATE table1 SET set_col = b
    WHERE id = 6;
    -- Update the base table with the new collection

END PROCEDURE;

```

This SPL routine declares two SET variables, **a** and **b**, each to hold a set of SMALLINT values. The first SELECT statement copies a SET column from one row of **table1** into variable **b**. The routine then declares a cursor called **cursor1** that copies one element at a time from **b** into SET variable **a**. When the cursor is positioned on the element whose value is 4, the DELETE statement removes that element from SET variable **b**. Finally, the UPDATE statement replaces the row of **table1** with the new collection that is stored in variable **b**.

For information on how to use collection variables in an SPL routine, see the *IBM Informix Guide to SQL: Tutorial*.

Example of Updating a Collection

Suppose that the **set_col** column of a table called **table1** is defined as a SET and that it contains the values {1,8,4,5,2}. The following Informix ESQL/C program changes the element whose value is 4 to a value of 10:

```

main
{
    EXEC SQL BEGIN DECLARE SECTION;
        int a;
        collection b;
    EXEC SQL END DECLARE SECTION;

    EXEC SQL allocate collection :b;
    EXEC SQL select set_col into :b from table1
        where int_col = 6;

    EXEC SQL declare set_curs cursor for
        select * from table(:b) for update;
    EXEC SQL open set_curs;
    while (SQLCODE != SQLNOTFOUND)
    {
        EXEC SQL fetch set_curs into :a;
        if (a = 4)
        {
            EXEC SQL update table(:b)(x)
                set x = 10 where current of set_curs;
            break;
        }
    }
    EXEC SQL update table1 set set_col = :b
        where int_col = 6;
    EXEC SQL deallocate collection :b;
    EXEC SQL close set_curs;
    EXEC SQL free set_curs;
}

```

After you execute this Informix ESQL/C program, the **set_col** column in **table1** contains the values {1,8,10,5,2}.

This Informix ESQL/C program defines two **collection** variables, **a** and **b**, and selects a SET from **table1** into **b**. The WHERE clause ensures that only one row is returned. Then the program defines a Collection cursor, which selects elements one at a time from **b** into **a**. When the program locates the element with the value 4, the first UPDATE statement changes that element value to 10 and exits the loop.

In the first UPDATE statement, **x** is a derived-column name used to update the current element in the collection-derived table. The second UPDATE statement updates the base table **table1** with the new collection.

For information on how to use **collection** host variables in an Informix ESQL/C program, see the discussion of complex data types in the *IBM Informix ESQL/C Programmer's Manual*.

Example of Inserting a Value into a Multiset Collection

Suppose the Informix ESQL/C host variable **a_multiset** has the following declaration:

```

EXEC SQL BEGIN DECLARE SECTION;
    client collection multiset(integer not null) a_multiset;
EXEC SQL END DECLARE SECTION;

```

The following INSERT statement adds a new MULTiset element of 142,323 to **a_multiset**:

```

EXEC SQL allocate collection :a_multiset;
EXEC SQL select multiset_col into :a_multiset from table1
    where id = 107;
EXEC SQL insert into table(:a_multiset) values (142323);

```

```
EXEC SQL update table1 set multiset_col = :a_multiset
      where id = 107;
```

```
EXEC SQL deallocate collection :a_multiset;
```

When you insert elements into a **client-collection** variable, you cannot specify a SELECT statement or an EXECUTE FUNCTION statement in the VALUES clause of the INSERT. When you insert elements into a **server-collection** variable, however, the SELECT and EXECUTE FUNCTION statements are valid in the VALUES clause. For more information on **client-** and **server-collection** variables, see the *IBM Informix ESQL/C Programmer's Manual*.

Accessing a Nested Collection

If the element of the collection is itself a complex type (**collection** or **row** type), the collection is a *nested collection*. For example, suppose the Informix ESQL/C **collection** variable, **a_set**, is a nested collection that is defined as follows:

```
EXEC SQL BEGIN DECLARE SECTION;
      client collection set(list(integer not null)) a_set;
      client collection list(integer not null) a_list;
      int an_int;
EXEC SQL END DECLARE SECTION;
```

To access the elements (or fields) of a nested collection, use a **collection** or **row** variable that matches the element type (**a_list** and **an_int** in the preceding code fragment) and a Select cursor.

Accessing a Row Variable

The TABLE keyword can make an Informix ESQL/C **row** variable a collection-derived table. That is, a row appears as a table in an SQL statement. For a **row** variable, think of the collection-derived table as a table of one row, with each field of the **row** type being a column of the row. Use the TABLE keyword in place of the name of a table, synonym, or view in these SQL statements:

- The FROM clause of the SELECT statement to access a field of the **row** variable
- The UPDATE statement to modify an existing field in the **row** variable

The DELETE and INSERT statements do not support a **row** variable in the collection-derived-table segment.

For example, suppose an ESQL/C host variable **a_row** has the following declaration:

```
EXEC SQL BEGIN DECLARE SECTION;
      row(x int, y int, length float, width float) a_row;
EXEC SQL END DECLARE SECTION;
```

The following ESQL/C code fragment adds the fields in the **a_row** variable to the **row_col** column of the **tab_row** table:

```
EXEC SQL update table(:a_row)
      set x=0, y=0, length=10, width=20;
EXEC SQL update rectangles set rect = :a_row;
```

Related Information

"DECLARE " on page 2-286, "DELETE " on page 2-303, "DESCRIBE" on page 2-311, "FETCH " on page 2-367, "INSERT " on page 2-430, "PUT " on page 2-479, "SELECT" on page 2-528, "UPDATE" on page 2-689, "DEFINE " on page 3-8, and "FOREACH " on page 3-22.

For information on how to use COLLECTION variables in an SPL routine, see the *IBM Informix Guide to SQL: Tutorial*.

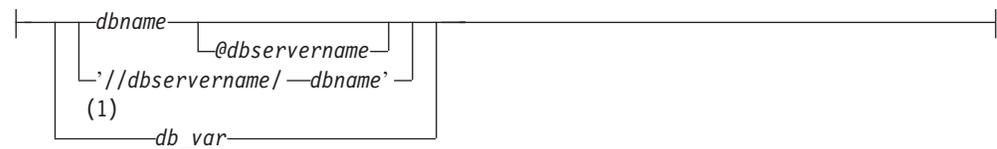
For information on how to use **collection** or **row** variables in an Informix ESQL/C program, see the chapter on complex data types in the *IBM Informix ESQL/C Programmer's Manual*.

Database Name

Use the Database Name segment to specify the name of a database. Use this segment when you see a reference to a database name in a syntax diagram.

Syntax

Database Name:



Notes:

- 1 ESQL/C only

Element	Description	Restrictions	Syntax
<i>dbname</i>	Database name (with no pathname nor database server name)	Must be unique among the names of databases of the database server	"Identifier " on page 5-21
<i>dbservername</i>	Database server on which the database <i>dbname</i> resides	Must exist. No blank space can separate @ from <i>dbservername</i> .	"Identifier " on page 5-21
<i>db_var</i>	Host variable whose value specifies a database environment	Variable must be a fixed-length character data type	Language specific

Usage

Database names are not case sensitive. You cannot use delimited identifiers for a database name.

The identifiers *dbname* and *dbservername* can each have a maximum of 128 bytes.

If the name of a database server is a delimited identifier or if it includes uppercase letters, that database server cannot participate in cross-server distributed DML operations. To avoid this restriction, use only undelimited names that include no uppercase letters when you declare the name or the alias of a database server.

In a nondefault locale, *dbname* can include alphabetic characters from the code set of the locale. In a locale that supports a multibyte code set, keep in mind that the maximum length of the database name refers to the number of bytes, not the number of characters. For more information on the GLS aspects of naming databases, see the *IBM Informix GLS User's Guide*.

Using Keywords as Table Names

You can choose a database on another database server as your current database by specifying a database server name. The database server that *dbservername* specifies must match the name of a database server that is listed in your **sqlhosts** information.

Using the @ Symbol

The @ symbol is a literal character. If you specify a database server name, blank spaces are not valid between the @ symbol and the database server name. Either put a blank space between *dbname* and the @ symbol, or omit the blank space.

The following examples show valid database specifications, qualified by the database server name:

```
empinfo@personnel  
empinfo @personnel
```

In these examples, **empinfo** is the name of the database and **personnel** is the name of the database server.

Using a Path-Type Naming Notation

If you specify a pathname, do not put blank spaces between the quotation marks, slashes, and names. The following example specifies a valid UNIX pathname:

```
'//personnel/empinfo'
```

Here **empinfo** is the *dbname* and **personnel** is the name of the database server.

Using a Host Variable

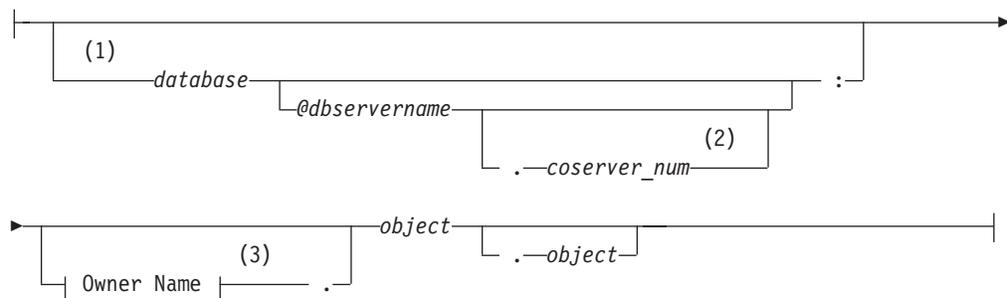
You can use a host variable within an Informix ESQL/C application to store a value that represents a database environment.

Database Object Name

Use the Database Object Name segment to specify the name of a database object, such as a column, table, view, or user-defined routine. Use this segment whenever you see a reference to a database object name.

Syntax

Database Object Name:



Notes:

- 1 Informix extension
- 2 Extended Parallel Server only

3 See “ Owner Name ” on page 5-45

Element	Description	Restrictions	Syntax
<i>coserver_num</i>	Coserver where <i>object</i> resides	Must exist.	“Literal Number ” on page 4-185
<i>database</i>	Database where <i>object</i> resides	Must exist.	“Database Name” on page 5-15
<i>dbservername</i>	Database server of <i>database</i>	Must exist. No space after @.	“Identifier ” on page 5-21
<i>object</i>	Name of a database object	See “Usage.”	“Identifier ” on page 5-21

Usage

A database object name can include qualifiers and separator symbols to specify a database, a server, a coserver (for XPS only), an owner, and (for some objects) another object of which the current database object is a component. For example, this expression specifies the **unit-price** column of the **stock** table, owned by user **informix**, in the **stores_demo** database of a database server called **butler**:

```
stores_demo@butler:informix.stock.unit_price
```

If you are creating or renaming a database object, the new name that you declare must be unique among objects of the same type in the database. Thus, the name of a new view must be unique among the names and synonyms of tables, views, and sequence objects that already exist in the same database. (But a view can have the same name as a view in a different database of the same server, or the same name as a trigger, for example, because these are different types of objects.)

In an ANSI-compliant database, the *owner.object* combination must be unique in the database for the type of object. A database object specification must include the owner name for a database object that you do not own. For example, if you specify a table that you do not own, you must also specify the owner of the table. The owner of all the system catalog tables is **informix**.

In Dynamic Server, the uniqueness requirement does not apply to the name of a user defined routine (UDR). For more information, see “Routine Overloading and Routine Signatures (IDS)” on page 5-19.

Characters from the code set of your database locale are valid in database object names. For more information, see *IBM Informix GLS User’s Guide*.

Specifying a Database Object in an External Database

You can specify a database object in either an external database on the local database server or in an external database on a remote database server.

Specifying a Database Object in a Cross-Database Query

To specify an object in another database of the local database server, you must qualify the identifier of the object with the name of the database (and of the owner, if the external database is ANSI compliant), as in this example:

```
corp_db:hrdirector.executives
```

In this example, the name of the external database is **corp_db**. The name of the owner of the table is **hrdirector**. The name of the table is **executives**. Here the colon (:) separator is required after the *database* qualifier.

In Dynamic Server, queries and other data manipulation language (DML) operations on other databases of the local database server can access most of the *built-in opaque data types*, as listed in “Data Types in Cross-Database Transactions” on page 2-536. DML operations can also access user-defined data types (UDTs) that can be cast to built-in types, as well as DISTINCT types that are based on built-in types, if each DISTINCT types and UDT is cast explicitly to a built-in type, and if all the DISTINCT types, UDTs, and casts are defined in all of the participating databases. The same data-type restrictions also apply to the arguments and to the returned values of a user-defined routine (UDR) that accesses other databases of the local Dynamic Server instance, if the UDR is defined in all of the participating databases.

Specifying a Database Object in a Cross-Server Query

To specify an object in a database of a remote database server, you must use a *fully-qualified identifier* that specifies the database, database server, and owner (if the external database is ANSI compliant) in addition to the database object name. For example, **hr_db@remoteoffice:hrmanager.employees** is a fully-qualified table name.

Here the database is **hr_db**, the database server is **remoteoffice**, the table owner is **hrmanager**, and the table name is **employees**. The at (@) separator, with no blank spaces, is required between the *database* and *database server* qualifiers. Cross-server queries can only access columns of built-in data types that are not opaque data types. (You cannot access UDTs, nor opaque, complex, or other extended data types in cross-server operations.)

In Dynamic Server, if a UDR exists on a remote database server, you must specify a fully-qualified identifier for the UDR. Like cross-server DML operations, a remote UDR is limited to built-in non-opaque data types for its arguments, parameters, and returned values.

You can refer to a remote database object in the following statements only. For information on the support in these statements across databases of the local server, or across database servers, refer to the *IBM Informix Guide to SQL: Tutorial*.

- CREATE DATABASE
- CREATE SYNONYM
- CREATE VIEW
- DATABASE
- DELETE
- EXECUTE FUNCTION
- EXECUTE PROCEDURE
- INFO
- INSERT
- LOAD
- LOCK TABLE
- SELECT
- UNLOAD
- UNLOCK TABLE
- UPDATE

If the name of a database server is a delimited identifier or if it includes uppercase letters, that database server cannot participate in distributed DML operations. To

avoid this restriction, use only undelimited names that include no uppercase letters when you declare the name or the alias of a database server.

Routine Overloading and Routine Signatures (IDS)

Because of routine overloading, the name of a user-defined routine does not need to be unique to the database. You can define more than one UDR with the same name, provided that the *routine signature* for each UDR is different.

UDRs are uniquely identified by their signatures. The signature of a UDR includes the following items of information:

- The type of routine (function or procedure)
- The identifier of the routine
- The cardinality, data type, and order of the parameters
- In an ANSI-compliant database, the owner name

For any given UDR, at least one item in the routine signature must be unique among all the UDRs registered in the database.

In a database that is not ANSI-compliant, two routines that have different owners cannot have the same signature, except in the special cases of the **sysdbopen()** and **sysdbclose()** routines. For information about the effects of these session configuration routines when their owners connect to or disconnect from a database where these routines are defined, see “**IFX_REPLACE_MODULE Function**” on page 6-54.

Specifying an Existing UDR

To reference an existing UDR by a name that does not uniquely identify the UDR, you must also specify the parameter data types after the UDR name, in the same order that they were declared when the UDR was created. Dynamic Server then uses routine resolution rules to identify the instance of the UDR to alter, drop, or execute. As an alternative, you can specify its *specific name*, if one was declared when the UDR was created. Specific names are described in the section “**Specific Name**” on page 5-73. For more details of routine resolution, see “**Comparing Arguments to the Parameter List**” on page 5-2, and *IBM Informix User-Defined Routines and Data Types Developer’s Guide*.

Owners of Objects Created by UDRs

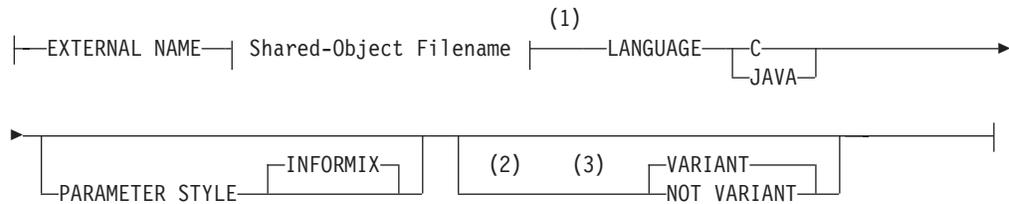
When a DDL statement within an owner-privileged UDR creates a new database object, the owner of the routine (rather than the user who executes it, if that user is not the owner of the routine) becomes the owner of the new database object. For a DBA-privileged UDR, however, the user who executes the routine (and who must hold DBA privilege) becomes the owner of any objects that the UDR creates.

External Routine Reference

Use an External Routine Reference when you write an external routine. Only Dynamic Server supports this syntax segment.

Syntax

External Routine Reference:



Notes:

- 1 See “ Shared-Object Filename ” on page 5-70
- 2 C
- 3 Java

Usage

If the IFX_EXTEND_ROLE configuration parameter is set to ON, authorization to use this segment is available only to the Database Server Administrator (DBSA), and to users whom the DBSA has granted the EXTEND role. By default, the DBSA is user **informix**.

This segment specifies the following information about an external routine:

- Pathname to the executable object code, stored in a shared-object file
 For C routines, this file is either a DLL or a shared library, depending on your operating system.
 For Java routines, this file is a jar file. Before you can create a UDR written in the Java language, you must assign a jar identifier to the external jar file with the **sqlj.install_jar** procedure. For more information, see “sqlj.install_jar” on page 6-58.
- The name of the programming language in which the UDR is written
- The parameter style of the UDR
 By default, the parameter style is INFORMIX. (This implies that if you specify OUT or INOUT parameters, the OUT or INOUT values are passed by reference.)
- The VARIANT or NOT VARIANT option, if you specify one. (This option is not available for SPL routines.)

VARIANT or NOT VARIANT Option

A function is *variant* if it can return different results when it is invoked with the same arguments or if it modifies the state of a database or of a variable. For example, a function that returns the current date or time is a variant function.

By default, user-defined functions are variant. If you specify NOT VARIANT when you create or modify a function, it cannot contain any SQL statements.

If the function is nonvariant, the database server might cache the return variant functions. For more information on functional indexes, see “CREATE INDEX ” on page 2-124.

To register a nonvariant function, add the NOT VARIANT option in this clause or in the Routine Modifier clause that is discussed in “Routine Modifier” on page 5-59

5-59. If you specify the modifier in both contexts, however, you must use the same modifier (either VARIANT or NOT VARIANT) in both clauses.

Example of a C User-Defined Function

The next example registers an external function named `equal()` that takes two `point` data type values as arguments. In this example, `point` is an opaque data type that specifies the `x` and `y` coordinates of a two-dimensional point.

```
CREATE FUNCTION equal( a point, b point ) RETURNING BOOLEAN;
  EXTERNAL NAME "/usr/lib/point/lib/libbtype1.so(point1_equal)"
  LANGUAGE C
END FUNCTION;
```

The function returns a single value of type `BOOLEAN`. The external name specifies the path to the C shared-object file where the object code of the function is stored. The external name indicates that the library contains another function, `point1_equal()`, which is invoked while `equal()` executes.

Examples

The following example includes an external routine reference for a Java language UDR. You must first register `demo_jar` using the procedure `install_jar(<absolute path><jar file name>,<internal registered name>)`.

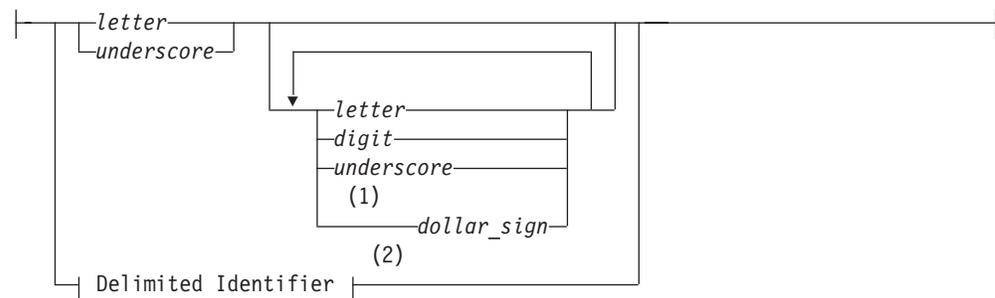
```
CREATE FUNCTION delete_order(int) RETURNING int
  EXTERNAL NAME 'informix.demo_jar:delete_order.delete_order()'
  LANGUAGE JAVA;
```

Identifier

An *identifier* specifies the unqualified name of a database object, such as an access method, aggregate, alias, blob space, cast, column, constraint, correlation, data type, index, operator class, partition, procedure, table, trigger, sequence, synonym, or view. Use the Identifier segment whenever you see a reference to an identifier in a syntax diagram.

Syntax

Identifier:



Notes:

- 1 Dynamic Server only
- 2 See “Delimited Identifiers” on page 5-23

Element	Description	Restrictions	Syntax
<i>digit</i>	Integer in range 0 to 9	Cannot be the first character	“Literal Number ” on page 4-185
<i>dollar_sign</i>	Dollar (\$) symbol	Cannot be the first character	Literal symbol entered from the keyboard.
<i>letter</i>	Upper- or lowercase letter of the alphabet	In the default locale, must be an ASCII character in the range A to Z or a to z	Literal symbol entered from the keyboard.
<i>underscore</i>	Underscore (_) character	Cannot substitute a space, hyphen, or other non-alphanumeric character	Literal symbol entered from the keyboard.

Usage

This is a logical subset of “Database Object Name ” on page 5-16, a segment that can specify the *owner*, *database*, and *database server* of external objects.

To include other non-alphanumeric symbols, such as a blank space (ASCII 32), in an identifier, you must use a delimited identifier. It is recommended that you do not use the dollar sign (\$) in identifiers, because this symbol is a special character whose inclusion in an identifier might cause conflicts with other syntax elements. For more information, see “Delimited Identifiers” on page 5-23.

An identifier must have a length of at least 1 byte, but no more than 128 bytes. For example, **employee_information** is valid as a table name. If you are using a multibyte code set, keep in mind that the maximum length of an identifier refers to the number of bytes, not to the number of logical characters.

For letter characters in nondefault locales, see “ Support for Non-ASCII Characters in Identifiers” on page 5-23. For further information on the GLS aspects of identifiers, see Chapter 3 of the *IBM Informix GLS User’s Guide*.

When you use ESQL/C with Dynamic Server, the database server checks the internal version number of the client application and the setting of the **IFX_LONGID** environment variable to determine whether a client application supports long identifiers (up to 128 bytes in length). For more information, see the *IBM Informix Guide to SQL: Reference*.

When the database server uses long identifiers, you might encounter error messages, warning messages, or other messages that truncate trailing characters in SQL identifiers or elsewhere in the message text. Truncation can usually be avoided, however, if identifiers have 18 or fewer bytes. Your code might be difficult to read or to maintain if identifiers of different SQL objects are identical in their first 18 characters.

Use of Uppercase Characters

You can specify the name of a database object with uppercase characters, but the database server shifts these to lowercase characters unless the **DELIMITED** environment variable is set and the identifier of the database object is enclosed between double (") quotation marks. In this case, the database server treats the name of the database object as a delimited identifier and preserves the uppercase characters in the name, as described in “Delimited Identifiers” on page 5-23.

If the name of a database server includes uppercase letters, that database server cannot participate in distributed DML operations. To avoid this restriction, use only un delimited names that include no uppercase letters when you declare the name or the alias of a database server.

Use of Keywords as Identifiers

Although you can use almost any word as an identifier, syntactic ambiguities can result from using keywords as identifiers in SQL statements. The statement might fail or might not produce the expected results. For a discussion of the syntactic ambiguities that can result from using keywords as identifiers and an explanation of workarounds for these problems, see “Potential Ambiguities and Syntax Errors” on page 5-26.

Delimited identifiers provide the easiest and safest way to use a keyword as an identifier without syntactic ambiguities. No workarounds are necessary for a keyword as a delimited identifier. For the syntax and usage of delimited identifiers, see “Delimited Identifiers.” Delimited identifiers require, however, that your code always use single (') quotation marks, rather than double (") quotation marks, to delimit character-string literals.

For the keywords of the implementation of SQL in Dynamic Server, see Appendix A, “Reserved Words for IBM Informix Dynamic Server,” on page A-1.

The keywords of SQL in Extended Parallel Server are listed in Appendix B, “Reserved Words for IBM Informix Extended Parallel Server,” on page B-1.

Tip: If an error message seems unrelated to the statement that caused the error, check to see if the statement uses a keyword as an un delimited identifier.

Support for Non-ASCII Characters in Identifiers

In a nondefault locale, you can use any alphabetic character that your locale recognizes as a *letter* in an SQL identifier. This feature enables you to use non-ASCII characters in the names of some database objects. For objects that support non-ASCII characters, see the *IBM Informix GLS User's Guide*.

Delimited Identifiers

By default, the character set of a valid SQL identifier is restricted to letters, digits, underscore, and dollar-sign symbols. If you set the **DELIMIDENT** environment variable, however, SQL identifiers can also include additional characters from the code set implied by the setting of the **DB_LOCALE** environment variable.

Delimited Identifier:



Element	Description	Restrictions	Syntax
<i>digit</i>	Integer in the range 0 to 9	Cannot be the first character	“Literal Number ” on page 4-185

Element	Description	Restrictions	Syntax
<i>letter</i>	Letter that forms part of the delimited identifier	Letters in delimited identifiers are case-sensitive	Literal value entered from the keyboard.
<i>other _character</i>	Nonalphanumeric character, such as #, \$, or blank space	Must be an element in the code set of the database locale	Literal value entered from the keyboard.
<i>underscore</i>	Underscore (<code>_</code>) symbol in the delimited identifier	Cannot include more than 128	Literal value entered from the keyboard.

If the database supports delimited identifiers, any double quotation marks (`"`) enclose an SQL identifier in your code, and only single (`'`) quotation marks, rather than double (`"`) quotation marks, delimit character-string literals.

Delimited identifiers enable you to declare names that are otherwise identical to SQL keywords, such as `TABLE`, `WHERE`, `DECLARE`, and so on. The only type of object for which you cannot specify a delimited identifier is a database name.

Letters in delimited identifiers are case sensitive. If you are using the default locale, *letter* must be an upper- or lowercase character in the range a to z or A to Z (in the ASCII code set). If you are using a nondefault locale, *letter* must be an alphabetic character that the locale supports. For more information, see “ Support for Non-ASCII Characters in Delimited Identifiers (GLS).”

Delimited identifiers are compliant with the ANSI/ISO standard for SQL.

When you create a database object, avoid including leading blank spaces or other white space characters between the first delimiting quotation mark and the first nonblank character of the delimited identifier. (Otherwise, you might not be able to reference the object in some contexts.)

If the name of a database server is a delimited identifier or if it includes uppercase letters, that database server cannot participate in distributed DML operations. To avoid this restriction, use only undelimited names that include no uppercase letters when you declare the name or the alias of a database server.

Support for Nonalphanumeric Characters

You can use delimited identifiers to specify nonalphanumeric characters in the names of database objects. You cannot use delimited identifiers, however, to specify nonalphanumeric characters in the names of storage objects such as `dbspaces`, `partitions`, and `blobspaces`.

Support for Non-ASCII Characters in Delimited Identifiers (GLS)

When you are using a nondefault locale whose code set supports non-ASCII characters, you can specify those non-ASCII characters in most delimited identifiers. The rule is that if you can specify non-ASCII characters in the undelimited form of the identifier, you can also specify non-ASCII characters in the delimited form of the same identifier. For a list of identifiers that support non-ASCII characters and for information on non-ASCII characters in delimited identifiers, see the *IBM Informix GLS User's Guide*.

Enabling Delimited Identifiers

To use delimited identifiers, you must set the `DELIMIDENT` environment variable. While `DELIMIDENT` is set, strings enclosed in double quotation marks (`"`) are treated as identifiers of database objects, and strings enclosed in single quotation

marks (') are treated as literal strings. If the **DELIMIDENT** environment variable is not set, however, strings enclosed in double quotation marks are also treated as literal strings.

If **DELIMIDENT** is set, the **SELECT** statement in the following example must be in single quotation marks in order to be treated as a quoted string:

```
PREPARE ... FROM 'SELECT * FROM customer';
```

If a delimited identifier is used in the **SELECT** statement that defines a view, then the **DELIMIDENT** environment variable must be set in order for the view to be accessed, even if the view name itself contains no special characters.

On UNIX and Linux systems, you can set **DELIMIDENT** by the procedures for setting environment variables that are described in *IBM Informix Guide to SQL: Reference*.

On Windows systems, you can set **DELIMIDENT** in various ways, which generally have the following descending order of precedence:

1. The setting of **DELIMIDENT** in the connection string when connecting
2. The setting of the **SQL_INFX_ATTR_DELIMIDENT** connection attribute before connecting
3. The setting of **DELIMIDENT** in **setnet32** with the **Use my settings** box selected
4. The setting of **DELIMIDENT** in **setnet32** with the **Use my settings** box cleared
5. The setting of **DELIMIDENT** on the command line before running the application
6. The setting of **DELIMIDENT** in Windows as a user variable
7. The setting of **DELIMIDENT** in Windows as a system variable
8. The default value (of no support for delimited identifiers).

This general order of precedence for Windows clients is sensitive, however, to the API through which you connect to the database, which can also affect the meaning of the setting and the default value. Refer to the documentation of your specific API for more information about the **DELIMIDENT** setting in Windows.

Examples of Delimited Identifiers

The next example shows how to create a table with a case-sensitive name:

```
CREATE TABLE "Proper_Ranger" (...);
```

The following example creates a table whose name includes a white space character. If the table name were not enclosed by double (") quotation marks, and if **DELIMIDENT** were not set, you could not use a blank space in the identifier.

```
CREATE TABLE "My Customers" (...);
```

The next example creates a table that has a keyword as the table name:

```
CREATE TABLE "TABLE" (...);
```

The following example for Dynamic Server shows how to delete all the rows from a table that is named **FROM** when you omit the keyword **FROM** in the **DELETE** statement:

```
DELETE "FROM";
```

Using Double Quotation Marks in a Delimited Identifier

To include a double quotation mark (") character within a delimited identifier, you must precede the double quotation mark (") with another double quotation mark ("). The following statement fragment specifies **My "Good" Data** as a table name:

```
CREATE TABLE "My ""Good"" Data" (...);
```

Potential Ambiguities and Syntax Errors

IBM does not recommend using any keyword of SQL as an identifier, because to do so tends to make your code more difficult to read and to maintain. If you ignore this potential problem for human readers, however, you can use almost any keyword as an SQL identifier, but various syntactic ambiguities can occur. An ambiguous statement might not produce the desired results. The following sections identify some potential ambiguities and workarounds when keywords are declared as identifiers, or when different database objects have the same identifier.

Using the Names of Built-In Functions as Column Names

The following two examples show a workaround for using a built-in function as a column name in a SELECT statement. This workaround applies to the built-in aggregate functions (AVG, COUNT, MAX, MIN, SUM) as well as the function expressions (algebraic, exponential and logarithmic, time, HEX, length, DBINFO, trigonometric, and TRIM functions).

Using **avg** as a column name causes the next example to fail because the database server interprets **avg** as an aggregate function rather than as a column name:

```
SELECT avg FROM mytab; -- fails
```

If the **DELIMITED** environment variable is set, you could use **avg** as a column name as the following example shows:

```
SELECT "avg" from mytab; -- successful
```

The workaround in the following example removes ambiguity by including a table name with the column name:

```
SELECT mytab.avg FROM mytab;
```

If you use the keyword **TODAY**, **CURRENT**, **SYSDATE**, or **USER** as a column name, ambiguity can occur, as the following example shows:

```
CREATE TABLE mytab (user char(10),
    CURRENT DATETIME HOUR TO SECOND,TODAY DATE);

INSERT INTO mytab VALUES('josh','11:30:30','1/22/2008');

SELECT user,current,today FROM mytab;
```

The database server interprets **user**, **current**, and **today** in the SELECT statement as the built-in functions **USER**, **CURRENT**, and **TODAY**. Thus, instead of returning **josh, 11:30:30,1/22/2008**, the SELECT statement returns the current user name, the current time, and the current date. The **SYSDATE** keyword has a similar effect in databases of Dynamic Server.

If you want to select the actual columns of the table, you must write the SELECT statement in one of the following ways:

```
SELECT mytab.user, mytab.current, mytab.today FROM mytab;

EXEC SQL select * from mytab;
```

Using Keywords as Column Names

Specific workarounds exist for using a keyword as a column name in a SELECT statement or other SQL statement. In some cases, more than one suitable workaround might be available.

Using ALL, DISTINCT, or UNIQUE as a Column Name

If you want to use the ALL, DISTINCT, or UNIQUE keywords as column names in a SELECT statement, you can take advantage of a workaround.

First, consider what happens when you try to use one of these keywords without a workaround. In the following example, using **all** as a column name causes the SELECT statement to fail because the database server interprets **all** as a keyword rather than as a column name:

```
SELECT all FROM mytab -- fails;
```

You must use a workaround to make this SELECT statement execute successfully. If the **DELIMIDENT** environment variable is set, you can use **all** as a column name by enclosing **all** in double quotation marks. In the following example, the SELECT statement executes successfully because the database server interprets **all** as a column name:

```
SELECT "all" from mytab; -- successful
```

The workaround in the following example uses the keyword ALL with the column name **all**:

```
SELECT ALL all FROM mytab;
```

The examples that follow show workarounds for using the keywords UNIQUE or DISTINCT as a column name in a CREATE TABLE statement.

The next example fails to declare a column named **unique** because the database server interprets **unique** as a keyword rather than as a column name:

```
CREATE TABLE mytab (unique INTEGER); -- fails
```

The following workaround uses two SQL statements. The first statement creates the column **mycol**; the second statement renames the column **mycol** to **unique**:

```
CREATE TABLE mytab (mycol INTEGER);

RENAME COLUMN mytab.mycol TO unique;
```

The workaround in the following example also uses two SQL statements. The first statement creates the column **mycol**; the second alters the table, adds the column **unique**, and drops the column **mycol**:

```
CREATE TABLE mytab (mycol INTEGER);

ALTER TABLE mytab
  ADD (unique INTEGER),
  DROP (mycol);
```

Using INTERVAL or DATETIME as a Column Name

The examples in this section show workarounds for using the keyword INTERVAL (or DATETIME) as a column name in a SELECT statement.

Using **interval** as a column name causes the following example to fail because the database server interprets **interval** as a keyword and expects it to be followed by an INTERVAL qualifier:

```
SELECT interval FROM mytab; -- fails
```

If the **DELIMITED** environment variable is set, you could use **interval** as a column name, as the following example shows:

```
SELECT "interval" from mytab; -- successful
```

The workaround in the following example removes ambiguity by specifying a table name with the column name:

```
SELECT mytab.interval FROM mytab;
```

The workaround in the following example includes an owner name with the table name:

```
SELECT josh.mytab.interval FROM josh.mytab;
```

Using rowid as a Column Name (IDS)

Every nonfragmented table has a virtual column named **rowid**. To avoid ambiguity, you cannot use **rowid** as a column name. Performing the following actions causes an error:

- Creating a table or view with a column named **rowid**
- Altering a table by adding a column named **rowid**
- Renaming a column to **rowid**

You can, however, use the term **rowid** as a table name.

```
CREATE TABLE rowid (column INTEGER, date DATE, char CHAR(20));
```

Important: It is recommended that you use primary keys as an access method, rather than exploiting the **rowid** column.

Using Keywords as Table Names

Examples in this section show workarounds that involve owner naming when the keyword **STATISTICS** or **OUTER** is a table name. (This workaround also applies to **STATISTICS** or **OUTER** as a view name or synonym.)

Using **statistics** as a table name causes the following example to fail because the database server interprets it as part of the UPDATE STATISTICS syntax rather than as a table name in an UPDATE statement:

```
UPDATE statistics SET mycol = 10;
```

The workaround in the following example specifies an owner name with the table name, to avoid ambiguity:

```
UPDATE josh.statistics SET mycol = 10;
```

Using **outer** as a table name causes the following example to fail because the database server interprets **outer** as a keyword for performing an outer join:

```
SELECT mycol FROM outer; -- fails
```

The following successful example uses owner naming to avoid ambiguity:

```
SELECT mycol FROM josh.outer;
```

Workarounds that Use the Keyword AS

In some cases, although a statement is not ambiguous and the syntax is correct, the database server returns a syntax error. The preceding pages show existing syntactic workarounds for several situations. You can use the AS keyword to provide a workaround for the exceptions.

You can use the AS keyword in front of column labels or table aliases.

The following example uses the AS keyword with a column label:

```
SELECT column_name AS display_label FROM table_name;
```

The following example uses the AS keyword with a table alias:

```
SELECT select_list FROM table_name AS table_alias;
```

Using AS with Column Labels

The examples in this section show workarounds that use the AS keyword with a column label. The first two examples show how you can use the keyword UNITS (or YEAR, MONTH, DAY, HOUR, MINUTE, SECOND, or FRACTION) as a column label.

Using **units** as a column label causes the next example to fail because the database server interprets it as part of an INTERVAL expression in which the **mycol** column is the operand of the UNITS operator:

```
SELECT mycol units FROM mytab;
```

The workaround in the following example includes the AS keyword:

```
SELECT mycol AS units FROM mytab;
```

The following examples use the AS or FROM keyword as a column label.

Using **as** as a column label causes the following example to fail because the database server interprets **as** as identifying **from** as a column label and thus finds no required FROM clause:

```
SELECT mycol as from mytab; -- fails
```

The following successful example repeats the AS keyword:

```
SELECT mycol AS as from mytab;
```

Using **from** as a column label causes the following example to fail because the database server expects a table name to follow the first **from**:

```
SELECT mycol from FROM mytab; -- fails
```

This example uses the AS keyword to identify the first **from** as a column label:

```
SELECT mycol AS from FROM mytab;
```

Using AS with Table Aliases

Examples in this section show workarounds that use the AS keyword with a table alias. The first pair shows how to use the ORDER, FOR, GROUP, HAVING, INTO, UNION, WITH, CREATE, GRANT, or WHERE keyword as a table alias.

Using **order** as a table alias causes the following example to fail because the database server interprets **order** as part of an ORDER BY clause:

```
SELECT * FROM mytab order; -- fails
```

The workaround in the following example uses the keyword AS to identify **order** as a table alias:

```
SELECT * FROM mytab AS order;
```

The next two examples show how to use the keyword WITH as a table alias.

Using **with** as a table alias causes the next example to fail because the database server interprets **with** as part of the WITH CHECK OPTION syntax:

```
EXEC SQL select * from mytab with; -- fails
```

The workaround in the following example uses the keyword AS to identify **with** as a table alias:

```
EXEC SQL select * from mytab as with; -- succeeds
```

The next two examples use the keyword CREATE as a table alias. Using **create** as a table alias causes the next example to fail because the database server interprets the keyword as part of the syntax to create a new database object, such as a table, synonym, or view:

```
EXEC SQL select * from mytab create; -- fails
```

```
EXEC SQL select * from mytab as create; -- succeeds
```

The workaround uses the keyword AS to identify **create** as a table alias. (Using **grant** as an alias would similarly fail, but is valid after the AS keyword.)

Fetching Cursors that have Keywords as Names

In a few situations, no workaround exists for the syntactic ambiguity that occurs when a keyword is used as an identifier in an SQL program.

In the following example, the FETCH statement specifies a cursor named **next**. The FETCH statement generates a syntax error because the preprocessor interprets **next** as a keyword, signifying the next row in the active set and expects a cursor name to follow **next**. This occurs whenever the keyword NEXT, PREVIOUS, PRIOR, FIRST, LAST, CURRENT, RELATIVE, or ABSOLUTE is used as a cursor name:

```
/* This code fragment fails */  
EXEC SQL declare next cursor for  
    select customer_num, lname from customer;  
EXEC SQL open next;  
EXEC SQL fetch next into :cnum, :lname;
```

Fetching Cursors that have Keywords as Names

If you use any of the following keywords as identifiers for variables in a user-defined routine (UDR), you can create ambiguous syntax:

- CURRENT
- DATETIME
- GLOBAL
- INTERVAL
- NULL
- OFF
- OUT
- PROCEDURE
- SELECT
- SYSDATE

Using CURRENT, DATETIME, INTERVAL, and NULL in INSERT

A UDR cannot insert a variable that was declared using the CURRENT, DATETIME, INTERVAL, or NULL keywords as the name. For example, if you declare a variable called **null**, when you try to insert the value **null** into a column, you receive a syntax error, as the following example shows:

```
CREATE PROCEDURE problem()
. . .
DEFINE null INT;
LET null = 3;
INSERT INTO tab VALUES (null); -- error, inserts NULL, not 3
```

Using NULL and SELECT in a Condition

If you declare a variable with the name **null** or **select**, including it in a condition that uses the IN keyword is ambiguous. The following example shows three conditions that cause problems: in an IF statement, in a WHERE clause of a SELECT statement, and in a WHILE condition:

```
CREATE PROCEDURE problem()
. . .
DEFINE x,y,select, null, INT;
DEFINE pname CHAR[15];
LET x = 3; LET select = 300;
LET null = 1;
IF x IN (select, 10, 12) THEN LET y = 1; -- problem if

IF x IN (1, 2, 4) THEN
SELECT customer_num, fname INTO y, pname FROM customer
  WHERE customer IN (select , 301 , 302, 303); -- problem in

WHILE x IN (null, 2)    -- problem while
. . .
END WHILE;
```

You can use the variable **select** in an IN list if you ensure it is not the first element in the list. The workaround in the following example corrects the IF statement that the preceding example shows:

```
IF x IN (10, select, 12) THEN LET y = 1; -- problem if
```

No workaround exists to using **null** as a variable name and attempting to use that variable in an IN condition.

Declaring Keywords or Routine Names as SPL Variables

If you declare a variable with the same name as a keyword or as the name of a routine, ambiguities can occur. Dynamic Server uses the following rules for resolving name conflicts among SPL variables, UDR names, and built-in SQL function names.

- Variable names that are declared in DEFINE statements take the highest precedence.
- User-defined routines defined in CREATE PROCEDURE or CREATE FUNCTION statements take precedence over built-in SQL functions.
- Procedures declared with the PROCEDURE keyword in the DEFINE statement take precedence over built-in SQL functions.
- Built-in SQL functions take precedence over SPL procedures that exist in the database but that are not explicitly identified as procedures in the DEFINE statement.

Do not use the name of a built-in SQL function as an SPL variable if you might need to invoke the SQL function. For example, do not declare a variable with the name **count** or **max**, if you might also need to call those aggregate functions.

Variables that Conflict with Column Names

If you use the same identifier for an SPL variable and a column name, then within the scope of reference of the variable, the database server interprets any instance of the unqualified identifier as a variable. To use the identifier to specify a column name, use *table.column* notation to qualify the column name with the table name. In the following example, the procedure variable **lname** is the same as the column name. In the following SELECT statement, **customer.lname** is a column in the database and **lname** is an SPL variable:

```
CREATE PROCEDURE table_test()
DEFINE lname CHAR(15);
LET lname = "Miller";
SELECT customer.lname FROM customer INTO lname
WHERE customer_num = 502;
```

This example is valid, but relying on the rules of precedence of Dynamic Server to resolve name conflicts between SPL variables and column names might make your code difficult for human readers to interpret and to maintain. An alternative to reusing the same identifier as a variable and as a column name is for the DEFINE statement to declare some prefix to the identifier, such as **v_lname** in this example, to indicate that this variable stores the value of the column **lname**.

Using ON, OFF, or PROCEDURE with TRACE

If you define an SPL variable called **on**, **off**, or **procedure**, and you attempt to use it in a TRACE statement, the value of the variable is not traced. Instead, the TRACE ON, TRACE OFF, or TRACE PROCEDURE statements execute. You can trace the value of the variable by specifying the variable in a more complex expression.

The following example shows both the ambiguous syntax and workarounds that use arithmetic or string expressions that evaluate to the variable:

```
DEFINE on, off, procedure INT;

TRACE on; --ambiguous
TRACE 0+ on; --ok
TRACE off; --ambiguous
TRACE '||off;--ok

TRACE procedure; --ambiguous
TRACE 0+procedure;--ok
```

Using GLOBAL as the Name of a Variable

If you attempt to define a variable with the name **global**, the define operation fails. The syntax that the following example shows conflicts with the syntax for defining global variables:

```
DEFINE global INT; -- fails;
```

If the **DELIMIDENT** environment variable is set, you could use **global** as a variable name, as the following example shows:

```
DEFINE "global" INT; -- successful
```

Important: Although workarounds that the preceding sections show can avoid compilation or runtime syntax conflicts from keywords used as identifiers, keep in mind that such identifiers tend to make code more difficult to understand and to maintain.

Using EXECUTE, SELECT, or WITH as Cursor Names

Do not use an EXECUTE, SELECT, or WITH keyword as the name of a cursor. If you try to use one of these keywords as the name of a cursor in a FOREACH statement, the cursor name is interpreted as a keyword in the FOREACH statement. No workaround exists.

The following example does not work:

```
DEFINE execute INT;
FOREACH execute FOR SELECT col1 -- error, looks to parser like
    INTO var1 FROM tab1;      -- 'FOREACH EXECUTE PROCEDURE'
```

SELECT Statements in WHILE and FOR Statements

If you use a SELECT statement in a WHILE or FOR loop, and if you need to enclose it in parentheses, enclose the entire SELECT statement in a BEGIN...END statement block. The SELECT statement in the first WHILE statement in the following example is interpreted as a call to the procedure `var1`; the second WHILE statement is interpreted correctly:

```
DEFINE var1, var2 INT;
WHILE var2 = var1
    SELECT col1 INTO var3 FROM TAB -- error, interpreted as call var1()
    UNION
    SELECT co2 FROM tab2;
END WHILE;

WHILE var2 = var1
    BEGIN
        SELECT col1 INTO var3 FROM TAB -- ok syntax
        UNION
        SELECT co2 FROM tab2;
    END
END WHILE;
```

SET Keyword in the ON EXCEPTION Statement

If you use a statement that begins with the keyword SET in ON EXCEPTION, you must enclose it in a BEGIN ... END statement block. The following list shows some of the SQL statements that begin with the keyword SET:

- SET ALL_MUTABLES
- SET AUTOFREE
- SET CONNECTION
- SET CONSTRAINTS
- SET DATASKIP
- SET DEBUG FILE
- SET Default Table Space
- SET Default Table Type
- SET DEFERRED_PREPARE
- SET DESCRIPTOR
- SET ENCRYPTION
- SET ENVIRONMENT
- SET EXPLAIN
- SET INDEX
- SET INDEXES
- SET ISOLATION

- SET LOCK MODE
- SET LOG
- SET OPTIMIZATION
- SET PDQPRIORITY
- SET PLOAD FILE
- SET ROLE
- SET SCHEDULE LEVEL
- SET SESSION AUTHORIZATION
- SET STATEMENT CACHE
- SET TABLE
- SET TRANSACTION
- SET TRIGGERS

The following examples show the incorrect and correct use of a SET LOCK MODE statement inside an ON EXCEPTION statement.

The following ON EXCEPTION statement returns an error because the SET LOCK MODE statement is not enclosed in a BEGIN ... END statement block:

```
ON EXCEPTION IN (-107)
  SET LOCK MODE TO WAIT; -- error, value expected, not 'lock'
END EXCEPTION;
```

The following ON EXCEPTION statement executes successfully because the SET LOCK MODE statement is enclosed in a BEGIN ... END statement block:

```
ON EXCEPTION IN (-107)
  BEGIN
    SET LOCK MODE TO WAIT; -- ok
  END
END EXCEPTION;
```

Related Information

For a discussion of owner naming, see your *IBM Informix Performance Guide*.

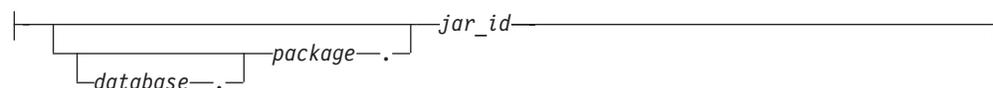
For a discussion of identifiers that support non-ASCII characters and a discussion of non-ASCII characters in delimited identifiers, see the *IBM Informix GLS User's Guide*.

Jar Name

Use the Jar Name segment to specify the name of a jar ID. Use this segment whenever you see a reference to Jar Name in a syntax diagram. Only Dynamic Server supports this syntax segment.

Syntax

Jar Name:



Element	Description	Restrictions	Syntax
<i>database</i>	Database in which to install or access <i>jar_id</i> . Default is the current database.	Fully qualified <i>database.package.jar_id</i> identifier must not exceed 255 bytes	"Database Name" on page 5-15
<i>jar_id</i>	The .jar file that contains the Java class to be accessed	File must exist in <i>database.package</i>	"Identifier " on page 5-21
<i>package</i>	Name of the package	Package must exist in <i>database</i>	"Identifier " on page 5-21

If a jar name is specified as a character string argument to the **sqlj.install_jar**, **sqlj.replace_jar**, or **sqlj.remove_jar** procedures, then any identifiers in the jar name that are delimited identifiers will include the surrounding double quotation mark characters.

Before you can access a *jar_id* in any way (including its use in a CREATE FUNCTION or CREATE PROCEDURE statement), it must be defined in the current database with the **install_jar()** procedure. For more information, see "EXECUTE PROCEDURE " on page 2-364.

Related Information

For information on how to update the three-part names of jar files after you rename the database, see the *J/Foundation Developer's Guide*.

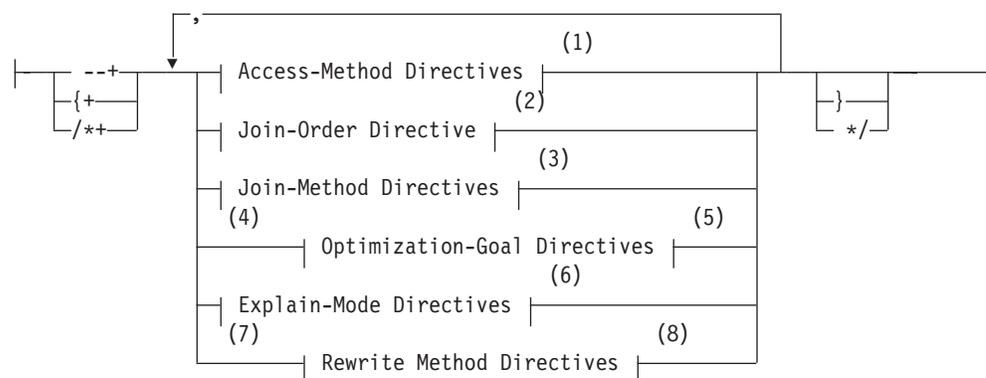
For descriptions of the **sqlj.install_jar**, **sqlj.replace_jar**, or **sqlj.remove_jar** procedures, see "SQLJ Driver Built-In Procedures" on page 6-57.

Optimizer Directives

The Optimizer Directives segment specifies keywords that you can use to partially or fully specify the query plan of the optimizer. Use this segment whenever you see a reference to Optimizer Directives in a syntax diagram.

Syntax

Optimizer Directives:



Notes:

- 1 See " Access-Method Directives" on page 5-37
- 2 See " Join-Order Directive" on page 5-40

- 3 See "Join-Method Directives " on page 5-40
- 4 Dynamic Server only
- 5 See " Optimization-Goal Directives (IDS)" on page 5-42
- 6 See "Explain-Mode Directives" on page 5-43
- 7 Extended Parallel Server only
- 8 See "Rewrite Method Directive (XPS)" on page 5-44

Usage

Use one or more optimizer directives to partially or fully specify the query plan of the optimizer. The scope of the directive is the current query only.

Directives are enabled by default. To obtain information about how specified directives are processed, view the output of the SET EXPLAIN statement. To disable directives, set the **IFX_DIRECTIVES** environment variable to 0, or set the **DIRECTIVES** parameter in the ONCONFIG file to 0.

Optimizer Directives as Comments

Optimizer directives require valid comment indicators as delimiters.

- If { is the opening delimiter, you must use } as the closing delimiter.
- If /* are the opening delimiters, you must use */ as the closing delimiters.
- If -- are the opening delimiters, then no closing delimiter is needed.

An optimizer directive or a list of optimizer directives immediately follows the DELETE, SELECT, UPDATE, or (for XPS) MERGE keyword in the form of a comment. After the comment symbol, the first character in an optimizer directive is always a plus (+) sign. No blank space or other white space character is allowed between the comment indicator and the plus sign.

You can use any of the following comment indicators:

- A double hyphen (--) delimiter
The double hyphen needs no closing symbol because it specifies only the remainder of the current line as comment. When you use this style, include the optimizer directive on only the current line.
- Braces ({ . . . }) delimiters
The comment extends from the left brace ({) until the next right (}) brace; this can be in the same line or in some subsequent line.
- C-language style slash and asterisk (/* . . . */) delimiters
The comment extends from the initial slash-asterisk (/*) pair until the next asterisk-slash (*/) characters in the same line or in some subsequent line.
In Informix ESQL/C, the **-keep** command option to the **esql** compiler must be specified when you use C-style comments.

For additional information, see "How to Enter SQL Comments" on page 1-3.

If you specify multiple directives in the same query, you must separate them with a blank space, a comma, or by any character that you choose. It is recommended that you separate successive directives with a comma.

If the query declares an alias for a table, use the alias (rather than the actual table name) in the optimizer directive specification. Because system-generated index names begin with a blank character, use quotation marks to delimit such names.

Syntax errors in an optimizer directive do not cause a valid query to fail. You can use the SET EXPLAIN statement to obtain information related to such errors.

In distributed queries, optimizer directives can reference objects in other databases of the same server instance by using the *database:table* or *database:owner.table* notation to qualify the name of a table in another database of the local database server.

Restrictions on Optimizer Directives

You can specify optimizer directives for any query in a DELETE, SELECT, or UPDATE statement, unless it includes any of the following syntax elements:

- A query accessing a table in a database of a remote database server instance
- In Informix ESQL/C, a statement with the WHERE CURRENT OF *cursor* clause

For queries that use ANSI/ISO-compliant syntax to specify a join, the query optimizer does not follow some directives:

- The join-method directives (USE_NL, AVOID_NL, USE_HASH, AVOID_HASH, /BUILD, and /PROBE) are ignored, except in cases where the optimizer rewrites the query so that it is no longer uses the ANSI/ISO syntax.
- The join-order directive (ORDERED) is ignored in ANSI-compliant joined queries that specify the RIGHT OUTER JOIN or FULL OUTER JOIN keywords.

Access-Method Directives

Use the access-method directives to specify the manner in which the optimizer should search the tables.

Access-Method Directives:

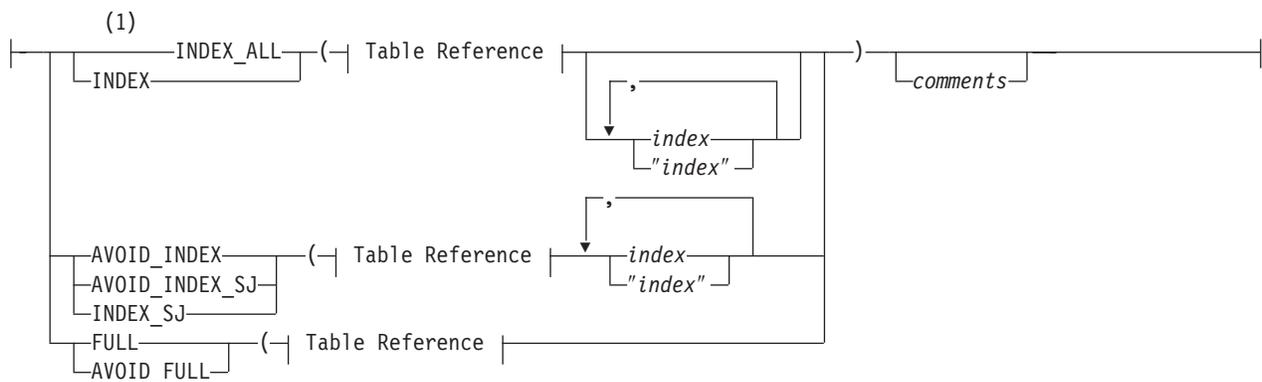


Table Reference:



Notes:

1 Extended Parallel Server only

Element	Description	Restrictions	Syntax
<i>alias</i>	Temporary alternative table name declared in the FROM clause	If an <i>alias</i> is declared, it must be used in the optimizer directive	"Identifier " on page 5-21
<i>comments</i>	Text that documents the optimizer directive	Must be outside the parenthesis but inside the comment symbols	Character string
<i>index</i>	Index for which to specify a query plan directive	Must exist. With AVOID_INDEX, at least one <i>index</i> is required	"Identifier " on page 5-21
<i>synonym, table</i>	Synonym or table in a query for which to specify a directive	Synonym and the table to which it points must exist	"Identifier " on page 5-21

Use commas or blank spaces to separate elements within the parentheses.

The following table describes each of the access-method directives and indicates how it affects the query plan of the optimizer.

Keywords	Effect	Optimizer Action
AVOID_FULL	No full-table scan on the listed table	The optimizer considers the various indexes it can scan. If no index exists, the optimizer performs a full-table scan.
AVOID_INDEX	Does not use any of the indexes listed	The optimizer considers the remaining indexes and a full-table scan. If all indexes for a table are specified, optimizer uses a full-table scan to access the table.
AVOID_INDEX_SJ	Does not use an index self-join path for the specified indexes	The optimizer does not consider the specified index for scanning the table in an index self-join path.
FULL	Performs a full-table scan	Even if an index exists on a column, the optimizer uses a full-table scan to access the table.
INDEX	Uses the index specified to access the table	If more than one index is specified, the optimizer chooses the index that yields the least cost. If no indexes are specified, then all the available indexes are considered.
INDEX_ALL	Access the table using all listed indexes (Multi-index scan)	(Extended Parallel Server only) If more than one index is specified, all are used. If only one is specified, optimizer follows an index self-join path using that index. If no index is specified, all available indexes are considered.
INDEX_SJ	Use the specified index to scan the table in an index self-join path.	The optimizer is forced to scan the table using an index self-join path with the specified index (or to choose the least costly index in a comma-separated list of indexes for an index self-join path).

Both the AVOID_FULL and INDEX keywords specify that the optimizer should avoid a full scan of a table. It is recommended, however, that you use the AVOID_FULL keyword to specify the intent to avoid a full scan on the table.

When you specify both the AVOID_FULL and AVOID_INDEX access-method directives, the optimizer avoids performing a full scan of the table and it avoids using the specified index or indexes. This combination of negative directives allows the optimizer to use indexes that are created after the access-method directives are specified.

Because the optimizer automatically considers the index self-join path if you specify the INDEX or AVOID_FULL directive, use the INDEX_SJ directive only to

force an index self-join path using the specified index (or choosing the least costly index in a comma-separated list of indexes). The INDEX_SJ directive can improve performance when a multicolumn index includes columns that provide only low selectivity as index key filters.

Specifying the INDEX_SJ directive circumvents the usual optimizer requirement for data distribution statistics on the lead keys of the index. This directive causes the optimizer to consider an index self-join path, even if data distribution statistics are not available for the leading index key columns. In this case, the optimizer only includes the minimum number of index key columns as lead keys to satisfy the directive.

For example, if an index is defined on columns c1, c2, c3, c4, and the query specifies filters on all four of these columns but no data distributions are available on any column, then specifying INDEX_SJ on this index will result in column c1 being used as the lead key in an index self-join path. If you want the optimizer to use an index but not to consider the index self-join path, then you must specify an INDEX or AVOID_FULL directive to choose the index, and you must also specify an AVOID_INDEX_SJ directive to prevent the optimizer from considering any other index self-join path.

In general, you can specify only one access-method directive per table. Only the following combinations of access method directives are valid for the same table in the same query:

- AVOID_FULL, AVOID_INDEX
- AVOID_FULL, AVOID_INDEX_SJ
- AVOID_INDEX, AVOID_INDEX_SJ
- AVOID_FULL, AVOID_INDEX, AVOID_INDEX_SJ
- INDEX, AVOID_INDEX_SJ

If AVOID_INDEX_SJ is used together with the INDEX directive, either as an explicit INDEX directive or as the equivalent AVOID_FULL and AVOID_INDEX combination, the indexes specified in the AVOID_INDEX_SJ directive must be a subset of the indexes specified in the INDEX directive. For more information about the effects of the INDEX_SJ and AVOID_INDEX_SJ directives, see the chapter of the *IBM Informix Performance Guide* that describes optimizer directives.

Suppose that you have a table named **emp** that contains the following indexes: **loc_no**, **dept_no**, and **job_no**. When you perform a SELECT that includes the **emp** table in the FROM clause, you might direct the optimizer to access the table in one of the following ways:

- Example using a positive directive:
SELECT {+INDEX(emp idx_dept_no)} ...

In this example the access-method directive forces the optimizer to scan the **idx_dept_no** index on the **dept_no** column.

- Example using a negative directive:
SELECT {+AVOID_INDEX(emp idx_loc_no, idx_job_no), AVOID_FULL(emp)} ...

This example includes multiple access-method directives. These access-method directives force a scan of the **idx_dept_no** index on the **dept_no** column by instructing the optimizer not to scan the **idx_loc_no** and **idx_job_no** indexes, and not to perform a full scan of the **emp** table. If a new **idx_emp_no** index, however, is created for table **emp**, these directives do not prevent the optimizer from considering it.

Note also that the term *negative directive* refers to the string "AVOID_" in an access method directive, and has nothing to do with the + symbol following the comment indicator that begins every optimizer directive.

Join-Order Directive

Use the ORDERED join-order directive to force the optimizer to join tables or views in the order in which they appear in the FROM clause of the query.

Join-Order Directive:



Element	Description	Restrictions	Syntax
<i>comments</i>	Text to document the directive	Must appear between comment symbols	Character string

For example, the following query forces the database server to join the **dept** and **job** tables and then join the result with the **emp** table:

```
SELECT --+ ORDERED
       name, title, salary, dname
FROM dept, job, emp WHERE title = 'clerk' AND loc = 'Palo Alto'
AND emp.dno = dept.dno
AND emp.job= job.job;
```

Because no predicates occur between the **dept** table and the **job** table, this query forces the database server to construct a Cartesian product.

When your query involves a view, the placement of the ORDERED join-order directive determines whether you are specifying a partial- or total-join order.

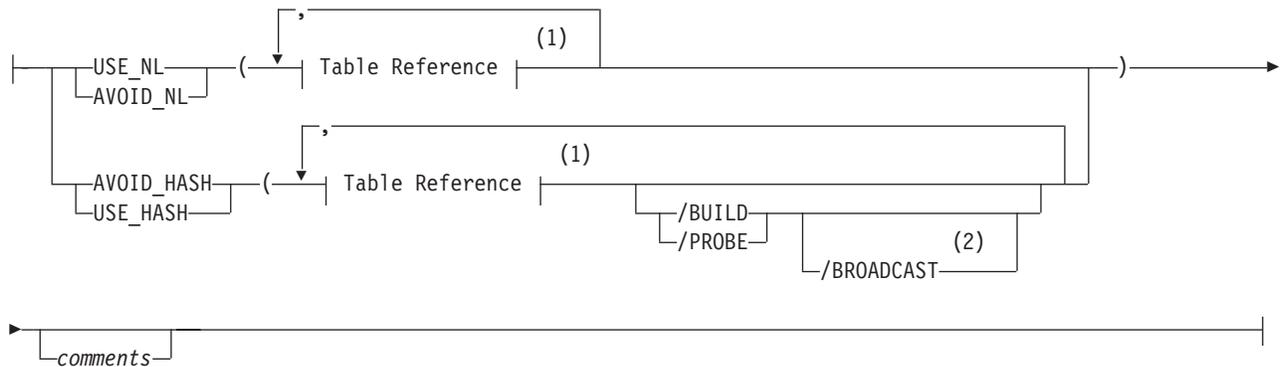
- Specifying partial-join order when you create a view
If you use the ORDERED directive when you create a view, the base tables are joined contiguously in the order of the view definition.
For all subsequent queries on the view, the database server joins the base tables contiguously in the order specified in the view definition. When used in a view, the ORDERED directive does not affect the join order of other tables named in the FROM clause in a query.
- Specifying total-join order when you query a view
When you specify the ORDERED join-order directive in a query that uses a view, all tables are joined in the order specified, even those tables that form views. If a view is included in the query, the base tables are joined contiguously in the order of the view definition. For examples of ORDERED with views, refer to your *IBM Informix Performance Guide*.

Because of ordering requirements for OUTER joins, in ANSI-compliant joined queries that specify the RIGHT OUTER JOIN or FULL OUTER JOIN keywords, the ORDERED join-order directive is ignored, but it is listed under *Directives Not Followed* in the **sqexplain.out** file.

Join-Method Directives

Use join-method directives to influence how tables are joined in an Informix-extension joined query.

Join-Method Directives:



Notes:

- 1 See " Access-Method Directives" on page 5-37
- 2 Extended Parallel Server only

Element	Description	Restrictions	Syntax
<i>comments</i>	Text to documents the directive	Must appear between comment symbols	Character string

This diagram is simplified: /BROADCAST is not valid with AVOID_HASH.

Use commas or blank spaces to separate the elements within the parentheses.

The following table describes each of the join-method directives.

Keyword

Effect

USE_NL

Uses the specified tables as the inner table in a nested-loop join

If n tables are specified in the FROM clause, then at most $(n-1)$ tables can be specified in the USE_NL join-method directive.

USE_HASH

Uses a hash join to access the specified table

You can also choose whether the table will be used to create the hash table or to probe the hash table.

AVOID_NL

Does not use the specified table as inner table in a nested loop join

A table listed with this directive can still participate in a nested loop join as the outer table.

AVOID_HASH

Does not access the specified table using a hash join

You can optionally use a hash join, but impose restrictions on the role of the table within the hash join.

A join-method directive takes precedence over the join method forced by the OPTCOMPIND configuration parameter.

When you specify the USE_HASH or AVOID_HASH directives (to use or avoid a hash join, respectively), you can also specify the role of each table:

- /BUILD

With the USE_HASH directive, this keyword indicates that the specified table be used to construct a hash table. With the AVOID_HASH directive, this keyword indicates that the specified table *not* be used to construct a hash table.

- /PROBE

With the USE_HASH directive, this keyword indicates that the specified table be used to probe the hash table. With the AVOID_HASH directive, this keyword indicates that the specified table *not* be used to probe the hash table. You can specify multiple probe tables as long as there is at least one table for which you do not specify PROBE.

- /BROADCAST

Only Extended Parallel Server supports the /BROADCAST directive, which can improve performance when small base tables or derived tables with few records (< 100K in data size) are involved in a join. In such cases, all the records of the small tables are sent to each instance of a given join operation. The BROADCAST directive must also include either /BUILD or /PROBE, but it is not valid with AVOID_HASH. If the /BUILD option is included, the USE_HASH table is broadcast. If the /PROBE option is included, the join in which the USE_HASH table participates is broadcast.

For the optimizer to find an efficient join query plan, you must at least run UPDATE STATISTICS LOW for every table that is involved in the join, so as to provide appropriate cost estimates. Otherwise, the optimizer might choose to broadcast the entire table to all instances, even if the table is large.

If neither the /BUILD nor the /PROBE keyword is specified, the optimizer uses cost estimates to determine the role of the table.

In this example, the USE_HASH directive forces the optimizer to construct a hash table on the **dept** table and consider only the hash table to join **dept** with the other tables. Because no other directives are specified, the optimizer can choose the least expensive join methods for the other joins in the query.

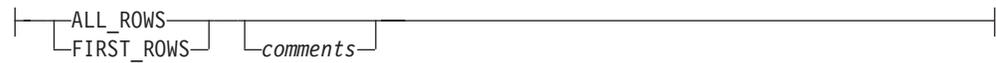
```
SELECT /*+ USE_HASH (dept /BUILD)
      The optimizer must use dept to construct a hash table */
      name, title, salary, dname
FROM emp, dept, job WHERE loc = 'Phoenix'
      AND emp.dno = dept.dno AND emp.job = job.job;
```

Join-method optimizer directives that you specify for an ANSI-compliant joined query are ignored, but they are listed under *Directives Not Followed* in the **sqexplain.out** file.

Optimization-Goal Directives (IDS)

Use optimization-goal directives to specify the measure that is used to determine the performance of a query result.

Optimization-Goal Directives:



Element	Description	Restrictions	Syntax
<i>comments</i>	Text documenting the directive	Must appear between comment symbols	Character string

The two optimization-goal directives are:

- **FIRST_ROWS**

This tells the optimizer to choose a plan that optimizes the process of finding only the first screenful of rows that satisfies the query. Use this option to decrease initial response time for queries that use an interactive mode or that require the return of only a few rows.

- **ALL_ROWS**

This directive tells the optimizer to choose a plan that optimizes the process of finding all rows that satisfy the query.

This form of optimization is the default.

An optimization-goal directive takes precedence over the **OPT_GOAL** environment variable setting and over the **OPT_GOAL** configuration parameter.

For information about how to set the optimization goal for an entire session, see the **SET OPTIMIZATION** statement.

You cannot use an optimization-goal directive in the following contexts:

- In a view definition
- In a subquery

The following query returns the names of the employees who earned the top fifty bonuses. The optimization-goal directive directs the optimizer to return the first screenful of rows as fast as possible.

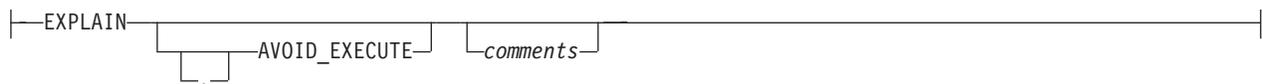
```

SELECT {+FIRST_ROWS
      Return the first screenful of rows as fast as possible}
LIMIT 50 fname, lname FROM employees ORDER BY bonus DESC;
  
```

Explain-Mode Directives

Use the explain-mode directives to test and debug query plans and to print information about the query plan to the **sqexplain.out** file.

Explain-Mode Directives:



Element	Description	Restrictions	Syntax
<i>comments</i>	Text documenting the directive	Must appear between comment symbols	Character string

The following table lists the effect of each explain-mode directive.

Keyword

Effect

EXPLAIN

Turns SET EXPLAIN ON for the specified query

AVOID_EXECUTE

Prevents the data manipulation statement from executing; instead, the query plan is printed to the **sqexplain.out** file

The EXPLAIN directive is primarily useful for testing and debugging query plans. It is redundant when SET EXPLAIN ON is already in effect. It is not valid in a view definition or in a subquery.

The next query executes and prints the query plan to the **sqexplain.out** file:

```
SELECT {+EXPLAIN}
      c.customer_num, c.lname, o.order_date
FROM customer c, orders o WHERE c.customer_num = o.customer_num;
```

The AVOID_EXECUTE directive prevents execution of a query on either the local or remote site, if a remote table is part of the query. This directive does not prevent nonvariant functions in a query from being evaluated.

The next query does returns no data, but writes its query plan to file **sqexplain.out**:

```
SELECT {+EXPLAIN, AVOID_EXECUTE} c.customer_num, c.lname, o.order_date
      FROM customer c, orders o WHERE c.customer_num = o.customer_num;
```

You must use both the EXPLAIN and AVOID_EXECUTE directives to see the query plan of the optimizer (in the **sqexplain.out** file) without executing the query. The comma (,) separating these two directives is optional.

If you omit the EXPLAIN directive when you specify the AVOID_EXECUTE directive, no error is issued, but no query plan is written to the **sqexplain.out** file and no DML statement is executed.

You cannot use the explain-mode directives in the following contexts:

- In a view definition
- In a trigger
- In a subquery

They are valid, however, in a SELECT statement within an INSERT statement.

Rewrite Method Directive (XPS)

By default, the database server attempts to unnest correlated subqueries. You can use the NESTED rewrite directive in the outermost query to prevent the query plan from being rewritten with unnested subqueries.

Rewrite Method Directive:

```
|---NESTED---|-----|
      |-----|
      |---comments---|
```

Element	Description	Restrictions	Syntax
<i>comments</i>	Text documenting the directive	Must appear between comment symbols	Character string

External Directives

You can use the `SAVE EXTERNAL DIRECTIVES` statement to store optimizer directives in the `sysdirectives` table of the system catalog. Dynamic Server applies these external directives automatically to subsequent queries and subqueries that match a specified `SELECT` statement.

The `EXT_DIRECTIVES` configuration parameter and the `IFX_EXTDIRECTIVES` environment variable can be set to control whether external directives are enabled or disabled for the database server instance or for the session. Setting either of these to zero disables external directives; setting both to 1 enables external directives.

You can also use the `EXTDIRECTIVES` option of the `SET ENVIRONMENT` statement to enable or disable external directives during a session. For more information, see “Enabling or Disabling External Directives for a Session” on page 2-523.

Related Information

For information about the `sqexplain.out` file, see `SET EXPLAIN`.

For information about how to create optimizer directives that the database server stores in the `sysdirectives` system catalog table for use with subsequent queries, see `SAVE EXTERNAL DIRECTIVES`.

For information about how to set optimization settings for an entire session, see the description of `SET OPTIMIZATION`.

For a discussion about optimizer directives and performance, see your *IBM Informix Performance Guide*.

For descriptions of the `IFX_DIRECTIVES` and `IFX_EXTDIRECTIVES` environment variables, see the *IBM Informix Guide to SQL: Reference*.

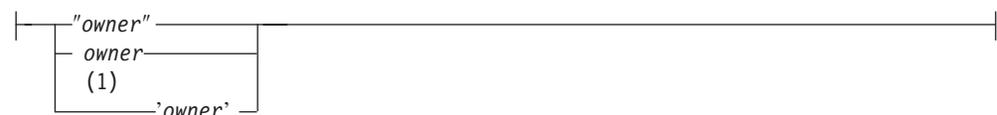
For information on the `DIRECTIVES` and `EXT_DIRECTIVES` parameters in the `onconfig` file, see your *IBM Informix Administrator's Reference*.

Owner Name

The owner name specifies the owner of a database object. Use this segment whenever you see a reference to Owner Name in a syntax diagram.

Syntax

Owner Name:



Notes:

1 Informix extension

Element	Description	Restrictions	Syntax
<i>owner</i>	User name of the owner of an object in a database	Maximum length is 32 bytes	Must conform to the rules of your operating system.

Usage

In an ANSI-compliant database, you must specify the *owner* of any database object that you do not own. In reference to the owner of database objects, the ANSI/ISO synonym for owner name is *authorization identifier*. (In reference to schema objects, however, the ANSI/ISO term for what the Dynamic Server documentation calls an *owner name* is *schema name*.)

In databases that are not ANSI-compliant, the *owner* name is optional. You do not need to specify *owner* when you create database objects or use data access statements. If you do not specify *owner* when you create a database object, the database server assigns your login name as the owner of the object, in most cases. For exceptions to this rule, see “Ownership of Created Database Objects” on page 2-121 in CREATE FUNCTION and “Ownership of Created Database Objects” on page 2-162 in CREATE PROCEDURE. When a DDL statement in an owner-privileged UDR creates a new database object, the owner of the routine (rather than the user who executes it, if that user is not the owner of the routine) becomes the owner of the new database object.

If you specify *owner* in data-access statements, the database server checks it for correctness. Without quotation marks, *owner* is case insensitive. The following four queries all can access data from the table **kaths.tab1**:

```
SELECT * FROM tab1;  
SELECT * FROM kaths.tab1;  
SELECT * FROM KATHS.tab1;  
SELECT * FROM KathS.tab1;
```

In an ANSI-compliant database, only the owner of the table, user **kaths**, can issue the first of these example queries, which specifies an unqualified table name, but any user who holds the Select privilege on **tab1** can issue that query in a database that is not ANSI-compliant. For more information about owner names in ANSI-compliant databases, see “ANSI-Compliant Database Restrictions and Case Sensitivity” on page 5-48.

A *role* that the CREATE ROLE statement declares is an authorization identifier, and is therefore subject to the syntax restrictions on owner names, but a role cannot be the owner of a database object. Similarly, the keyword PUBLIC, which specifies the group of all users, cannot be the owner of a database object, except in the special cases of the **sysdbopen()** and **sysdbclose()** procedures. For more information about these built-in session configuration UDRs, see “IFX_REPLACE_MODULE Function” on page 6-54.

Using Quotation Marks

Within quotation marks, *owner* is case sensitive. Quotation marks instruct the database server to read or store the name exactly as typed when you create or

access a database object. For example, suppose that you have a table whose owner is Sam. You can use either one of the following two statements to access data in the table:

```
SELECT * FROM table1;
SELECT * FROM 'Sam'.table1;
```

The first query succeeds because the owner name is not required. The second query succeeds because the specified owner name matches the owner name as it is stored in the database.

Referencing Tables Owned by User **informix**

If you use the owner name as one of the selection criteria to access database object information from one of the system catalog tables, the owner name is case sensitive. To preserve lettercase, you must enclose *owner* in single or double quotation marks, and you must type the owner name exactly as it is stored in the system catalog table. Of the following two examples, only the second successfully accesses information on the table **Kaths.table1**.

```
SELECT * FROM systables WHERE tabname = 'tab1' AND owner = 'kaths';
SELECT * FROM systables WHERE tabname = 'tab1' AND owner = 'Kaths';
```

User **informix** is the owner of the system catalog tables, and in an ANSI-compliant database you must specify **informix** as a qualifier when SQL statements reference system catalog tables, unless you are user **informix**:

```
SELECT * FROM "informix".systables WHERE tabname = 'tab1' AND owner = 'Kaths';
```

Dynamic Server accepts any of the following notations to specify a system catalog table of an ANSI-compliant database:

- "informix".*system_table*
- *informix.system_table*
- 'informix'.*system_table*

Of these three formats, however, only the first, where the *owner* is specified as a delimited identifier, is directly interoperable with most other database servers. For the format with no delimiters, the ANSI/ISO standard for SQL upshifts the lowercase letters to INFORMIX, and the same standard does not support single (') quotation marks as valid delimiters for owner names or for schema names.

In contrast, Dynamic Server treats the name **informix** as a special case, and preserves lowercase letters when **informix** is specified, with or without delimiters, whether or not the database is ANSI-compliant. To write SQL code that is portable to non-Informix database servers, however, you should always delimit the owner names of database objects between double (") quotation marks.

The following SQL examples use undelimited owner names:

```
CREATE TABLE informix.t1(i SERIAL NOT NULL);
CREATE TABLE someone.t1(i SERIAL NOT NULL);
```

If these statements execute successfully, the first table has **informix** registered in **systables** as the owner, and the second has **SOMEONE** registered as the owner. When the owner name is delimited by quotation marks in SQL statements, the specified lettercase of *owner* is preserved, but the lettercase does not matter when the owner name is undelimited, because Dynamic Server upshifts most undelimited owner names, but downshifts the undelimited **informix** (or **INFORMIX**) owner name to **informix**.

For example, suppose that after the previous two CREATE TABLE statements execute successfully, user **informix** issues the following statement:

```
CREATE TABLE INFORMIX.t1(i SERIAL NOT NULL);
```

This statement fails, because the combination of *owner* name and *table* name is not unique, if the previously registered table **t1** that is owned by **informix** already exists in the database.

Tip: The USER operator returns the login name of the current user exactly as it is stored on the system. If the owner name is stored differently from the login name (for example, a mixed-case owner name and an all lowercase login name), the owner = USER syntax fails.

ANSI-Compliant Database Restrictions and Case Sensitivity

The following table describes how the database server reads and stores *owner* when you create, rename, or access a database object.

Owner Name Specification	What the ANSI-Compliant Database Server Does
Omitted	Reads or stores <i>owner</i> exactly as the login name is stored in the system, but returns an error if the user is not the <i>owner</i> .
Specified without quotation marks	Reads or stores <i>owner</i> in uppercase letters
Enclosed between quotation marks	Reads or stores <i>owner</i> exactly as entered. See also "Using Quotation Marks" on page 5-46 and "Referencing Tables Owned by User informix" on page 5-47.

If you specify the owner name when you create or rename a database object in an ANSI-compliant database, you must include the owner name in data access statements. You must include the owner name when you access a database object that you do not own.

Because the database server automatically shifts *owner* to uppercase letters if not between quotation marks, case-sensitive errors can cause queries to fail. For example, if you are user **nancy** and you use the following statement, the resulting view has the name **nancy.njcust**:

```
CREATE VIEW 'nancy'.njcust AS
  SELECT fname, lname FROM customer WHERE state = 'NJ';
```

The following SELECT statement fails because it tries to match the name **NANCY.njcust** to the actual owner and table name of **nancy.njcust**:

```
SELECT * FROM nancy.njcust;
```

In a Dynamic Server distributed query, if the owner name is not between quotation marks, the remote database follows the lettercase convention of the local database. If the local database is ANSI-compliant, then the remote database processes the owner name in *uppercase*. If the local database is not ANSI compliant, then the remote database processes the owner name in *lowercase*.

Tip: When you use the owner name as one of the selection criteria in a query (for example, WHERE owner = 'kaths'), make sure that the quoted string matches the owner name exactly as it is stored in the database. If the database server cannot

find the database object or database, you might need to modify the query so that the quoted string uses uppercase letters (for example, `WHERE owner = 'KATHS'`).

Because owner name is an authorization identifier, rather than an SQL identifier, you can enclose *owner* between single-quotation marks (`'`) in SQL statements of a database where the `DELIMITIDENT` environment variable specifies support for delimited identifiers, thereby requiring double-quotation marks (`"`) around SQL identifiers.

Setting ANSIOWNER for an ANSI-Compliant Database

The default behavior of an ANSI-compliant database is to replace any lowercase letters with uppercase letters in any *owner* specification that is not enclosed in quotation marks. You can prevent this by setting the `ANSIOWNER` environment variable to 1 before the database server is initialized. This preserves whatever lettercase you use when you specify the *owner* string without quotation marks.

Default Owner Names

If you create a database object without explicitly specifying an owner name in a database that is not ANSI-compliant, your authorization identifier (as the default owner of the object) is stored in the system catalog of the database as if you had specified your authorization identifier within quotation marks (that is, preserving the lettercase).

If you create a database object without explicitly specifying an owner name in a database that is ANSI-compliant, any lowercase letters in your authorization identifier (as the default owner of the object) are stored in the system catalog of the database in uppercase characters, unless the `ANSIOWNER` environment variable was set to 1 before the database server was initialized. If `ANSIOWNER` was set to 1, however, the database stores the default owner of the object as your authorization identifier, with its lettercase preserved.

Summary of Lettercase Rules for Owner Names

To create a database object, such as a table called `mytab`, a user whose login name is `Otho` can declare the name of the new database object in any of the following ways:

1. `CREATE TABLE mytab . . .`
2. `CREATE TABLE Otho.mytab . . .`
3. `CREATE TABLE "Otho".mytab . . .`

The format in which an undelimited *owner* name (as in the second example) is stored in the `owner` column of the `systables` system catalog table is dependent on whether or not the local database is an ANSI-compliant database.

- In case 1, no *owner* name is specified. The implicit owner of the table is `Otho`, the user who created the table, and that owner name is stored in the `systables` table in the same format (`Otho`) as the user ID of the owner, independent of the ANSI-compliance status of the database.
- In case 2, an undelimited *owner* name is specified. The `systables` table stores all letters in the *owner* name in lowercase (here as `otho`) for databases that are not ANSI-compliant databases. For ANSI-compliant databases in which `ANSIOWNER` is not set to 1, `systables` table stores all *owner* name letters in uppercase (here as `OTHO`). If `ANSIOWNER` is set to 1, however, the name is stored in the same lettercase as specified in the DDL statement (here as `Otho`).

- In case 3, the delimited *owner* name is stored in the **systables** table in the same format in which it was specified (here as **Otho**), independent of the ANSI-compliance status of the database.

Note that user identifiers are case sensitive, but database object names are case insensitive. Therefore, the same user cannot own both a table **tab** and a table **TAB**.

In addition to the CREATE TABLE statement in these examples, all SQL statements and SPL statements follow these rules where a table name can be specified. For example, when using DROP TABLE, the format in which owner name appears while the statement is being processed is dependent upon the same conditions:

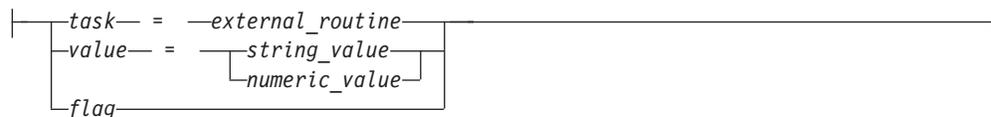
- whether an explicit owner name is specified.
- if an explicit owner name is specified, whether quotation marks delimit the owner name.
- if an explicit owner name is not delimited by quotation marks, whether or not the database is ANSI compliant.
- if the database is ANSI compliant, whether or not **ANSIOWNER** was set to 1 before the database was initialized.

Purpose Options

The CREATE ACCESS_METHOD, CREATE XADATASOURCE TYPE, and ALTER ACCESS_METHOD statements of Dynamic Server can specify purpose options for user-defined routines with the following syntax.

Syntax

Purpose Options:



Element	Description	Restrictions	Syntax
<i>external_routine</i>	User-defined routine that performs a <i>task</i>	Must be registered in the database	"Database Object Name " on page 5-16
<i>flag</i>	Keyword indicating which feature a flag enables	The interface specifies flag names	<i>Flag</i> Purpose Category in the table in " Purpose Functions, Methods, Flags, and Values" on page 5-52.
<i>numeric_value</i>	A value of a real number	Must be within the range of a numeric data type	"Literal Number " on page 4-185
<i>string_value</i>	A value that is expressed as one or more characters	Characters must be from the code set of the database	"Quoted String" on page 4-189.
<i>task</i>	Keyword that identifies a purpose function	Keywords to which you can assign a function (whose name cannot match the keyword)	<i>Task</i> Purpose Category in the table in " Purpose Functions, Methods, Flags, and Values" on page 5-52.
<i>value</i>	Keyword that identifies configuration information	Predefined configuration keywords to which you can assign values	<i>Value</i> Purpose Category in the table in " Purpose Functions, Methods, Flags, and Values" on page 5-52.

Usage

Dynamic Server supports purpose options in two contexts:

- Defining or modifying primary and secondary access methods for local or remote tables, views, and indexes
- Defining access methods for XA-compliant external data sources.

Purpose Options for Access Methods

A registered access method is a set of attributes, including a name and options called *purpose options*, that you can use to accomplish the following tasks:

- Specify which functions perform data access and manipulation tasks, such as opening, reading, and closing a data source.
- Set configuration options, such as a storage-space type.
- Set flags, such as enabling rowid interpretation.

You specify purpose options when you create an access method with the CREATE ACCESS_METHOD statement. To change the purpose options of an access method, use the ALTER ACCESS_METHOD statement.

Each *task*, *value*, or *flag* keyword corresponds to a column name in the **sysams** system catalog table. The keywords let you set the following attributes:

- Purpose function

A *purpose-function attribute* maps the name of a user-defined function or method to a *task* keyword, such as **am_create**, **am_beginscan**, or **am_getnext**. For a complete list of these keywords, see the “Task” category in the table in “Purpose Functions, Methods, Flags, and Values” on page 5-52. The *external_routine* specifies the corresponding function (C) that you supply for the access method. Example setting:

```
am_create = FS_create
```

- Purpose flag

A *purpose flag* indicates whether an access method supports a given SQL statement or keyword. Example setting:

```
am_rowids
```

- Purpose value

These string, character, or numeric values provide configuration information that a flag cannot supply. Example setting:

```
am_sptype = 'X'
```

To enable a user-defined function or method as a purpose function, you must first register the C function or Java method that performs the appropriate tasks, using the CREATE FUNCTION statement, and then set the purpose keyword equal to the registered function or method name. This creates a new access method. An example on page “Example” on page 2-6 adds a purpose method to an existing access method.

To enable a purpose flag, specify the name without a corresponding value.

To clear a purpose-option setting in the **sysams** table, use the DROP clause of the ALTER ACCESS_METHOD statement.

Purpose Functions, Methods, Flags, and Values

The following table describes the possible settings for the **sysams** columns that contain purpose functions or methods, flags, and values. The entries appear in the same order as the corresponding **sysams** columns.

Keyword	Explanation	Category	Default
am_sptype	<p>A character that specifies from what type of storage space a primary or secondary-access method can access data. The am_sptype character can have any of the following settings:</p> <ul style="list-style-type: none"> • 'X' indicates the method accesses only extspaces. • 'S' indicates the method accesses only sbspaces. • 'A' indicates the method can access extspaces and sbspaces. <p>Valid only for a new access method. You cannot change or add an am_sptype value with ALTER ACCESS_METHOD. Do not set am_sptype to 'D' or attempt to store a virtual table in a dbspace.</p>	Value	Virtual-Table Interface (C): 'A'
am_defopclass	The default operator class for a secondary-access method. The access method must exist before you can define its operator class, so you set this value in the ALTER ACCESS_METHOD statement.	Value	None
am_keyscan	A flag that, if set, indicates that am_getnext returns rows of index keys for a secondary-access method. If a query selects only the columns in the index key, the database server uses the row of index keys that the secondary-access method puts in shared memory, without reading the table.	Flag	Not set
am_unique	A flag to set if a secondary-access method checks for unique keys	Flag	Not set
am_cluster	A flag that you set if a primary- or secondary-access method supports clustering of tables	Flag	Not set
am_rowids	A flag that you set if a primary-access method can retrieve a row from a specified address	Flag	Not set
am_readwrite	<p>A flag to set if a primary-access method supports data changes. The default setting, not set, indicates that the virtual data is read-only. For the C Virtual-Table Interface, set this flag if your application will write data, to avoid the following problems:</p> <ul style="list-style-type: none"> • An INSERT, DELETE, UPDATE, or ALTER FRAGMENT statement causes an SQL error. • Function am_insert, am_delete, or am_update is not executed. 	Flag	Not set
am_parallel	<p>A flag that the database server sets to indicate which purpose functions or methods can execute in parallel in a primary or secondary-access method. If set, the hexadecimal am_parallel bitmap contains one or more of the following bit settings:</p> <ul style="list-style-type: none"> • The 1 bit is set for parallelizable scan. • The 2 bit is set for parallelizable delete. • The 4 bit is set for parallelizable update. • The 8 bit is set for parallelizable insert. <p>Insertions, deletions, and updates are not supported in the Java Virtual-Table Interface.</p>	Flag	Not set
am_costfactor	A value by which the database server multiplies the cost that the am_scancost purpose function or method returns for a primary or secondary-access method. An am_costfactor value from 0.1 to 0.9 reduces the cost to a fraction of the value that am_scancost calculates. An am_costfactor value of 1.1 or greater increases the am_scancost value.	Value	1.0

Keyword	Explanation	Category	Default
am_create	A keyword that you associate with a user-defined function or method (UDR) name that creates a virtual table or virtual index	Task	None
am_drop	A keyword that you associate with the name of a UDR that drops a virtual table or virtual index	Task	None
am_open	A keyword that you associate with the name of a UDR that makes a fragment, extspace, or sbospace available	Task	None
am_close	A keyword that you associate with the name of a UDR that reverses the initialization that am_open performs	Task	None
am_insert	A keyword that you associate with the name of a UDR that inserts a row or an index entry	Task	None
am_delete	A keyword that you associate with the name of a UDR that deletes a row or an index entry	Task	None
am_update	A keyword that you associate with the name of a UDR that changes the values in a row or key	Task	None
am_stats	A keyword that you associate with the name of a UDR that builds statistics based on the distribution of values in storage spaces	Task	None
am_scancost	A keyword that you associate with the name of a UDR that calculates the cost of qualifying and retrieving data	Task	None
am_check	A keyword that you associate with the name of a UDR that tests the physical structure of a table or performs an integrity check on an index	Task	None
am_beginscan	A keyword that you associate with the name of a UDR that sets up a scan	Task	None
am_endscan	A keyword that you associate with the name of a UDR that reverses the setup that am_beginscan initializes	Task	None
am_rescan	A keyword that you associate with the name of a UDR that scans for the next item from a previous scan to complete a join or subquery	Task	None
am_getnext	A keyword that you associate with the name of the required UDR that scans for the next item that satisfies a query	Task	None
am_getbyid	A keyword that you associate with the name of a UDR that fetches data from a specific physical address; am_getbyid is available only for primary-access methods	Task	None
am_truncate	A keyword that you associate with the name of a UDR that deletes all rows of a virtual table (primary-access method) or that deletes all corresponding keys in a virtual index (secondary-access method)	Task	None

The following rules apply to the purpose-option specifications in the CREATE ACCESS_METHOD and ALTER ACCESS_METHOD statements:

- To specify multiple purpose options in one statement, separate them with commas.
- The CREATE ACCESS_METHOD statement must specify a user-defined function or method name that corresponds to the **am_getnext** keyword.
The ALTER ACCESS_METHOD statement cannot drop the function or method name that corresponds to **am_getnext** but can modify it.
- The ALTER ACCESS_METHOD statement cannot add, drop, or modify the **am_sptype** value.
- You can specify the **am_defopclass** value only with the ALTER ACCESS_METHOD statement.

You must first register a secondary-access method with the CREATE ACCESS_METHOD statement before you can assign a default operator class.

Purpose Options for XA Data Source Types

The CREATE XADATASOURCE TYPE statement specifies purpose functions that provide access to data from external data sources that comply with the X/Open XA standards. These functions also enable external data to be processed in accordance with the transactional semantics of Dynamic Server. Only databases that use transaction logging, such as ANSI-compliant databases and Dynamic Server databases that support explicit transactions, can support transaction coordination.

The following statement creates a new XA data source type called **MQSeries**, owned by user **informix**.

```
CREATE XADATASOURCE TYPE 'informix'.MQSeries(  
    xa_flags      = 1,  
    xa_version    = 0,  
    xa_open       = informix.mqseries_open,  
    xa_close      = informix.mqseries_close,  
    xa_start      = informix.mqseries_start,  
    xa_end        = informix.mqseries_end,  
    xa_rollback   = informix.mqseries_rollback,  
    xa_prepare    = informix.mqseries_prepare,  
    xa_commit     = informix.mqseries_commit,  
    xa_recover    = informix.mqseries_recover,  
    xa_forget     = informix.mqseries_forget,  
    xa_complete   = informix.mqseries_complete);
```

These values represents the fields in the XA Switch Structure, as listed in the file **\$INFORMIXDIR/incl/public/xa.h**. The order of specifications in this example follows the order of column names in the **sysxasourcetypes** system catalog table, but they can be listed in any order, provided that no item is repeated. The **xa_flags** and **xa_version** values must be numbers; the rest must be names of UDRs that the Transaction Manager can invoke. These UDRs must already exist in the database before you can issue a CREATE XADATASOURCE TYPE statement that references them among its purpose option specifications.

The DROP FUNCTION or DROP ROUTINE statement cannot drop a UDR that is listed among the purpose options of a CREATE XADATASOURCE TYPE statement until all of the XA datasource types that were defined using the UDR are dropped.

For information about how to use the UDRs in the previous example to coordinate transactions with external XA data sources, see the *IBM Informix Datablade API Programmer's Guide*.

For information about the MQDataBlade module, see the *IBM Informix Database Extensions User's Guide*.

Related Information

Related statements: "ALTER ACCESS_METHOD " on page 2-5, "CREATE ACCESS_METHOD " on page 2-90, "CREATE OPCLASS" on page 2-149, and "CREATE XADATASOURCE TYPE " on page 2-280

For the following topics, see the *IBM Informix Virtual-Table Interface Programmer's Guide* (for C):

- Managing storage spaces, executing in parallel, and calculating statement costs
- Registering the access method and purpose functions
- Purpose-function reference

For the following topics, see the *IBM Informix Virtual-Index Interface Programmer's Guide* (for C):

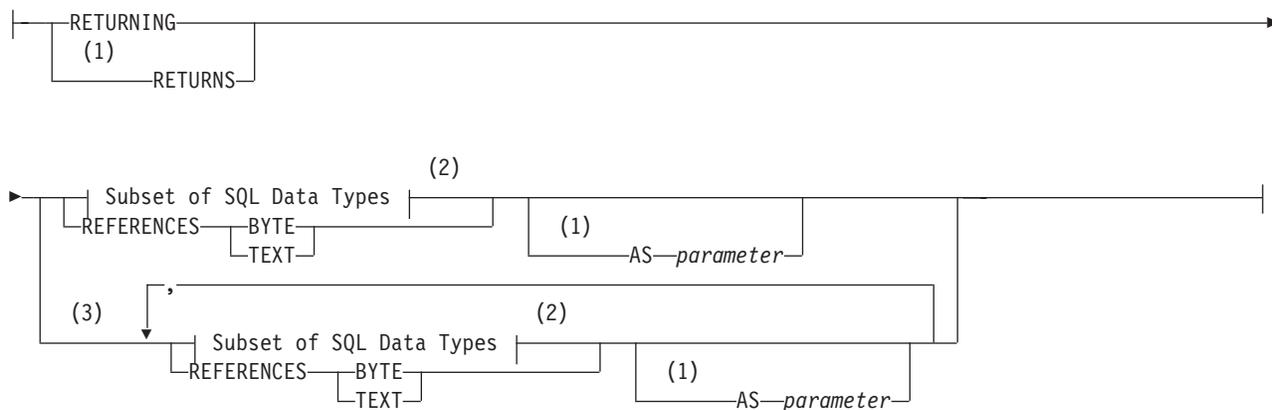
- Managing storage spaces, executing in parallel, calculating statement costs, bypassing table scans, and enforcing unique-index constraints
- Registering the access method and purpose functions
- Purpose-function reference

Return Clause

The Return clause specifies the data type of a value or values that a user-defined function returns. You can use this segment in UDR definitions.

Syntax

Return Clause:



Notes:

- 1 Dynamic Server only
- 2 See "Subset of SQL Data Types" on page 5-56
- 3 Stored Procedure Language only

Element	Description	Restrictions	Syntax
<i>parameter</i>	Name that you declare here for a returned parameter of the UDR	Must be unique among returned <i>parameter</i> names of the UDR. If any returned value of the UDR has a name, then all must have names.	"Identifier " on page 5-21

Usage

For compatibility with earlier Dynamic Server releases, you can create SPL functions with the `CREATE PROCEDURE` statement. (That is, you can include a Return clause in `CREATE PROCEDURE` statements.) Use `CREATE FUNCTION`, however, to create new SPL routines that return one or more values.

After the Return clause has indicated what data types are to be returned, you can use the RETURN statement of SPL at any point in the statement block to return SPL variables that correspond to the values in the Return clause.

Limits on Returned Values

An SPL function can specify more than one data type in the Return clause.

An external function (a function written in the C or the Java language) can specify only one data type in the Return clause, but an external function can return more than one row of data if it is an iterator function. For more information, see "ITERATOR" on page 5-63.

Subset of SQL Data Types

Not all data types are valid in the Return clause. For more information, see the table that follows. See also "Data Type " on page 4-21.

A user-defined function of Extended Parallel Server can return values of any built-in data type except SERIAL, SERIAL8, TEXT, or BYTE.

In Dynamic Server, a UDF can return values of any built-in data type except the complex, serial, and large object types that are not blank in the following table.

Data Type	C	Java	SPL
BIGSERIAL	X	X	X
BLOB	X		X
CLOB	X		X
BYTE	X	X	X
TEXT	X	X	X
COLLECTION		X	
LIST		X	
MULTISET		X	
ROW		X	
SET		X	
SERIAL	X	X	X
SERIAL8	X	X	X

In Dynamic Server, if you use a complex data type in the Return clause, the calling user-defined routine must define variables of the appropriate complex types to hold the values that the C or SPL user-defined function returns.

User-defined functions can return a value of opaque or distinct data types that are defined in the database.

The default precision of a DECIMAL value that an SPL function returns is 16 digits. For a function to return a DECIMAL with a different number of significant digits, you must specify the returned precision explicitly in the data type specification of the Return clause.

Using the REFERENCES Clause to Point to a Simple Large Object

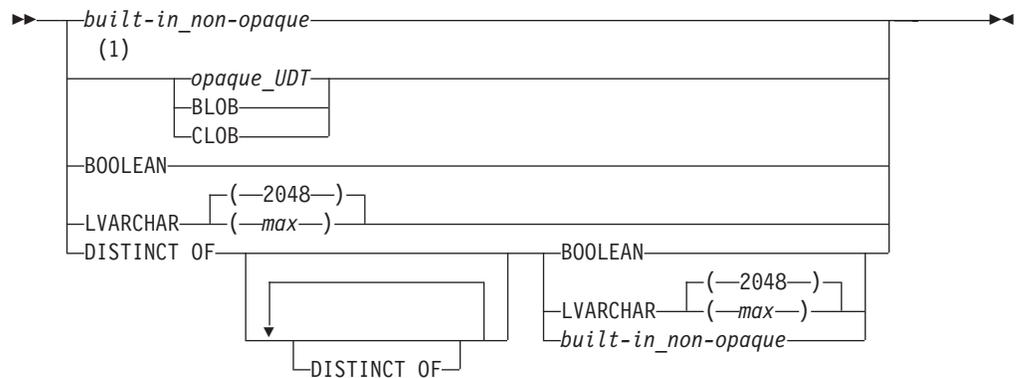
A user-defined function cannot return a BYTE or TEXT value (collectively called *simple large objects*) directly. A user-defined function can, however, use the REFERENCES keyword to return a descriptor that contains a pointer to a BYTE or TEXT object. The following example shows how to select a TEXT column within an SPL routine and then return the value:

```
CREATE FUNCTION sel_text()
  RETURNING REFERENCES text;
  DEFINE blob_var REFERENCES text;
  SELECT blob_col INTO blob_var
    FROM blob_table WHERE key_col = 10;
  RETURN blob_var;
END FUNCTION;
```

Returning a Value from Another Database (IDS)

The data types that a user-defined routine of Dynamic Server can return from another database have the following syntax:

For UDRs that access tables or views outside the local database, only the following data types are valid as return values:



Notes:

- 1 Not valid in cross-server operations

Element	Description	Restrictions	Syntax
<i>built-in _non-opaque</i>	Name of a built-in data type that is not opaque	Cannot be BIGSERIAL, BYTE, SERIAL, SERIAL8, or TEXT	"Data Type " on page 4-21
<i>max</i>	Maximum size in bytes. Default is 2048.	Must be an integer, where $1 \leq max \leq 32,739$	"Literal Number " on page 4-185
<i>opaque_UDT</i>	Name of a user-defined opaque data type	Must be cast explicitly to a built-in type by a cast defined in every participating database	"Identifier " on page 5-21

If the Return clause to return a value (or multiple values, in the case of an SPL function) from another database of the local Dynamic Server instance, the following data types are supported as the returned data type:

- Built-in data type that are not opaque
- Most of the *built-in opaque data types*, as listed in "Data Types in Cross-Database Transactions" on page 2-536

- Any DISTINCT type based on one of the built-in types that are identified in this list
- Any DISTINCT type based on one of the DISTINCT types in this list
- Any user-defined type (UDT) that is cast explicitly to one of the data types in this list.

The UDF and all of the DISTINCT types, opaque UDTs, data type hierarchies, and casts must have exactly the same definitions in each of the participating databases. The same data-type restrictions apply to a value that an external function returns from another database of the local Dynamic Server instance. For more information about data types that are supported in distributed operations across two or more databases of the same database server, see “Data Types in Cross-Database Transactions” on page 2-536. For the data type hierarchies that are valid for DISTINCT data types in distributed transactions, see “DISTINCT Types in Distributed Operations” on page 4-33.

From databases of other Dynamic Server instances, however, UDFs can specify only the following as a parameter or as a returned data type:

- Built-in data types that are not opaque
- BOOLEAN
- LVARCHAR
- DISTINCT of built-in types that are not opaque
- DISTINCT of BOOLEAN
- DISTINCT of LVARCHAR
- DISTINCT of the DISTINCT types in this list.

The definitions of the UDF and of any data type hierarchies, casts, and DISTINCT types must be exactly the same in each of the participating databases. Except for the BOOLEAN, DISTINCT, and LVARCHAR data types that are identified in the previous list, UDFs cannot return any other built-in opaque data type or opaque UDTs in cross-server function calls.

For more information about data types that are supported in distributed operations across two or more Dynamic Server instances, see “Data Types in Cross-Server Transactions” on page 2-537. For the data type hierarchies that are valid for DISTINCT data types in distributed transactions, see “DISTINCT Types in Distributed Operations” on page 4-33.

Named Return Parameters (IDS)

You can declare names for the returned parameters of an SPL routine, or a name for the single value that an external function can return.

If an SPL routine returns more than one value, you must either declare names for all of the returned parameters, or else none of them can have names. The names must be unique. Here is an example of named parameters:

```
CREATE PROCEDURE p (inval INT DEFAULT 0)
RETURNING INT AS serial_num,
           CHAR(10) AS name,
           INT AS points;
RETURN (inval + 1002), "Newton", 100;
END PROCEDURE;
```

Executing this UDR would return:

serial_num	name	points
1002	Newton	100

There is no relationship between the names of returned parameters and the names of any variables in the body of the routine. For example, you can define a function to return an INTEGER as `xval`, but in the body of the same function, a variable declared as `xval` could be of the data type INTERVAL YEAR TO MONTH.

Cursor and Noncursor Functions

A *cursor* function can fetch returned values one by one by iterating the generated result set of returned values. Such a function is an *implicitly iterated function*.

A function that returns only one set of values (such as one or more columns from a single row of a table) is a *noncursor* function.

The Return clause is valid in a cursor function or in a noncursor function. In the following example, the Return clause can return zero (0) or one value in a noncursor function. In a cursor function, however, it returns more than one row from a table, and each returned row contains zero or one value:

```
RETURNING INT;
```

In the following example, the Return clause can return zero (0) or two values if it occurs in a noncursor function. In a cursor function, however, it returns more than one row from a table, and each returned row contains zero or two values:

```
RETURNING INT, INT;
```

In both of the preceding examples, the receiving function or program must be written appropriately to accept the information that the function returns.

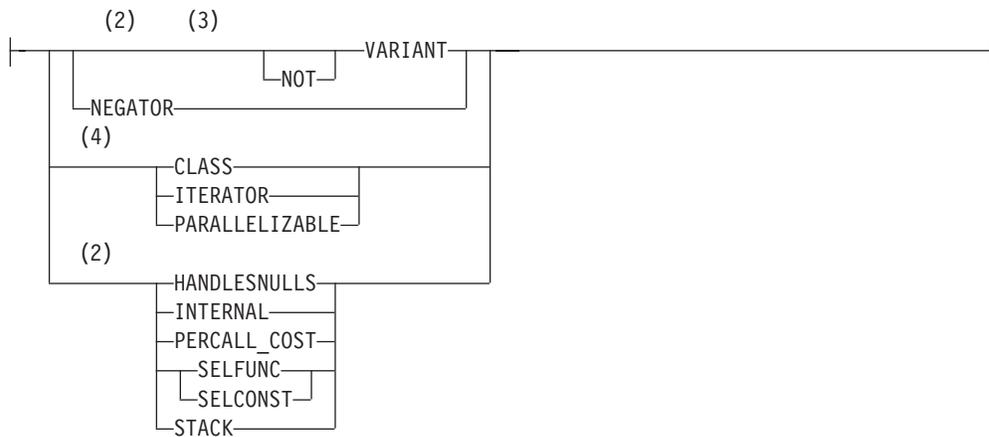
Routine Modifier

A routine modifier specifies characteristics of how a user-defined routine (UDR) behaves. Only Dynamic Server supports routine modifiers.

Syntax



Dropping a Routine Modifier:



Notes:

- 1 See “Adding or Modifying a Routine Modifier ”
- 2 C only
- 3 SPL only
- 4 External routines only

Element	Description	Restrictions	Syntax
<i>parameter</i>	Name that you declare here for a returned parameter of the UDR	Must be unique among returned <i>parameters</i> of UDRs. If any returned value of the UDR has a name, then all must have names.	“Identifier ” on page 5-21

Usage

If you drop an existing modifier in an ALTER FUNCTION, ALTER PROCEDURE, or ALTER ROUTINE statement, the database server sets the value of the modifier to the default value, if a default exists.

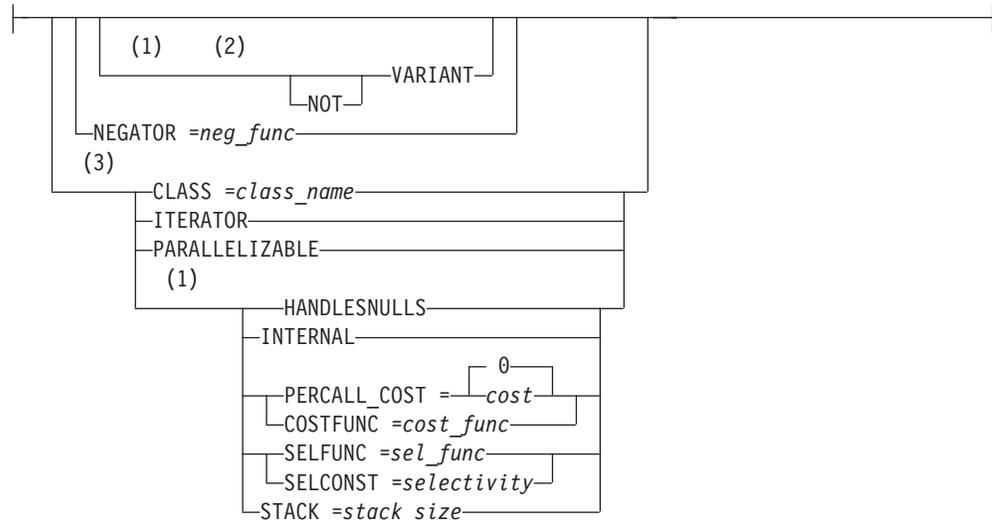
Some modifiers are available only with user-defined functions. For information on whether a specific routine modifier applies only to user-defined functions (that is, if it does not apply to user-defined procedures), see the description of the modifier in the sections that follow. In these sections, as elsewhere in this document, *external* refers to UDRs written in the C or Java languages. Features valid for only one language are so designated in the previous diagrams.

Except for VARIANT and NOT VARIANT modifiers, none of the options in this segment are valid for SPL routines.

Adding or Modifying a Routine Modifier

Use this segment in the ALTER FUNCTION, ALTER PROCEDURE, or ALTER ROUTINE statement to add or modify values for routine modifiers of a UDR.

Adding or Modifying a Routine Modifier:



Notes:

- 1 C language
- 2 Stored Procedure Language
- 3 External routines only

Element	Description	Restrictions	Syntax
<i>class_name</i>	Virtual processor (VP) class in which to run the external routine	Any C UDR must run in the CPU VP or in a user-defined VP class	“Quoted String” on page 4-189.
<i>cost</i>	CPU use cost for each invocation of a C language UDR. Default is 0.	Integer; $1 \leq cost \leq 2^{31}-1$ (highest cost).	“Literal Number ” on page 4-185
<i>cost_func</i>	Name of a companion user-defined cost function to invoke	Must have same owner as the UDR. Execute privilege needed to invoke	“Identifier ” on page 5-21
<i>neg_func</i>	Negator function that can be invoked instead of the UDR	Must have same owner as the UDR. Execute privilege needed to invoke	“Identifier ” on page 5-21
<i>sel_func</i>	Name of a companion user-defined selectivity function to invoke	Must have same owner as the UDR. Execute privilege needed to invoke	“Identifier ” on page 5-21
<i>selectivity</i>	CPU use cost for each invocation of a C language UDR. Default is 0.	See “Concept of Selectivity” on page 5-65.	“Literal Number ” on page 4-185
<i>stack_size</i>	Size (in bytes) of stack of the thread that executes the C-language UDR	Must be a positive integer	“Literal Number ” on page 4-185

You can add these modifiers in any order. If you list the same modifier more than once, the last setting overrides any previous values.

Modifier Descriptions

The following sections describe the modifiers that you can use to help the database server optimally execute a UDR.

CLASS

Use the CLASS modifier to specify the name of a virtual-processor (VP) class in which to run an external routine. A user-defined VP class must be defined before the UDR can be invoked.

You can execute C UDRs in the following types of VP classes:

- The CPU virtual-processor class (CPU VP)
- A user-defined virtual-processor class.

If you omit the CLASS modifier to specify a VP class for a UDR written in C, the UDR runs in the CPU VP. User-defined VP classes protect the database server from ill-behaved C UDRs. An ill-behaved C UDR has at least one of the following characteristics:

- It runs in the CPU VP for a long time without yielding.
- It is not thread safe.
- It calls an unsafe operating-system routine.

A well-behaved C UDR has none of these characteristics. Execute only well-behaved C UDRs in the CPU VP.

Warning: Execution of an ill-behaved C UDR in the CPU VP can cause serious interference with the operation of the database server, and the UDR might not produce correct results. For a discussion of ill-behaved UDRs, see the *IBM Informix Datablade API Programmer's Guide*.

By default, a UDR written in Java runs in a Java virtual processor class (JVP). Therefore, the CLASS modifier is optional for a UDR written in Java. However, use the CLASS modifier when you register a UDR written in Java to improve readability of your SQL statements.

COSTFUNC

Use the COSTFUNC modifier to specify the cost of a C UDR. The *cost* of the UDR is an estimate of the time required to execute it.

Occasionally, the cost of a UDR depends on its inputs. In that case, you can use a user-defined function to calculate a cost that depends on input values.

To execute *cost_func*, you must have Execute privilege on it and on the UDR.

HANDLESNULLS

Use the HANDLESNULLS modifier to specify that a C UDR can handle NULL values that are passed to it as arguments. If you do not specify HANDLESNULLS for a C language UDR, and if you pass to it an argument that has a NULL value, the UDR does not execute and returns a NULL value.

By default, a C language UDR does *not* handle NULL values.

The HANDLESNULLS modifier is not available for SPL routines because SPL routines handle NULL values by default.

INTERNAL

Use the INTERNAL modifier with an external routine to specify that an SQL or SPL statement cannot call the external routine. An external routine that is specified as INTERNAL is not considered during routine resolution. Use the INTERNAL modifier for external routines that define access methods, language managers, and so on.

By default, an external routine is *not* internal; that is, an SQL or SPL statement can call the routine.

ITERATOR

Use the ITERATOR modifier with external functions to specify that the function is an *iterator function*. An iterator function is a function that returns a single element per function call to return a set of data; that is, it is called with an initial call and zero or more subsequent calls until the set is complete.

By default, an external C or Java language function is *not* an iterator function.

An SPL iterator function requires the RETURN WITH RESUME statement, rather than the ITERATOR modifier.

In ESQL/C, an iterator function requires a cursor. The cursor allows the client application to retrieve the values one at a time with the FETCH statement.

For more information on how to write iterator functions, see *IBM Informix User-Defined Routines and Data Types Developer's Guide* and the *IBM Informix Datablade API Programmer's Guide*.

For information about using an iterator function with a virtual table interface in the FROM clause of a query, see "Iterator Functions (IDS)" on page 2-551.

NEGATOR

Use the NEGATOR modifier with UDRs that return Boolean values.

The NEGATOR modifier designates another user-defined function, called a *negator function*, as a companion to the current function. A negator function takes the same arguments as its companion function, in the same order, but returns the Boolean complement.

That is, if a function returns TRUE for a given set of arguments, its negator function returns FALSE when passed the same arguments, in the same order. For example, the following functions are negator functions:

```
equal(a,b)  
notequal(a,b)
```

Both functions take the same arguments, in the same order, but return complementary Boolean values. When it is more efficient to do so, the query optimizer can use the negator function instead of the function that you specify.

To invoke a user-defined function that has a negator function, you must have the Execute privilege on both functions. In addition, the function must have the same owner as its negator function.

PARALLELIZABLE

Use the PARALLELIZABLE modifier to indicate that an external routine can be executed in parallel in the context of a parallelizable data query (PDQ).

By default, an external routine is non-parallelizable; that is, it executes in sequence.

If your UDR has a complex or smart large object data type as either a parameter or a returned value, you cannot use the PARALLELIZABLE modifier.

If you specify the PARALLELIZABLE modifier for an external routine that cannot be parallelizable, the database server returns a runtime error.

A C language UDR that calls only PDQ thread-safe DataBlade API functions is parallelizable. These categories of DataBlade API functions are PDQ thread safe:

- Data handling
An exception in this category is that collection manipulation functions (**mi_collection_***) are not PDQ thread safe.
- Session, thread, and transaction management
- Function execution
- Memory management
- Exception handling
- Callbacks
- Miscellaneous

For details of the DataBlade API functions that are included in each category, see the *IBM Informix Datablade API Function Reference*.

If your C language UDR calls a function that is not included in one of these categories, it is not PDQ thread safe and is therefore not parallelizable.

To parallelize Java language UDR calls, the database server must have multiple instances of JVPs. UDRs written in the Java language and that open a JDBC connection are not parallelizable.

PERCALL_COST (C)

Use the PERCALL_COST modifier to specify the approximate CPU usage cost that a C language UDR incurs each time it executes. The optimizer uses the cost you specify to determine the order in which to evaluate SQL predicates in the UDR for best performance. For example, the following query has two predicates joined by a logical AND:

```
SELECT * FROM tab1 WHERE func1() = 10 AND func2() = 'abc';
```

In this example, if one predicate returns FALSE, the optimizer need not evaluate the other predicate.

The optimizer uses the specified cost to order the predicates so that the least expensive predicate is evaluated first. The CPU usage cost must be an integer between 1 and $2^{31}-1$, with 1 the lowest cost and $2^{31}-1$ the most expensive.

To calculate an approximate cost per call, add the following two figures:

- The number of lines of code executed each time the C UDR is called
- The number of predicates that require an I/O access

The default cost per execution is 0. When you drop the PERCALL_COST modifier, the cost per execution returns to 0.

SELCONST (C)

Use the SELCONST modifier to specify the selectivity of a C UDR. The *selectivity* of the UDR is an estimate of the fraction of the rows that the query will select.

The value of selectivity constant, **selconst**, is a floating-point number between 0 and 1 that represents the fraction of the rows for which you expect the UDR to return TRUE.

SELFUNC (C)

Use the SELFUNC modifier with a C UDR to name a companion user-defined function, called a *selectivity function*, to the current UDR. The selectivity function provides selectivity information about the current UDR to the optimizer.

The *selectivity* of a UDR is an estimate of the fraction of the rows that the query will select. That is, it is an estimate of the number of times the UDR will execute.

To execute *sel_func*, you must have Execute privilege on it and on the UDR.

Concept of Selectivity

Selectivity is an attribute of queries that performs a search based on an equality condition. The selectivity of the query depends inversely on the proportion of qualifying rows. The smaller the proportion of qualifying rows among all the rows in FROM clause table objects, the more selective is the query.

For example, the following query has a search condition based on the **customer_num** column in the **customer** table:

```
SELECT * FROM customer WHERE customer_num = 102;
```

Because each row in the table has a different customer number, this query is highly selective. In contrast, the following query has low selectivity:

```
SELECT * FROM customer WHERE state = 'CA';
```

Because most of the rows in the **customer** table are for customers in California, more than half of the rows in the table would be returned.

Restrictions on the SELFUNC Modifier

The selectivity function that you specify must satisfy the following criteria:

- It must take the same number of arguments as the current C UDR.
- The data type of each argument must be SELFUNCARGS.
- It must return a value of type FLOAT between 0 and 1, which represents the percentage of selectivity of the function. (1 is highly selective; 0 is not at all selective.)
- It can be written in any language that the database server supports.

A user who invokes the C UDR must have the Execute privilege both on that UDR and on the selectivity function that the SELFUNC modifier specifies.

Both the C UDR and the selectivity function must have the same owner.

For information on how to use the **mi_funcarg*** functions to extract information about the arguments of a selectivity function, see the *IBM Informix Datablade API Programmer's Guide*.

STACK (C)

Use the STACK modifier with a C UDR to override the default stack size that the STACKSIZE configuration parameter specifies.

The STACK modifier specifies the size (in bytes) of the thread stack, which a user thread that executes the UDR uses to hold information such as routine arguments and returned values from functions.

A UDR needs to have enough stack space for all its local variables. For a particular UDR, you might need to specify a stack size larger than the default size to prevent stack overflow.

When a UDR that includes the STACK modifier executes, the database server allocates a thread-stack size of the specified number of bytes. Once the UDR completes execution, subsequent UDRs execute in threads with a stack size that the STACKSIZE configuration parameter specifies (unless any of these subsequent UDRs have also specified the STACK modifier).

For more information about the thread stack, see your *IBM Informix Administrator's Guide* and the *IBM Informix Datablade API Function Reference*.

VARIANT and NOT VARIANT

Use the VARIANT and NOT VARIANT modifiers with C user-defined functions and SPL functions. A function is *variant* if it returns different results when it is invoked with the same arguments or if it modifies a database or variable state. For example, a function that returns the current date or time is a variant function.

By default, user-defined functions are variant. If you specify NOT VARIANT when you create or modify a user-defined function, the function cannot contain any SQL statements.

If the user-defined function is nonvariant, the database server might cache the returned values of expensive functions. You can create functional indexes only on nonvariant functions. For more information on functional indexes, see "CREATE INDEX " on page 2-124.

In ESQ/C, you can specify VARIANT or NOT VARIANT in this clause or in the EXTERNAL Routine Reference. For more information, see " External Routine Reference" on page 5-19. If you specify the modifier in both places, however, you must use the same modifier in both clauses.

Examples

The following statement includes an external routine reference for a Java language UDR. You must first register *demo_jar* using the procedure **install_jar**(*absolute path*><*jar file name*>,<*internal registered name*>).

```
CREATE FUNCTION delete_order(int) RETURNING int
  WITH (NOT VARIANT)
  EXTERNAL NAME 'informix.demo_jar:delete_order.delete_order()'
  LANGUAGE JAVA;
```

Related Information

For more information on user-defined routines, see *IBM Informix User-Defined Routines and Data Types Developer's Guide* and the *IBM Informix Datablade API Programmer's Guide*.

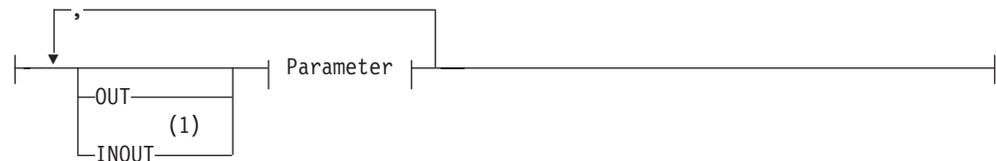
For more information about how these modifiers can affect performance, see your *IBM Informix Performance Guide*.

Routine Parameter List

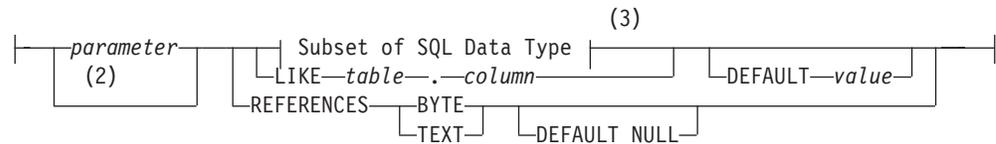
Use the appropriate part of the Routine Parameter List segment whenever you see a reference to a Routine Parameter List in a syntax diagram.

Syntax

Routine Parameter List:



Parameter:



Notes:

- 1 Dynamic Server only
- 2 External routines only
- 3 See "Subset of SQL Data Types" on page 5-68

Element	Description	Restrictions	Syntax
<i>column</i>	Name of a column whose data type is declared for <i>parameter</i>	Must exist in the specified table	"Database Object Name" on page 5-16
<i>parameter</i>	Name of a parameter of the UDR	Name is required for SPL routines	"Identifier" on page 5-21
<i>table</i>	Table that contains <i>column</i>	The table must exist in the database	"Identifier" on page 5-21
<i>value</i>	Default used if UDR is called with no value for <i>parameter</i>	Must be a literal, of the same data type as <i>parameter</i> . For opaque types, an input function must be defined	"Literal Number" on page 4-185

Usage

A *parameter* is a formal argument in the declaration of a UDR. (When you subsequently invoke a UDR that has parameters, you must substitute an actual argument for the parameter, unless the parameter has a default value.)

The name of the parameter is optional for external routines of Dynamic Server.

When you create a UDR, you declare a *name* and *data type* for each parameter. You can specify the data type directly, or use the LIKE or REFERENCES clause to specify the data type. You can optionally specify a default value.

You can define any number of SPL routine parameters, but the total length of all parameters passed to an SPL routine must be less than 64 kilobytes.

No more than nine arguments to a UDR written in the Java language can be DECIMAL data types of SQL that the UDR declares as BigDecimal data types of the Java language.

Any C language UDR that returns an opaque data type must specify **opaque_type** in the var binary declaration of the C host variable.

Subset of SQL Data Types

Serial and large-object data types are not valid as parameters. A UDR can declare a parameter of any other data type defined in the database, including any built-in data types except BIGSERIAL, BLOB, BYTE, CLOB, SERIAL, SERIAL8, or TEXT.

On Dynamic Server, a parameter can also be a complex data type or a UDT, but complex data types are not valid for parameters of external UDRs written in the Java language.

For information about the data types of Dynamic Server that are valid as parameters or return values of routines that access tables or views outside the local database, see “**Returning a Value from Another Database (IDS)**” on page 5-57.

Using the LIKE Clause

Use the LIKE clause to specify that the data type of a parameter is the same as a column defined in the database. If the ALTER TABLE statement changes the data type of the column, the data type of the parameter also changes.

In Dynamic Server, if you use the LIKE clause to declare any parameter, you cannot overload the UDR. For example, suppose you create the following user-defined procedure:

```
CREATE PROCEDURE cost (a LIKE tableX.colY, b INT)
. . .
END PROCEDURE;
```

You cannot create another procedure named **cost()** in the same Dynamic Server database with two arguments. You can, however, create a procedure named **cost()** with a number of arguments other than two. (Another way to circumvent this restriction on the LIKE clause is through user-defined data types.)

Using the REFERENCES Clause

Use the REFERENCES clause to specify that a parameter contains BYTE or TEXT data. The REFERENCES keyword allows you to use a pointer to a BYTE or TEXT

object as a parameter. If you use the DEFAULT NULL option in the REFERENCES clause, and you call the UDR without a parameter, a NULL value is used as the default value.

Using the DEFAULT Clause

Use the DEFAULT keyword followed by an expression to specify a default value for a parameter. If you provide a default value for a parameter, and the UDR is called with fewer arguments than were defined for that UDR, the default value is used. If you do not provide a default value for a parameter, and the UDR is called with fewer arguments than were defined for that UDR, the calling application receives an error.

The following example shows a CREATE FUNCTION statement that specifies a default value for a parameter. This function finds the square of the *i* parameter. If the function is called without specifying the argument for the *i* parameter, the database server uses the default value 0 for the *i* parameter.

```
CREATE FUNCTION square_w_default
  (i INT DEFAULT 0) {Specifies default value of i}
  RETURNING INT; {Specifies return of INT value}
  DEFINE j INT; {Defines routine variable j}
  LET j = i * i; {Finds square of i and assigns it to j}
  RETURN j; {Returns value of j to calling module}
END FUNCTION;
```

In Extended Parallel Server, to re-create this example, use the CREATE PROCEDURE statement instead of the CREATE FUNCTION statement.

Extended Parallel Server supports no more than a single OUT parameter. If one is specified, it must be the last item in the parameter list.

Warning: When you specify a date value as the default value for a parameter, make sure to specify 4 digits instead of 2 digits for the year. When you specify a 2-digit year, the **DBCENTURY** environment variable setting can affect how the database server interprets the date value, so the UDR might not use the default value that you intended. For more information, see the *IBM Informix Guide to SQL: Reference*.

Specifying OUT Parameters for User-Defined Routines

When you register a user-defined routine of Dynamic Server, you can use the OUT keyword to specify that any parameter in the list is an OUT parameter. Each OUT parameter corresponds to a value the routine returns indirectly, through a pointer. The value that the routine returns through the pointer is an extra value, in addition to any values that it returns explicitly.

After you have registered a user-defined function that has one or more OUT parameters, you can use the function with a statement-local variable (SLV) in a SELECT statement. (For information about statement-local variables, see “Statement-Local Variable Expressions (IDS)” on page 4-164.)

If you specify any OUT parameters, and you use Informix-style parameters, the arguments are passed to the OUT parameters by reference. The OUT parameters are not significant in determining the routine signature.

For example, the following declaration of a C user-defined function allows you to return an extra value through the *y* parameter:

```
int my_func( int x, int *y );
```

Register the C function with a CREATE FUNCTION statement similar to this:

```
CREATE FUNCTION my_func( x INT, OUT y INT )
  RETURNING INT
  EXTERNAL NAME "/usr/lib/local_site.so"
  LANGUAGE C
END FUNCTION;
```

In the next example, this Java method returns an extra value by passing an array:

```
public static String allVarchar(String arg1, String[] arg2)
  throws SQLException
{
  arg2[0] = arg1;
  return arg1;
}
```

To register this as a UDF, use a statement similar to the following example:

```
CREATE FUNCTION all_varchar(VARCHAR(10), OUT VARCHAR(7))
  RETURNING VARCHAR(7)
  WITH (class = "jvp")
  EXTERNAL NAME 'informix.testclasses.jlm.Param.allVarchar(java.lang.String,
  java.lang.String[ ])'
  LANGUAGE JAVA;
```

Specifying INOUT Parameters for a User-Defined Routine (UDR)

UDRs of Dynamic Server that are written in the SPL, C, or Java languages can also support INOUT parameters. When the UDR is invoked, a value for each INOUT parameter is passed by reference as an argument to the UDR. When the UDR completes execution, it can return a modified value for the INOUT parameter to the calling context. The INOUT parameter can be of any data type that Dynamic Server supports, including user-defined and complex data types.

In the following example, the CREATE PROCEDURE statement registers a C routine that has a single INOUT parameter:

```
CREATE PROCEDURE CALC ( INOUT param1 float )
  EXTERNAL NAME "$INFORMIXDIR/etc/myudr.so(calc)"
  LANGUAGE C;
```

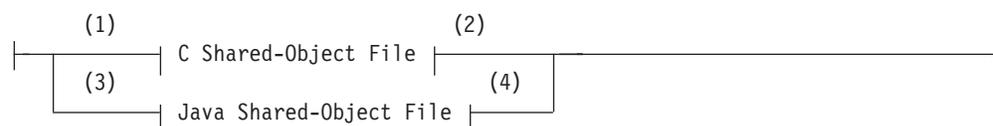
You can assign INOUT parameters to statement-local variables (SLVs), which the section "Statement-Local Variable Expressions (IDS)" on page 4-164 describes.

Shared-Object Filename

Use a shared-object filename to specify a pathname to an executable object file when you register or alter an external routine. Only Dynamic Server supports this syntax segment.

Syntax

Shared-Object File:



Notes:

- 1 C only
- 2 See "C Shared-Object File"
- 3 Java only
- 4 See "Java Shared-Object File" on page 5-72

Usage

If the IFX_EXTEND_ROLE configuration parameter is set to 1 or to ON, only users to whom the DBSA has granted the built-in EXTEND role are authorized to use this segment.

The Database Server Administrator should include in the DB_LIBRARY_PATH configuration parameter settings every file system where the security policy authorizes DataBlade modules and UDRs to reside. Unless DB_LIBRARY_PATH is absent or has no setting, the database server cannot access a file that this segment specifies unless its pathname begins with a string that exactly matches one of the values of DB_LIBRARY_PATH.

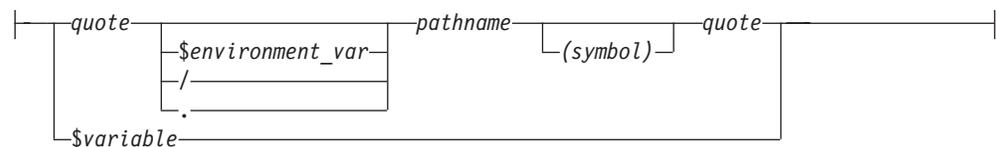
For example, if "\$INFORMIXDIR/extend" is one of the DB_LIBRARY_PATH values on a Linux system, then shared-object files can have pathnames within the \$INFORMIXDIR/extend file system or its subdirectories. (This is also the file system where built-in DataBlade modules reside, and the default location where the DataBlade Developers Kit creates user-defined DataBlade modules.)

The syntax by which you specify a shared-object filename depends on whether the external routine is written in the C language or in the Java language. Sections that follow describe each of these external languages.

C Shared-Object File

To specify the location of a C shared-object file, specify the path to the dynamically loaded executable file within a quoted pathname or as a variable.

C Shared-Object File:



Element	Description	Restrictions	Syntax
<i>environment_var</i>	Platform-independent indicator	Must begin with a dollar sign (\$)	"Identifier " on page 5-21
<i>pathname</i>	Pathname to the file	See notes that follow this table	Must conform to operating system conventions
<i>quote</i>	Either single (') or double (") quotation mark symbol	Opening and closing quotation mark symbols must match	Literal symbol (either ' or ")

Element	Description	Restrictions	Syntax
<i>symbol</i>	Entry point to the file	Must be enclosed in parentheses	Must conform to operating system conventions
<i>variable</i>	Platform-independent indicator	Must begin with a dollar sign (\$)	Must conform to C language conventions

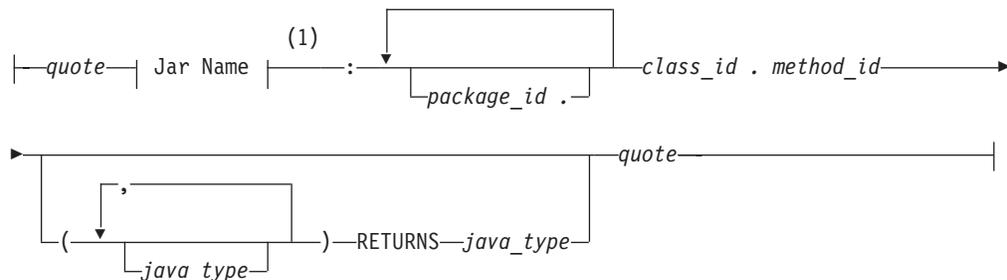
The following rules affect *pathname* and filename specifications in C:

- A filename (with no *pathname*) can specify an internal function.
- You can omit the period (.) symbol if *pathname* is relative to the current directory when the CREATE or ALTER statement is run.
- On UNIX, an absolute pathname must begin with a slash (/) symbol, and each directory name must end with a slash (/) symbol.
- On Windows, an absolute pathname must begin with a backslash (\) symbol, and each directory name must end with a backslash (\) symbol.
- The filename at the end of *pathname* must have the .so file extension and must refer to an executable file in a shared object library.
- Use a *symbol* only if the entry point to the dynamically loadable executable object file has a different name from the UDR that you are registering with CREATE FUNCTION or CREATE PROCEDURE.
- If you specify a *variable*, it must contain the full pathname to the executable file.
- You can include white space characters, such as blank spaces or tab characters, within a quoted pathname.

Java Shared-Object File

To specify the name of a Java shared-object file, specify the name of the static Java method to which the UDR corresponds and the location of the Java binary that defines the method.

Java Shared-Object File:



Notes:

- 1 See "Jar Name" on page 5-34

Element	Description	Restrictions	Syntax
<i>class_id</i>	Java class whose method implements the UDR	Class must exist in the .jar file that Jar Name identifies	Must conform to rules for Java identifiers

Element	Description	Restrictions	Syntax
<i>java_type</i>	Java data type for a parameter in the Java-method signature	Must be defined in a JDBC class or by an SQL-to-Java mapping	Must conform to rules for Java identifiers
<i>method_id</i>	Name of the Java method that implements the UDR	Must exist in the Java class that <i>java_class_name</i> specifies	Must conform to rules for Java identifiers
<i>package_id</i>	Name of package that contains the Java class	Must exist	Must conform to rules for Java identifiers
<i>quote</i>	Single (') or double (") quotation mark delimiters	Opening and closing quotation marks must match	Literal symbol (' or ") entered at the keyboard

Before you can create a UDR written in the Java language, you must assign a jar identifier to the external jar file with the **sqlj.install_jar** procedure. (For more information, see “sqlj.install_jar” on page 6-58.) You can include the Java signature of the method that implements the UDR in the shared-object filename.

- If you do *not* specify the Java signature, the routine manager determines the *implicit* Java signature from the SQL signature in the CREATE FUNCTION or CREATE PROCEDURE statement.
It maps SQL data types to the corresponding Java data types with the JDBC and SQL-to-Java mappings. For information on mapping user-defined data types to Java data types, see “sqlj.setUDTextName” on page 6-61.
- If you do specify the Java signature, the routine manager uses this *explicit* Java signature as the name of the Java method to use.

For example, if the Java method **explosiveReaction()** implements the Java UDR **sql_explosive_reaction()** as discussed in “sqlj.install_jar” on page 6-58, its shared-object filename could be:

```
course_jar:Chemistry.explosiveReaction
```

The preceding shared-object filename provides an implicit Java signature. The following shared-object filename is the equivalent with an explicit Java signature:

```
course_jar:Chemistry.explosiveReaction(int)
```

Related Information

See the “ External Routine Reference” on page 5-19 segment for the context in which a shared-object filename appears within EXTERNAL NAME clause of the ALTER FUNCTION, ALTER PROCEDURE, ALTER ROUTINE, CREATE FUNCTION, and CREATE PROCEDURE statements. For further information on the DB_LIBRARY_PATH parameter in the ONCONFIG file, see the *IBM Informix Administrator's Reference*.

Specific Name

Use a specific name to declare an identifier for a UDR that is unique in the database or name space. Use the Specific Name segment whenever you see a reference to a specific name in a syntax diagram. Only Dynamic Server supports specific names.

Syntax

Specific Name:



Element	Description	Restrictions	Syntax
<i>owner</i>	Owner of the UDR	No more than 32 bytes. Must be same as <i>owner</i> of function or procedure name of this UDR. See also "Restrictions on the Owner Name."	" Owner Name " on page 5-45
<i>specific_id</i>	Unique name of the UDR	Must be no more than 128 bytes long. See also "Restrictions on the Specific Name."	"Identifier " on page 5-21

Usage

A *specific name* is a unique identifier that the CREATE PROCEDURE or CREATE FUNCTION statement declares as an alternative name for a UDR.

Because you can overload routines, a database can have more than one UDR with the same name and different parameter lists. You can assign a UDR a specific name that uniquely identifies the specific UDR.

If you declare a specific name when you create the UDR, you can later use that name when you alter, drop, grant, or revoke privileges, or update statistics on that UDR. Otherwise, you need to include the parameter data types with the UDR name, if the name alone does not uniquely identify the UDR.

Restrictions on the Owner Name

When you declare a specific name, the *owner* must be the same *authorization identifier* that qualifies the function name or procedure name of the UDR that you create. That is, whether or not you specify the owner name to qualify either the UDR name or the specific name or both, the names of the *owner* must match.

When you specify no owner name in the DDL statement that creates a UDR, Dynamic Server uses the login name of the user who creates the UDR. Therefore, if you specify the owner name in one location and not the other, the owner name that you specify must match your user ID.

Restrictions on the Specific Name

In a database that is not ANSI-compliant, *specific_id* must be unique among routine names within the database. Two UDRs cannot have the same *specific_id*, even if they have different owners.

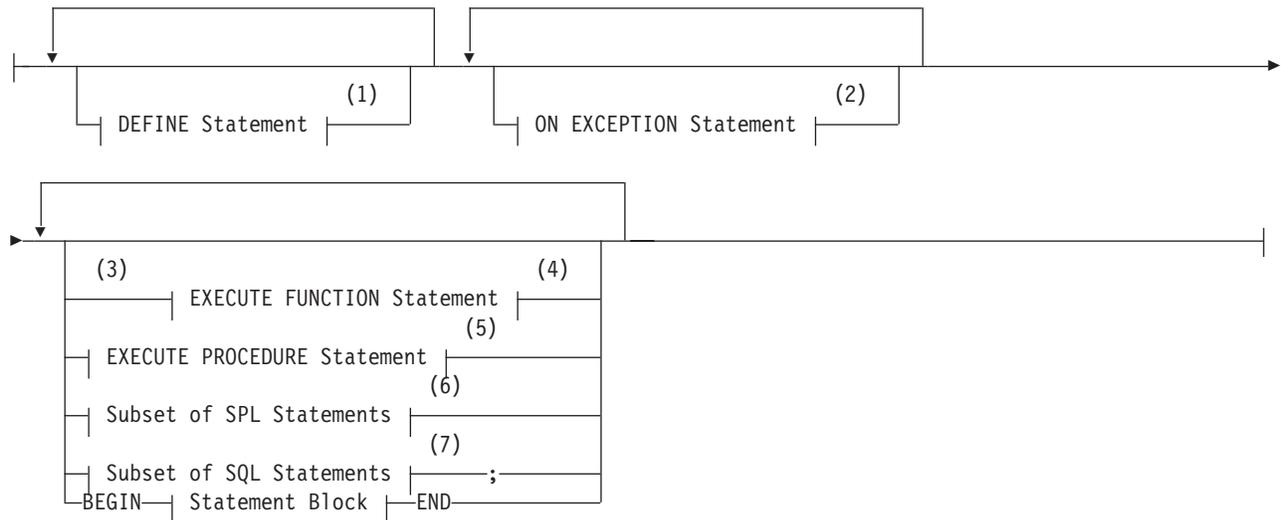
In an ANSI-compliant database, the combination *owner.specific_id* must be unique. That is, the specific name must be unique among UDRs that have the same owner.

Statement Block

Use a statement block to specify SPL and SQL operations to take place when an SPL statement that includes this segment is executed.

Syntax

Statement Block:



Notes:

- 1 See "DEFINE " on page 3-8
- 2 See "ON EXCEPTION " on page 3-38 and "Identifier " on page 5-21
- 3 Dynamic Server only
- 4 See "EXECUTE FUNCTION " on page 2-356
- 5 See "EXECUTE PROCEDURE " on page 2-364
- 6 See "Subset of SPL Statements Valid in the Statement Block"
- 7 See " SQL Statements Valid in SPL Statement Blocks" on page 5-76

Usage

SPL and SQL statements can appear in a *statement block*, a set of zero or more statements that can define the scope of a variable or of the ON EXCEPTION statement. If a statement block is empty, no operation takes place when control of execution within the SPL routine passes to the empty SPL statement block.

Subset of SPL Statements Valid in the Statement Block

The diagram for the "Statement Block " on page 5-74 refers to this section. You can use any of the following SPL statements in the statement block:

- <<Label >>
- CALL
- CONTINUE
- EXIT
- FOR
- FOREACH
- GOTO
- IF

- LET
- LOOP
- RAISE EXCEPTION
- RETURN
- SYSTEM
- TRACE
- WHILE

GOTO and << *Label* >>, however, are not valid in ON EXCEPTION statement blocks.

SQL Statements Valid in SPL Statement Blocks

The diagram for the “Statement Block ” on page 5-74 refers to this section. Most SQL statements are valid in SPL statement blocks, except for the statements that are listed below. The following SQL statements are *not* valid in an SPL statement block:

- CLOSE DATABASE
- CONNECT
- CREATE DATABASE
- CREATE FUNCTION
- CREATE FUNCTION FROM
- CREATE PROCEDURE
- CREATE PROCEDURE FROM
- CREATE ROUTINE FROM DATABASE
- DISCONNECT
- EXECUTE
- FLUSH
- INFO
- LOAD
- OUTPUT
- PUT
- SET AUTOFREE
- SET CONNECTION
- UNLOAD

For example, you cannot close the current database or connect to a new database within an SPL routine. Similarly, you cannot drop the current SPL routine within the same routine. You can, however, drop another SPL routine.

Only two forms of the SELECT statement are valid in queries within SPL routines:

- You can use the INTO TEMP clause to put the results of the SELECT statement into a temporary table.
- You can use the SELECT ... INTO form of the SELECT statement to put the resulting values into SPL variables.

When you include the ORDER BY clause in the SELECT ... INTO TEMP or the SELECT ... INTO *variable* statement, you imply that the query returns more than

one row. The database server issues an error if you specify the ORDER BY clause without a FOREACH loop to process the returned rows individually within the SPL routine.

If an SPL routine is called as part of a data-manipulation language (DML) statement, additional restrictions exist. For more information, see “Restrictions on SPL Routines in Data-Manipulation Statements” on page 5-78.

Nested Statement Blocks

You can use the BEGIN and END keywords to delimit a statement block that is nested within another statement block.

Scope of Reference of SPL Variables and Exception Handlers

The BEGIN and END keywords can limit the scope of SPL variables and exception handlers. Declarations of variables and definitions of exception handlers inside a BEGIN and END statement block are local to that statement block and are not visible from outside the statement block. The following code uses a BEGIN and END statement block to delimit the scope of reference of variables:

```
CREATE DATABASE demo;
CREATE TABLE tracker (
  who_submitted CHAR(80), -- Show what code was running.
  value INT,              -- Show value of the variable.
  sequential_order SERIAL -- Show order of statement execution.
);
CREATE PROCEDURE demo_local_var()
DEFINE var1, var2 INT;
  LET var1 = 1;
  LET var2 = 2;
  INSERT INTO tracker (who_submitted, value)
  VALUES ('var1 param before sub-block', var1);
BEGIN
  DEFINE var1 INT; -- same name as global parameter.
  LET var1 = var2;
  INSERT INTO tracker (who_submitted, value)
  VALUES ('var1 var defined inside the "IF/BEGIN".', var1);
END
INSERT INTO tracker (who_submitted, value)
  VALUES ('var1 param after sub-block (unchanged!)', var1);
END PROCEDURE;
EXECUTE PROCEDURE demo_local_var();
SELECT sequential_order, who_submitted, value FROM tracker
  ORDER BY sequential_order;
```

This example declares three variables, two of which are named **var1**. (Name conflicts are created here to illustrate which variables are visible. Using the same name for different variables is generally not recommended, because conflicting names of variables can make your code more difficult to read and to maintain.)

Because of the statement block, only one **var1** variable is in scope at a time.

The **var1** variable that is declared inside the statement block is the only **var1** variable that can be referenced from within the statement block.

The **var1** variable that is declared outside the statement block is not visible within the statement block. Because it is out of scope, it is unaffected by the change in value to the **var1** variable that takes place inside the statement block. After all the statements run, the outer **var1** still has a value of 1.

The `var2` variable is visible within the statement block because it was not superseded by a name conflict with a block-specific variable.

Restrictions on SPL Routines in Data-Manipulation Statements

If you call the SPL routine in a SQL statement that is not a data-manipulation language (DML) statement (namely EXECUTE FUNCTION or EXECUTE PROCEDURE), the SPL routine can execute any statement that is not listed in the section “SQL Statements Valid in SPL Statement Blocks” on page 5-76.

In Dynamic Server, if you call the SPL routine as part of a DML statement (namely, an INSERT, UPDATE, DELETE, or SELECT statement), then the routine cannot execute any of the following SQL statements:

- ALTER ACCESS_METHOD
- ALTER FRAGMENT
- ALTER INDEX
- ALTER OPTICAL CLUSTER
- ALTER SEQUENCE
- ALTER TABLE
- BEGIN WORK
- COMMIT WORK
- CREATE ACCESS_METHOD
- CREATE AGGREGATE
- CREATE DISTINCT TYPE
- CREATE OPAQUE TYPE
- CREATE OPCLASS
- CREATE ROLE
- CREATE ROW TYPE
- CREATE SEQUENCE
- CREATE TRIGGER
- DELETE
- DROP ACCESS_METHOD
- DROP AGGREGATE
- DROP INDEX
- DROP OPCLASS
- DROP OPTICAL CLUSTER
- DROP ROLE
- DROP ROW TYPE
- DROP SEQUENCE
- DROP SYNONYM
- DROP TABLE
- DROP TRIGGER
- DROP TYPE
- DROP VIEW
- INSERT
- RENAME COLUMN
- RENAME DATABASE
- RENAME SEQUENCE

- RENAME TABLE
- ROLLBACK WORK
- SET CONSTRAINTS
- TRUNCATE
- UPDATE

In Extended Parallel Server, if you call the SPL routine as part of a DML statement (namely, an INSERT, UPDATE, DELETE, MERGE, or SELECT statement), then the routine can execute *only* the following statements of SQL:

- SELECT
- SET PLOAD FILE
- SET DEBUG FILE TO
- SET EXPLAIN
- SET OPTIMIZATION

For both Dynamic Server and Extended Parallel Server, these restrictions do not apply to an SPL routine that is invoked by a trigger, because in this case the SPL routine is not called by the DML statement, and therefore can include any SQL statement, such as UPDATE, INSERT and DELETE, that is not listed among the “SQL Statements Valid in SPL Statement Blocks” on page 5-76.

Transactions in SPL Routines

In a database that is not ANSI-compliant, you can use the BEGIN WORK and COMMIT WORK statements in an SPL statement block to start a transaction, to finish a transaction, or start and finish a transaction in the same SPL routine. If you start a transaction in a routine that is executed remotely, you must finish the transaction before the routine exits.

As previously noted, however, the ROLLBACK WORK statement is not valid in an SPL statement block.

Support for Roles and User Identity

You can use roles with SPL routines. You can execute role-related statements (CREATE ROLE, DROP ROLE, GRANT, REVOKE, and SET ROLE) and a user who holds the SETSESSIONAUTH privilege can issue SET SESSION AUTHORIZATION statements within an SPL routine. You can also grant privileges or grant other roles to roles with the GRANT or (for Dynamic Server only) GRANT FRAGMENT statement within an SPL routine, or use REVOKE to cancel the privileges (or roles) of roles. Privileges that a user has acquired by enabling a role or by a SET SESSION AUTHORIZATION statement are not automatically relinquished after an SPL routine completes execution.

For further information about roles, see the CREATE ROLE, DROP ROLE, GRANT, REVOKE, and SET ROLE statements in Chapter 2.

Chapter 6. Built-In Routines

In This Chapter

This chapter describes routines whose identifiers are known to the database server when a database is created, but whose usage differs in various ways from that of the built-in SQL functions that “Expression” on page 4-40 describes.

The built-in routines that this chapter describes can be classified within the following categories, according to the tasks that these routines perform:

- SQL Admin API functions
 - `admin()`
 - `task()`
- Session configuration procedures
 - `sysdbclose()`
 - `sysdbopen()`
- The `explain_sql` routine
- UDR definition routines
 - `ifx_replace_module()`
 - `ifx_unload_module()`
 - `jvpcontrol()`
 - `sqlj.alter_java_path()`
 - `sqlj.install_jar()`
 - `sqlj.remove_jar()`
 - `sqlj.replace_jar()`
 - `sqlj.setUDTExtName()`
 - `sqlj.unsetUDTExtName()`
 - `sysibm.Metadata()`
 - `sysibm.sqlcaMessage()`

The `ADMIN` and `TASK` functions exist only in the `sysadmin` database. The `Metadata` and `sqlcaMessage` routines are created automatically if the database server is configured to support the DRDA protocol. The `sysdbopen` and `sysdbclose` routines can be defined in every Dynamic Server database.

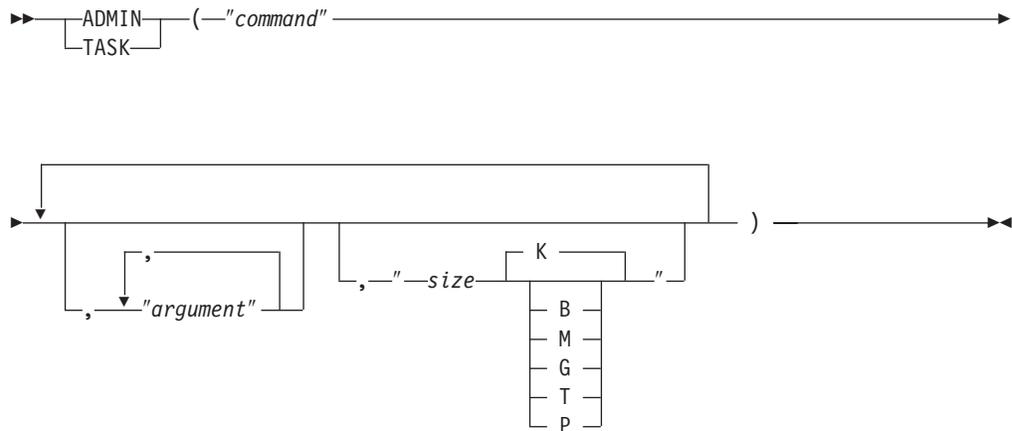
Sections that follow describe these categories and the individual routines within each category.

SQL Admin API Functions

Use the built-in SQL Admin API functions `ADMIN` or `TASK` to accomplish administrative tasks remotely through `EXECUTE FUNCTION` statements of SQL that emulate various administrative command-line utilities of Dynamic Server.

Syntax

SQL Administration API Functions:



Element	Description	Restrictions	Syntax
<i>argument</i>	Argument that follows the <i>command</i> . These are interpreted as arguments to the <i>command</i> .	No more than 28 arguments. For <i>command</i> strings that correspond to utilities, see your <i>Administrator's Guide</i> .	"Quoted String" on page 4-189
<i>command</i>	The initial argument: the mnemonic name of a task, or else ONMODE.	See the tables of <i>command</i> strings in the Usage section.	"Quoted String" on page 4-189
<i>size</i>	A number, if <i>argument</i> specifies an offset, storage object, or buffer	Blank space is optional between <i>size</i> and the initial letter of its storage unit abbreviation	"Literal Number " on page 4-185

Usage

The **ADMIN** and **TASK** functions provide an SQL interface to emulate many administrative command-line utilities of Dynamic Server, and to perform additional operations that no utility supports. Because these are SQL functions, you can invoke them within SQL statements that can include an expression, or you can use the EXECUTE FUNCTION statement to call them. For example, the following query uses the **TASK** function to check the data portion of all of the tables in the **root** dbspace:

```

SELECT task("check data", partnum)
  FROM sysmaster:systables
 WHERE trunc(partnum/1048575)=1 ;
    
```

These built-in functions are defined only in the **sysadmin** database of each Dynamic Server instance. By default, they can be invoked only by user **informix**. If Connect privilege on the **sysadmin** database is granted to user **root** or to **DBSA** group members, they too can invoke the **ADMIN** and **TASK** functions of the SQL Admin API when they are connected directly or remotely to the **sysadmin** database. For more information about the **sysadmin** database, see the *IBM Informix Administrator's Reference*.

Every call to the **ADMIN** or **TASK** function attempts to produce two results:

- To perform a *command* task, typically emulating an administrative utility
- To insert a new row into the **command_history** table of the **sysadmin** database.

The operation that the initial *command* argument to **ADMIN** or **TASK** specifies occurs in a separate transaction from the insertion of the new row into the **command_history** table. If the command executes successfully, but the insertion into **command_history** fails, the command takes effect, but an **online.log** error entry indicates the problem.

These two functions differ primarily in their names and in their returned value, which indicates what occurred when the function was called:

- **TASK** returns the value of the **cmd_ret_msg** column in the new row that it inserts into the **command_history** table. This LVARCHAR value indicates the result (or the failure) of the command.
- **ADMIN** returns an integer, showing the serial value of the **cmd_number** column in the new row that **ADMIN** inserts into the **command_history** table.
 - If this value is greater than zero, the command succeeded, and a new row was inserted into the **command_history** table.
 - If this value is zero, the command succeeded, but Dynamic Server could not insert a new row into the **command_history** table.
 - If this value is less than zero, the command failed, but a new row was inserted into the **command_history** table.

The *command* specification and any additional arguments define the administrative task for the **ADMIN** or **TASK** function. For example, this SQL statement, equivalent to the **oncheck -ce** command, instructs the database server to check all the extents:

```
EXECUTE FUNCTION admin('check extents');
```

This example also check blobspaces, smart-large-object extents, and user-data and metadata information in sbspace chunks. If a second argument had specified a valid dbspace number, only the extents of that dbspace would have been checked.

If the **command_history.cmd_number** serial counter is 200 when this function is called, and the command succeeds, then Dynamic Server executes the command and returns the integer 201. If the command fails, this example returns the value -201.

Suppose the **TASK** function had executed the same command:

```
EXECUTE FUNCTION task('check extents');
```

This instructs the database server to check the extents, and returns a message indicating whether the command succeeded or failed.

ADMIN and TASK Arguments and Size Specifications

Most calls to the **ADMIN** or **TASK** function require one or more additional specifications as quoted strings that follow *command* in the comma-separated list of arguments to the **ADMIN** or **TASK** function. These typically correspond to arguments to the command-line utility that *command* emulates, but with each *argument* delimited by a pair of single (') quotation marks or double (") quotation marks.

Arguments that follow the initial *command* argument are resolved positionally, rather than through named parameters. In commands for which the SQL Admin

API defines default values, you can omit parameters for which a default exists, but the argument list must not include any argument that has no default value and that follows an omitted parameter.

For example, if a command recognizes up to five input parameters, and only the third has a default value, and the last two parameters are optional, then a valid argument list can include either two, three, four, or five arguments.

- If you supply only two arguments, the **ADMIN** or **TASK** function executes as if three parameters were supplied, including a default value for the third parameter.
- If you supply values for the first two input parameters, but none for the third, you must not include argument values for the fourth and fifth parameters. If you do, the function call will either fail with an error, or else return an unpredictable result by misinterpreting the positional order of the last two arguments, which will be processed as values for the third and fourth input parameters, rather than for the fourth and fifth.

This is the expected behavior of any function call whose argument list does not use named parameters.

The following example illustrates a command task, equivalent to the **onspaces -c -d** command, that includes additional arguments and a *size* specification:

```
EXECUTE FUNCTION admin( 'create dbspace',  
    'dbspace2', '/usr/informix/dbspace2', '20MB', '5MB');
```

This example used a pathname to specify the storage location of the new dbspace, but you can also specify pathnames that begin with an environment variable whose value is set to a pathname, as in this example:

```
EXECUTE FUNCTION admin( 'create dbspace',  
    'dbspace2', '$INFORMIXDIR/dbspace2', '20MB', '5MB');
```

In processing the example above, the value of the **INFORMIXDIR** environment variable replaces the **\$INFORMIXDIR** portion of the pathname, and the full path would be `/usr/informix/dbspace2` (assuming that **INFORMIXDIR** is set to `/usr/informix`).

This **ADMIN** function call corresponds to the following **onspaces** command:

```
onspaces -c -d dbspace2 -p /usr/informix/dbspace2 -o 20000 -s 5000
```

If the **command_history.cmd_number** serial counter is 201 when this function is called, and the command fails, then the returned value is -202. Suppose that the next SQL administration API function that the DBSA invokes is this:

```
EXECUTE FUNCTION admin('create dbspace',  
    'dbspace2', '/work/CHUNKS/dbspace2', "20M");
```

If in this case the command succeeds, the returned value is 203. The DBSA can use the following query to examine the two rows of the **command_history** table that these calls to the **ADMIN** function inserted:

```
SELECT * FROM command_history WHERE cmd_number IN (202,203);
```

This query returns two rows:

```
cmd_number      202  
cmd_exec_time   2007-11-17 16:26:14  
cmd_user        informix  
cmd_hostname    olympia  
cmd_executed    create dbspace
```

```

cmd_ret_status  -1
cmd_ret_msg     Unable to create file /work/dbspace2

cmd_number      203
cmd_exec_time   2007-11-17 16:26:15
cmd_user        informix
cmd_hostname    olympia
cmd_executed    create dbspace
cmd_ret_status  0
cmd_ret_msg     created dbspace number 2 named dbspace2

```

Here the *size* units are specified as "M" for megabytes. If no units are specified, the default units are kilobytes, which in this example would produce a dbspace of probably insufficient size. The same notation for units and the "K" default also applies to arguments to commands of **ADMIN** or **TASK** that specify offsets. Dynamic Server recognizes the following specifications for units in *size* arguments to the **ADMIN** or **TASK** function:

Notation

Corresponding Units

'B' or 'b'
Bytes (= 2 to the power 0)

'K' or 'k'
Kilobytes (= 2 to the power 10)

'M' or 'm'
Megabytes (= 2 to the power 20)

'G' or 'g'
Gigabytes (= 2 to the power 30)

'T' or 't'
Terabytes (= 2 to the power 40)

'P'
Petabytes (= 2 to the power 50)

The lettercase of these characters is ignored. You can replace the double (") quotation mark delimiters with single (') quotation marks, but the opening and closing delimiters of a string must be the same. The initial letter in the input specification ('B', 'K', 'M', 'G', 'T' or 'P') determines the units of measure, and any trailing characters are ignored. (This closely resembles the syntax of the **FORMAT_UNITS** function of SQL, except that this function requires PB to specify petabytes, and interprets any other string starting with "P" (such as "PA", "pc", "PhD", "papyrus" and so forth) as specifying *pages*, rather than *petabytes*. Unlike the **FORMAT_UNITS** function, however, **ADMIN** and **TASK** support no abbreviation for pages as a unit.)

If no units are specified for *size* arguments, the default unit is kilobytes, even if the corresponding command-line utility argument has some different default unit of measure. When a *size* argument is omitted, the default size for that object applies by default, based either on the setting of a configuration parameter, or on the system default if no parameter is set. Storage spaces, for example, have a default size of 100 megabytes.

Any white space that separates the *size* specification and the units abbreviation in the same argument to the **ADMIN** or **TASK** function is ignored. For example, the specifications '128M' and '128 m' are both interpreted as 128 megabytes.

See also the description of the built-in “FORMAT_UNITS Function” on page 4-158, which **ADMIN** and **TASK** call internally to process size specification arguments.

Mnemonic Initial Commands

As their first argument, the SQL Admin API functions accept the following mnemonic *command* specifications that identify some operation, in many cases emulating one of the Dynamic Server command-line utilities. Each row in the following table describes one of the mnemonic commands.

- The first column lists the first argument of the command, You can omit a required argument that has a default value, or an optional argument, but no additional arguments can follow an omitted argument.
- The second column tersely describes the effect of the command.
- The third column identifies any corresponding command-line utility, configuration parameter, or SQL statement that produces similar results.

See the *IBM Informix Administrator's Reference* for information about the configuration parameters that can define some of these defaults, and for information about the semantics of the command-line utilities that some of these mnemonic commands emulate.

Table 6-1. Mnemonic SQL Admin API Commands

Initial Command	Effect of Successful Operation	Utility or Configuration Parameter That Can Produce a Similar Result
ADD BUFFERPOOL	Creates a new buffer pool.	onparams -b -g BUFFERPOOL parameter
ADD CHUNK	Adds a chunk to a <i>dbspace</i> or <i>blobospace</i> in a specified <i>pathname</i> .	onspaces -a ON-Monitor
ADD LOG	Adds a log to a <i>dbspace</i> .	onparams -a -d ON-Monitor
ADD MEMORY	Increases the virtual portion of shared memory.	onmode -a SHMADD configuration parameter
ADD MIRROR	Adds a mirror chunk to a <i>dbspace</i> of a specified <i>path</i> and <i>offset</i> .	onspaces -m
ALTER CHUNK OFFLINE	Changes the status of a mirrored chunk in a <i>dbspace</i> , a non-primary chunk within a noncritical <i>dbspace</i> , a <i>blobospace</i> or an <i>sbspace</i> to offline.	onspaces -s
ALTER CHUNK ONLINE	Changes the status of a mirrored chunk in a <i>dbspace</i> , a non-primary chunk within a noncritical <i>dbspace</i> , a <i>blobospace</i> or an <i>sbspace</i> to online.	onspaces -s
ALTER LOGMODE	Changes database logging mode to buffered, unbuffered, ANSI, or non-logging.	ondblog, ontape
ALTER PLOG	Changes the location and size of the physical log.	onparams -p
ARCHIVE FAKE	Archives a fake backup operation.	ontape -F

Table 6-1. Mnemonic SQL Admin API Commands (continued)

Initial Command	Effect of Successful Operation	Utility or Configuration Parameter That Can Produce a Similar Result
CHECK DATA	Checks the data portion of every page in a partition.	oncheck -cd
CHECK EXTENTS	Checks the extents of one or all dbspaces.	oncheck -ce
CHECK PARTITION	Checks the data portion of a partition.	oncheck -pt
CHECKPOINT	Force a checkpoint.	onmode -c
CLEAN SBSPACE	Releases from the <i>sbspace</i> any unreferenced BLOB or CLOB objects.	onspaces -cl
CREATE BLOBSPACE	Creates a blobspace.	onspaces -c -b
CREATE CHUNK	Adds a chunk to a dbspace or to a blobspace.	onspaces -a ON-Monitor
CREATE DBSPACE	Creates a dbspace in specified path, and (if necessary) creates path.	onspaces -c -d
CREATE SBSPACE	Creates an sbspace in specified path and (if necessary) creates path.	onspaces -c -S
CREATE TEMPDBSPACE	Creates a temporary dbspace in specified path and (if necessary) creates path.	onspaces -c -d -t
DROP BLOBSPACE	Drops the specified blobspace.	onspaces -d ON-Monitor
DROP CHUNK	Drops the specified chunk.	onspaces -d ON-Monitor
DROP DBSPACE	Drops the specified dbspace.	onspaces -d ON-Monitor
DROP LOG	Drops the specified log.	onparams -d
DROP SBSPACE	Drops the specified sbspace.	onspaces -d ON-Monitor
DROP TEMPDBSPACE	Drops the specified temporary dbspace.	onspaces -d ON-Monitor
PRINT ERROR	Prints the error message associated with an error number or the headers of a specified partition number, optionally including index-specific information and page-allocation information by page type for dbspaces.	finderr
PRINT PARTITION	Prints the headers of a specified partition number. TASK returns the same output as oncheck -pt , but ADMIN returns a count of rows.	oncheck -pt
PRINT PARTITION FULL	Same as PRINT PARTITION, but also displays index-specific information and page-allocation information by page type (for dbspaces).	oncheck -pT
RENAME SPACE	Renames a dbspace, blobspace, sbspace, or extspace.	onspaces -ren

Table 6-1. Mnemonic SQL Admin API Commands (continued)

Initial Command	Effect of Successful Operation	Utility or Configuration Parameter That Can Produce a Similar Result
RESET SYSADMIN	Moves the sysadmin database to a specified dbspace.	No comparable utility commands.
SCHEDULER SHUTDOWN	Stops the scheduler and deallocates all its resources..	No comparable utility command.
SCHEDULER START	Starts the scheduler.	No comparable utility command.
SET CHUNK OFFLINE	Changes the status of a dbspace, blobspace, or sbspace from online to offline. Chunk must be in a mirrored pair, or a non-primary chunk within a noncritical dbspace.	onspaces -s -p
SET CHUNK ONLINE	Changes the status of a dbspace, blobspace, or sbspace from offline to online. Chunk must be in a mirrored pair, or a non-primary chunk within a noncritical dbspace.	onspaces -s
SET DATASKIP OFF	Disables the DATASKIP setting.	DATASKIP configuration parameter
SET DATASKIP ON	Specifies that the database server should skip designated dbspaces.	DATASKIP configuration parameter SQL statement SET DATASKIP ON
SET INDEX COMPRESSION	Modifies the level at which two partially used index pages are merged. The pages are merged if the data on those pages totals a set level.	onmode -C BTSCANNER configuration parameter
SET ONCONFIG	Changes the value of the specified configuration parameter.	onmode -wm
SET ONCONFIG MEMORY	Changes the value of the specified configuration parameter in memory.	onmode -wf
SET ONCONFIG PERMANENT	Changes the value of the specified configuration parameter permanently.	onmode -wm
SET SBSPACE ACESSTIME OFF	Disables the tracking of the time of access to all smart large objects stored in the sbspace.	onspaces -ch
SET SBSPACE ACESSTIME ON	Tracks the time of access to all smart large objects stored in the sbspace.	onspaces -ch
SET SBSPACE AVG_LO_SIZE	Specifies an expected average sbspace size.	onspaces -ch
SET SBSPACE LOGGING OFF	Specifies that the database server does not log changes to the user data area of the sbspace.	onspaces -c -S
SET SBSPACE LOGGING ON	Specifies that the database server logs changes to the user data area of the sbspace.	onspaces -c -S
SET SQL TRACING	Enables global SQL tracing with default values.	No comparable utility commands.
SET SQL TRACING DATABASE ADD	Specifies tracing for one or more databases, rather than for all databases. The default is all databases. The maximum number of database names that can be set is 16.	No comparable utility commands.

Table 6-1. Mnemonic SQL Admin API Commands (continued)

Initial Command	Effect of Successful Operation	Utility or Configuration Parameter That Can Produce a Similar Result
SET SQL TRACING DATABASE CLEAR	Clears all databases from the list of database being traced. Returns tracing back to the default of all databases.	No comparable utility commands.
SET SQL TRACING DATABASE LIST	Lists the databases which are currently being traced.	No comparable utility commands.
SET SQL TRACING DATABASE REMOVE	Removes a single database from the list of databases being traced	No comparable utility commands.
SET SQL TRACING INFO	Returns a single status line about the global SQL tracing state.	No comparable utility commands.
SET SQL TRACING OFF	Turns off all SQL tracing and frees all associated resources.	No comparable utility commands.
SET SQL TRACING ON	Sets global SQL tracing to specified or default values.	No comparable utility commands.
SET SQL TRACING RESUME	Restarts SQL tracing when it has been suspended.	No comparable utility commands.
SET SQL TRACING SESSION	Starts SQL tracing for the current session. This command behaves in the same manner as SET SQL USER TRACING.	No comparable utility commands.
SET SQL TRACING SESSION CLEAR	Clears all sessions from the SQL tracing for the current session. This command behaves in the same manner as SET SQL USER TRACING.	No comparable utility commands.
SET SQL TRACING SESSION OFF	Turns on SQL tracing for the current session. This command behaves in the same manner as SET SQL USER TRACING.	No comparable utility commands.
SET SQL TRACING SESSION ON	Sets SQL tracing for the current session. This command behaves in the same manner as SET SQL USER TRACING.	No comparable utility commands.
SET SQL TRACING SUSPEND	Pauses SQL tracing without deallocating any resources.	No comparable utility commands.
SET SQL TRACING USER ADD	Adds a user to the user tracing list.	No comparable utility commands.
SET SQL TRACING USER CLEAR	Removes all users from the user tracing list.	No comparable utility commands.
SET SQL TRACING USER LIST	Lists all users on the user trace list.	No comparable utility commands.
SET SQL TRACING USER REMOVE	Removes a user from the user trace list.	No comparable utility commands.
SET SQL USER TRACING	Changes the mode of global SQL tracing for a specified user session.	No comparable utility commands.
SET SQL USER TRACING CLEAR	Clears any user tracing flags for the specified user session so that it adheres to the global tracing policy.	No comparable utility commands.
SET SQL USER TRACING OFF	Disables SQL tracing for a user session, even if the global mode is ON.	No comparable utility commands.
SET SQL USER TRACING ON	Enables user SQL tracing for a user session. Even if the global tracing mode is OFF, this user will trace SQL statements.	No comparable utility commands.

Table 6-1. Mnemonic SQL Admin API Commands (continued)

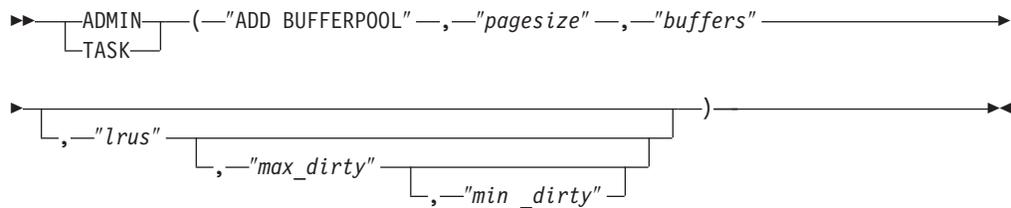
Initial Command	Effect of Successful Operation	Utility or Configuration Parameter That Can Produce a Similar Result
START MIRRORING	Starts mirroring the specified dbspace, blobspace, or sbspace.	onspaces -m
STOP MIRRORING	Stops mirroring the specified dbspace, blobspace, or sbspace.	onspaces -r

For information about SQL Admin API commands for high-availability, see “High Availability Initial Commands” on page 6-44. For information about SQL Admin API commands for Enterprise Replication, see “Enterprise Replication Initial Commands” on page 6-48.

ADD BUFFERPOOL SQL Admin API Command

Creates a new buffer pool of a specified *pagesize*. If you specify no additional arguments, this creates 1000 buffers, using default values from BUFFERPOOL configuration parameter.

ADD BUFFERPOOL



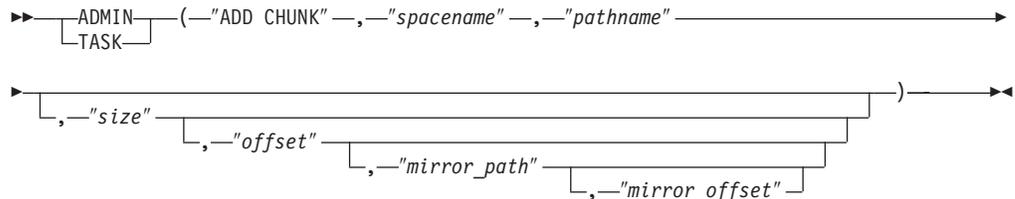
Argument	Description	Restrictions
<i>pagesize</i>	The page size in kilobytes.	The page size must be an integral multiple of the default page size, and cannot be greater than 16 kilobytes. On Windows, the page size is always 4 kilobytes.
<i>buffers</i>	Number of buffers. Each buffer is the size of the operating system page. The default is 10000.	
<i>lrus</i>	Number of LRU queues.	The maximum for 32-bit platforms is 128, and for 64-bit platforms is 512. The default is 8.
<i>max_dirty</i>	Specifies the percentage of modified pages in the LRU queues at which the queue is cleaned. If a field is specified out of the range of values, the default of 60.00 percent is set.	

Argument	Description	Restrictions
<i>min_dirty</i>	Specifies the percentage of modified pages in the LRU queues at which page cleaning is no longer mandatory. Page cleaners might continue cleaning beyond this point under some circumstances. If a field is specified out of the range of values, the default of 80.00 percent is set.	

ADD CHUNK SQL Admin API Command

Adds a chunk to a dbspace or blobspace in a specified pathname.

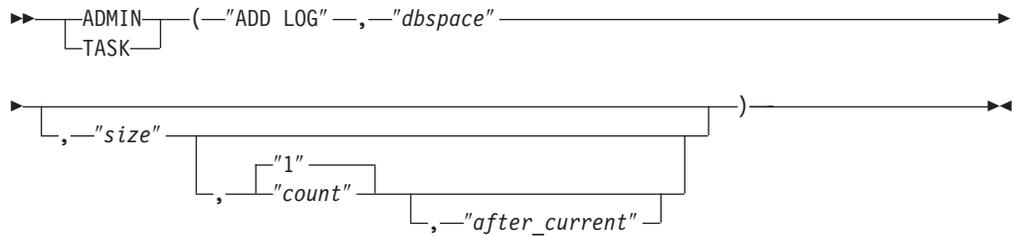
The size of the chunk must be equal to or greater than 1000 KB and a multiple of the page size. The starting offset plus the chunk size cannot exceed the maximum chunk size. The maximum offset is 4 TB.



Argument	Description	Restrictions
<i>spacename</i>	The name of the dbspace, blobspace, or sbspace to which you are adding disk space	
<i>pathname</i>	The path of the added disk space	
<i>size</i>	The amount of disk space to add in kilobytes	
<i>offset</i>	The location of the new chunk	
<i>mirror_path</i>	The path to the mirror chunk. If you are adding a chunk to a mirrored storage space, you must also add a mirror chunk.	
<i>mirror_offset</i>	The location of the mirror chunk	

ADD LOG SQL Admin API Command

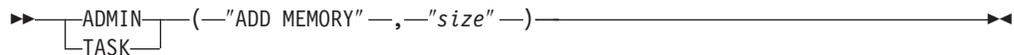
Adds a log to a dbspace. If you specify no additional arguments, this creates one log file at the end of the list. The size defaults to the LOGSIZE configuration parameter.



Argument	Description	Restrictions
<i>dbspace</i>	The name of the dbspace to add a log to	
<i>size</i>	The size in kilobytes of the new logical-log file	
<i>count</i>	The number of log files to create	
<i>after_current</i>	Inserts the logical-log file after the current log file	

ADD MEMORY SQL Admin API Command

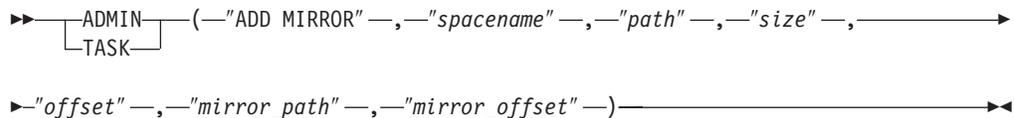
Adds to the virtual portion of shared memory a positive size in kilobytes. This size defaults to the SHMADD configuration parameter.



Argument	Description	Restrictions
<i>size</i>	The number of kilobytes of new virtual shared-memory segment	This value must not exceed the operating system limit on the size of shared-memory segments

ADD MIRROR SQL Admin API Command

Adds a mirror chunk to a dbspace of a specified path and offset.



Argument	Description	Restrictions
<i>spacename</i>	The name of a dbspace, blobspace, or sbspace to mirror	
<i>path</i>	The disk partition or unbuffered device of the initial chunk of the dbspace, blobspace, or sbspace that you want to mirror	

Argument	Description	Restrictions
<i>size</i>	The amount of storage to use for the mirrored dbspace, blobspace, or sbspace	
<i>offset</i>	The offset into the disk partition or into the unbuffered device in kilobytes to reach the initial chunk of the newly mirrored dbspace, blobspace, or sbspace	
<i>mirror_path</i>	The disk partition or unbuffered device of the initial chunk of the dbspace, blobspace, or sbspace that performs the mirroring	
<i>mirror_offset</i>	The offset to reach the mirrored chunk of the newly mirrored dbspace, blobspace, or sbspace	

ALTER CHUNK SQL Admin API Command

Brings online or takes offline a chunk in a dbspace, blobspace, or sbspace whose path and offset you specify. The chunk must be in a mirrored pair, or a non-primary chunk within a noncritical dbspace.

```

▶▶ [ADMIN TASK] ("ALTER CHUNK [OFFLINE|ONLINE] ", "path" , "offset" ) ▶▶

```

Argument	Description	Restrictions
ALTER CHUNK OFFLINE	Indicates that you want to bring the chunk offline	
ALTER CHUNK ONLINE	Indicates that you want to restore the chunk and bring it online	
<i>path</i>	Indicates the disk partition or unbuffered device of the chunk	
<i>offset</i>	Indicates, in kilobytes, the offset into the disk partition or unbuffered device to reach the chunk. Defaults to 0.	

ALTER LOGMODE SQL Admin API Command

Changes database logging mode to ANSI, buffered, non-logging, or unbuffered.

Unlike **ondblog** or **ontape**, the database remains accessible, and no level-0 backup is required. This command fails if any other session is active.

```

▶▶ [ADMIN] (—"ALTER LOGMODE" —,—"database" —,—"A" —)
   [TASK] [ "B" ]
          [ "N" ]
          [ "U" ]
▶▶

```

Argument	Description	Restrictions
<i>database</i>	The name of the database whose logging mode you want to alter	
A	Changes database logging to be ANSI compliant.	After you create or convert a database to ANSI mode, you cannot change it back to any of the other logging modes.
B	Sets the logging mode so that transaction information is written to a buffer before it is written to a logical log	
N	Sets the logging mode so that no database transactions are logged	
U	Sets the logging mode so that data is not written to a buffer before it is written to a logical log	

ALTER PLOG SQL Admin API Command

Changes the location and size of the physical log. To change only the size, specify current dbspace of the physical log.

```

▶▶ [ADMIN] (—"ALTER PLOG" —,—"dbspace" — [—"size" —])
   [TASK]
▶▶

```

Argument	Description	Restrictions
<i>dbspace</i>	Changes the location of the physical log to the specified dbspace	
<i>size</i>	Changes the size, specified in kilobytes, of the physical log	

ARCHIVE FAKE SQL Admin API Command

Performs a fake backup operation.

```

▶▶ [ADMIN] (—"ARCHIVE FAKE" —)
   [TASK]
▶▶

```

For more information on fake archives, see *IBM Informix Backup and Restore Guide*.

Argument	Description	Restrictions
UNBLOCK	Unblocks the database server. When the database server is unblocked, data transactions and normal database server operations can resume. Use this option after you complete an external backup on Dynamic Server.	

CLEAN SBSPACE SQL Admin API Command

Releases from the sbspace any unreferenced BLOB or CLOB objects.

```

>> ADMIN
   TASK  ("CLEAN SBSPACE" —,—"sbspace" —)

```

Argument	Description	Restrictions
<i>sbspace</i>	Name of the sbspace to clean	

CREATE BLOBSPACE SQL Admin API Command

Creates a blobspace.

```

>> ADMIN
   TASK  ("CREATE
   WITH_CHECK  BLOBSPACE" —,—"blobspace" —,—"path" —
   ,"size"
   ,"offset"
   ,"page_size"
   )

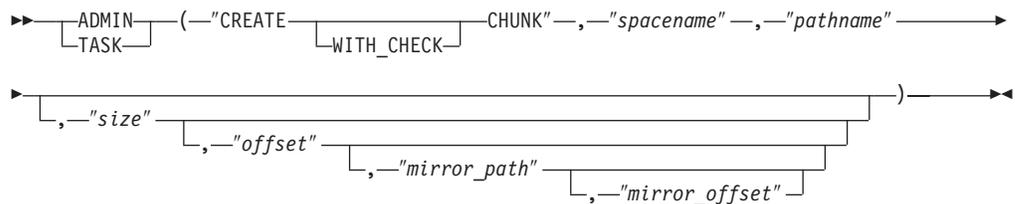
```

Argument	Description	Restrictions
WITH_CHECK	Verifies that the path exists	
<i>blobspace</i>	Name of the blobspace to be created	
<i>path</i>	The disk partition or device of the initial chunk of the blobspace or dbspace that you are creating	
<i>size</i>	The size, in kilobytes, of the initial chunk of the new blobspace or dbspace	
<i>offset</i>	The offset, in kilobytes, into the disk partition or into the device to reach the initial chunk of the new blobspace, dbspace, or sbspace	

Argument	Description	Restrictions
<i>page_size</i>	The blobpage size in terms of <i>page_unit</i> , the number of disk pages in each blobpage. Valid page sizes are 2, 4, 6, 8, 10, 12, or 16 KB for a 2 KB computer, and 4, 8, 12, or 16 KB for a 4 KB computer.	

CREATE CHUNK SQL Admin API Command

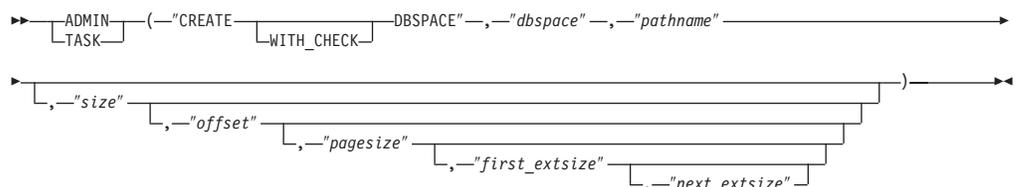
Creates a chunk in a dbspace or to a blobpage.



Argument	Description	Restrictions
<i>spacename</i>	The name of the dbspace, blobpage, or sbspace to which you are adding disk space	
<i>pathname</i>	The path of the added disk space	
<i>size</i>	The amount of disk space to add in kilobytes	
<i>offset</i>	The location of the new chunk	
<i>mirror_path</i>	The path to the mirror chunk. If you are adding a chunk to a mirrored storage space, you must also add a mirror chunk.	
<i>mirror_offset</i>	The location of the mirror chunk	
WITH_CHECK	Returns an error if <i>pathname</i> does not exist	

CREATE DBSPACE SQL Admin API Command

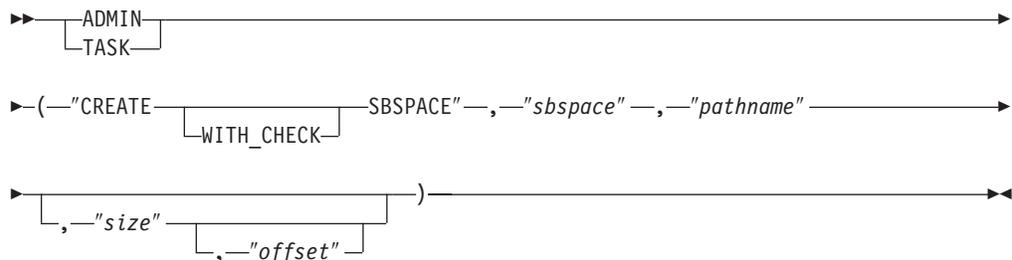
Creates a dbspace in specified path.



Argument	Description	Restrictions
<i>dbspace</i>	Names the dbspace to be created	
<i>first_extsize</i>	Indicates, in kilobytes, the size of the first extent for the tblspace tblspace	
<i>next_extsize</i>	Indicates, in kilobytes, the size of the next extents in the tblspace tblspace	
<i>offset</i>	Indicates, in kilobytes, the offset into the disk partition or into the device to reach the initial chunk of the new dbspace	
<i>pagesize</i>	Indicates, in kilobytes, the non-default page size for the new dbspace	
<i>pathname</i>	Indicates the disk partition or device of the initial chunk of the dbspace that you are creating. Valid page sizes are 2, 4, 6, 8, 10, 12, or 16 KB for a 2 KB computer, and 4, 8, 12, or 16 KB for a 4 KB computer.	
<i>size</i>	Indicates, in kilobytes, the size of the initial chunk of the new dbspace	
WITH_CHECK	Returns an error if <i>pathname</i> does not exist	

CREATE SBSPACE SQL Admin API Command

Creates an sbSPACE in specified path.

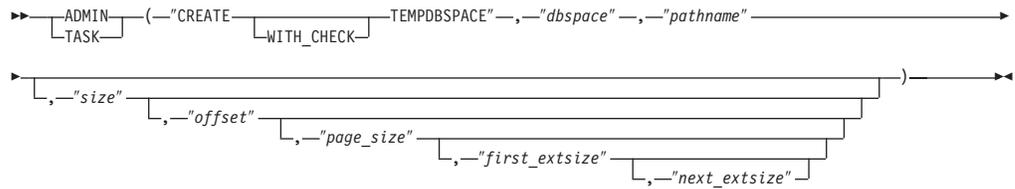


Argument	Description	Restrictions
<i>offset</i>	Indicates, in kilobytes, the offset into the disk partition or into the device to reach the initial chunk of the sbSPACE	
<i>pathname</i>	Indicates the disk partition or unbuffered device of the initial chunk of the sbSPACE	

Argument	Description	Restrictions
<i>sbspace</i>	Names the sbspace to be created	
<i>size</i>	Indicates, in kilobytes, the size of the initial chunk of the new sbspace	
WITH_CHECK	Returns an error if <i>pathname</i> does not exist	

CREATE TEMPDBSPACE SQL Admin API Command

Creates a temporary dbspace



Argument	Description	Restrictions
<i>dbspace</i>	Names the temporary dbspace to be created	
<i>first_extsize</i>	Indicates, in kilobytes, the size of the first extent for the tblspace <i>tblspace</i>	
<i>next_extsize</i>	Indicates, in kilobytes, the size of the next extent in the tblspace <i>tblspace</i>	
<i>offset</i>	Indicates, in kilobytes, the offset into the disk partition or unbuffered device to reach the chunk	
<i>page_size</i>	Indicates in kilobytes, the non-default page size for the new temporary dbspace. Valid page sizes are 2, 4, 6, 8, 10, 12, or 16 KB for a 2 KB computer, and 4, 8, 12, or 16 KB for a 4 KB computer.	
<i>pathname</i>	Indicates the disk partition or device of the initial chunk of the tempdbspace that you are creating	
<i>size</i>	Indicates, in kilobytes, the size of the initial chunk of the new tempdbspace	
WITH_CHECK	Returns an error if <i>pathname</i> does not exist	

Argument	Description	Restrictions
PRINT PARTITION	Prints a tblspace report for given table or fragment whose name and database you specify when you execute oncheck at the command line	
PRINT PARTITION FULL	Includes index-specific information and page-allocation information by page type for dbspaces	
<i>partition_number</i>	The partition number	

RENAME SPACE SQL Admin API Command

Renames a dbspace, blobspace, sbspace, or extspace.

```

▶▶ [ADMIN] (—"RENAME SPACE" —,—"space" —,—"new_name" —)
   [TASK]

```

Argument	Description	Restrictions
<i>space</i>	The name of the dbspace, blobspace, sbspace, or extspace to rename	
<i>new_name</i>	The new name of the space	

RESET SYSADMIN SQL Admin API Command

Moves the **sysadmin** database to a specified dbspace

```

▶▶ [ADMIN] (—"RESET SYSADMIN" —,—"dbspace" —)
   [TASK]

```

Argument	Description	Restrictions
<i>dbspace</i>	The name of the dbspace to which this command applies	

SCHEDULER SQL Admin API Command

Starts or stops the scheduler

```

▶▶ [ADMIN] (—"SCHEDULER" [SHUTDOWN] [START])
   [TASK]

```

Argument	Description	Restrictions
SCHEDULER SHUTDOWN	Stops the scheduler and deallocates all its resources	
SCHEDULER START	Starts the scheduler	

SET CHUNK SQL Admin API Command

Changes the status of a dbspace, blobspace, or sbspace to online or offline. The chunk must be in a mirrored pair or a non-primary chunk within a noncritical dbspace.

```

▶▶ [ADMIN] (—"SET CHUNK" [OFFLINE] ,—"path" —,—"offset" —)
   [TASK]
  
```

Argument	Description	Restrictions
SET CHUNK OFFLINE	Change the status of the specified chunk to offline	
SET CHUNK ONLINE	Change the status of the specified chunk to online	
<i>path</i>	Indicates the disk partition or unbuffered device of the chunk	
<i>offset</i>	Indicates, in kilobytes, the offset into the disk partition or unbuffered device to reach the chunk	

SET DATASKIP SQL Admin API Command

Specifies that the database server can skip a designated dbspace.

```

▶▶ [ADMIN] (—"SET DATASKIP" [OFF] ,—"dbspace" —)
   [TASK]
  
```

Argument	Description	Restrictions
SET DATASKIP OFF	Stop skipping the specified dbspace	
SET DATASKIP ON	Begin skipping the specified dbspace	
<i>dbspace</i>	The name of the dbspace to which this command applies	

SET INDEX COMPRESSION SQL Admin API Command

Modifies the level at which two partially used index pages are merged. The pages are merged if the data on those pages totals a set level.

```

▶▶ [ADMIN]
   [TASK]
▶▶ (—"SET INDEX COMPRESSION" —,—"partition_number" —,—"compression_level" —)
  
```

Argument	Description	Restrictions
<i>partition_number</i>	The partition number to compress	

Argument	Description	Restrictions
<i>compression_level</i>	Valid values for the level are low, med (medium), high, and default. The system default value is med.	

SET ONCONFIG SQL Admin API Commands

Sets ONCONFIG file values

```

▶▶ [ADMIN] (—"SET ONCONFIG" [PERMANENT" | " | MEMORY"]
   [TASK]
▶▶ ,—"config_param" —,—"value" —)

```

Argument	Description	Restrictions
SET ONCONFIG MEMORY	The default. Changes the value of the specified configuration parameter in memory.	
SET ONCONFIG PERMANENT	Changes the value of the specified configuration parameter in the configuration file.	
<i>config_param</i>	Specifies a configuration parameter. For a list of dynamically configurable configuration parameters, see the topic on onmode -wf, -wm in the <i>IBM Informix Administrator's Reference</i> .	
<i>value</i>	Specifies the value of the configuration parameter.	

SET SBSPACE ACESSTIME SQL Admin API Command

Tracks the time of access to all smart large objects stored in the sbspace

```

▶▶ [ADMIN] (—"SET SBSPACE ACESSTIME" [OFF" | ON"] ,—"sbspace" —)
   [TASK]
▶▶

```

Argument	Description	Restrictions
SET SBSPACE ACESSTIME OFF	Turns off tracking of access time	
SET SBSPACE ACESSTIME ON	Turns on tracking of access time to all smart large objects stored in the sbspace	
<i>sbspace</i>	The name of the sbspace to which this command applies	

SET SBSPACE AVG_LO_SIZE SQL Admin API Command

Specifies an expected average sbspace size

```

▶▶ [ADMIN] (—"SET SBSPACE AVG_LO_SIZE" —,—"sbspace" —)
   [TASK]
▶▶

```

Argument	Description	Restrictions
<i>sbspace</i>	The name of the sbspace to which this average applies	

SET SBSPACE LOGGING SQL Admin API Command

Specifies that the database server logs changes to the user data area of the sbspace

```

▶▶ [ADMIN] (—"SET SBSPACE LOGGING [OFF]" ,—"sbspace" —)
   [TASK] [ON]

```

Argument	Description	Restrictions
SET SBSPACE LOGGING OFF	Specifies that the database server does not log changes to the user data area of the sbspace	
SET SBSPACE LOGGING ON	Specifies that the database server logs changes to the user data area of the sbspace	
<i>sbspace</i>	The name of the sbspace to which the command is applied	

SET SQL USER TRACING SQL Admin API Command

Sets the mode of global SQL tracing for a specified user session

```

▶▶ [ADMIN] (—"SET SQL USER TRACING [CLEAR]" ,—"session_id" —)
   [TASK] [OFF]
           [ON]

```

Argument	Description	Restrictions
SET SQL USER TRACING CLEAR	Clear user tracing flags for the specified user session so that it adheres to the global tracing policy	
SET SQL USER TRACING OFF	Disables SQL tracing for a user session even if the global mode is ON	
SET SQL USER TRACING ON	Enables user SQL tracing for a user session. Even if the global tracing mode is OFF, SQL statements for this user session are traced.	
<i>session_id</i>	The session ID to which the command applies	

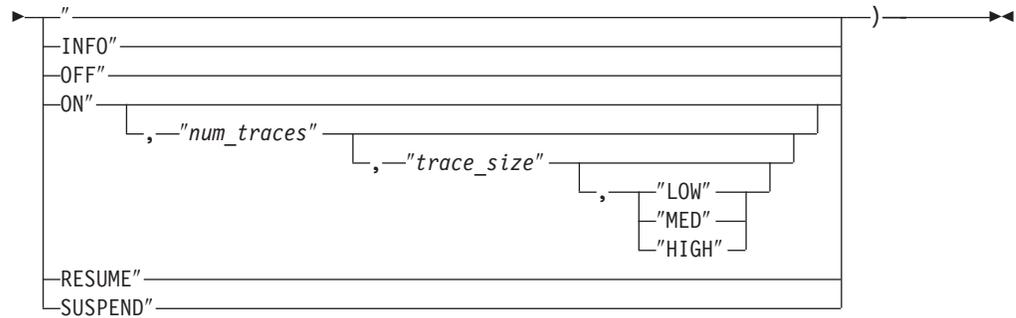
SET SQL TRACING SQL Admin API Command

Sets global SQL tracing

```

▶▶ [ADMIN] (—"SET SQL TRACING" —)
   [TASK]

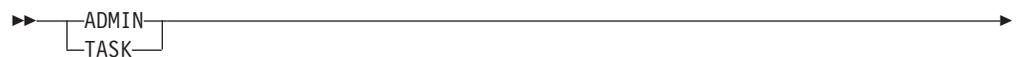
```



Argument	Description	Restrictions
SET SQL TRACING	Enables global SQL tracing with default values	
SET SQL TRACING INFO	Return the state of global SQL tracing	
SET SQL TRACING OFF	Turns off global SQL tracing	
SET SQL TRACING ON	Sets global SQL tracing	
SET SQL TRACING RESUME	Restarts SQL tracing when it has been suspended	
SET SQL TRACING SUSPEND	Pauses SQL tracing without deallocating any resources.	
<i>num_traces</i>	The number of SQL statements to trace	
<i>trace_size</i>	The number of kilobytes for the size of the trace buffer. If this buffer size is exceeded, the database server discards saved data	
LOW	Captures statement statistics, statement text, and statement iterators	
MED	Captures all of the information included in low-level tracing, plus table names, the database name, and stored procedure stacks	
HIGH	Captures all of the information included in medium-level tracing, plus host variables	

SET SQL TRACING DATABASE SQL Admin API Commands

Sets SQL tracing for a database



► (—"SET SQL TRACING DATABASE" —,—"database_name"—)

- "ADD" —,—"database_name"—
- "CLEAR" —
- "LIST" —
- "REMOVE" —,—"database_name"—

Argument	Description	Restrictions
ADD	Specify tracing for one or more databases, rather than for all databases. The default is all databases. Specify up to six arguments in a statement. The maximum number of database names that can be set is 16.	
CLEAR	Clear all databases from the list of database being traced. Returns tracing back to the default of all databases.	
LIST	List the databases which are currently being traced.	
<i>database_name</i>	The name of the database to which the command applies	
REMOVE	Remove a single database from the list of databases being traced	

SET SQL TRACING SESSION SQL Admin API Command

Sets SQL tracing for the current session. This command behaves in the same manner as SET SQL USER TRACING.

► ADMIN TASK

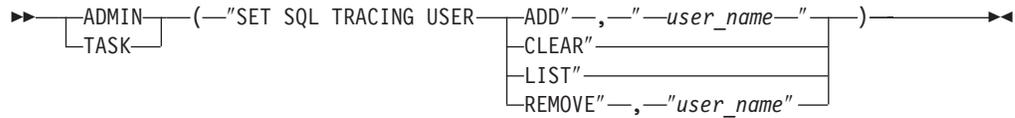
► (—"SET SQL TRACING SESSION" —,—"CLEAR" —,—"OFF" —,—"ON" —,—"current_session_id"—,—"session_id"—)

Argument	Description	Restrictions
CLEAR	Clear user tracing flags for the specified user session so that it adheres to the global tracing policy	
<i>current_session_id</i>	The current session ID to which the command applies	
OFF	Disables SQL tracing for a user session if the global mode is ON	
ON	Enables user SQL tracing for a user session. Even if the global tracing mode is OFF, SQL statements for this user session are traced.	

Argument	Description	Restrictions
<i>session_id</i>	The session ID to which the command applies	

SET SQL TRACING USER SQL Admin API Commands

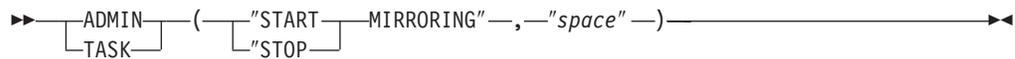
Sets SQL tracing for a user



Argument	Description	Restrictions
ADD	Add a user to the user tracing list.	
CLEAR	Remove all users from the user tracing list.	
LIST	List all users on the user trace list.	
REMOVE	Remove a user from the user trace list.	
<i>user_name</i>	The name of the user to trace.	

MIRRORING SQL Admin API Command

Sets mirroring for a specified dbspace, blobspace, or sbspace



Argument	Description	Restrictions
START MIRRORING	Adds mirroring for an existing dbspace, blobspace, or sbspace	
STOP MIRRORING	Ends mirroring for an existing dbspace, blobspace, or sbspace	
<i>space</i>	Names the blobspace, dbspace, or sbspace that you want to mirror	

The ONMODE Initial Command

As the Table 6-1 on page 6-6 list of mnemonic initial commands shows, the **ADMIN** and **TASK** functions accept a descriptive word or phrase as the initial *command* argument for many operations. To emulate some operations of the **onmode** utility, however, these functions accept "ONMODE" as the initial argument. Subsequent arguments that follow "ONMODE" are the corresponding **onmode** key letter flags, but with no preceding hyphen (-) symbol, and any additional input specifications. For a few **onmode** operations, either a mnemonic

command (such as "ADD MEMORY") or an "ONMODE" command can begin the argument list of SQL Admin API function calls that both have the same effect.

Each argument that follows "ONMODE" must appear in the same order as the corresponding specification in the equivalent **onmode** command line, but as a quoted string. Arguments that are literal numbers do not require quotation marks as delimiters. Just as with the mnemonic initial commands that support additional arguments, the arguments that follow "ONMODE" are interpreted positionally, rather than as named input parameters.

The **ADMIN** and **TASK** functions do not support interactivity; for this reason, the "-y" flag of some **onmode** command-line operations is unavailable.

In the following example of "ONMODE" as the initial SQL Admin API command argument, the EXECUTE FUNCTION statement is equivalent to the **onmode -p -3 encrypt** command:

```
EXECUTE FUNCTION ADMIN("ONMODE", "p", "-3", "encrypt");
```

If this SQL statement executes successfully, it drops three encryption virtual processors of the **Encrypt** VP class.

The next example is equivalent to the **onmode -d secondary cleveland** command:

```
EXECUTE FUNCTION ADMIN("ONMODE", "d", "secondary", "cleveland");
```

If this function executes successfully, it designates the Dynamic Server instance called **cleveland** as a secondary server for data replication. Just as with the **onmode** utility, this operation is valid only if the database server from which this command is issued is in quiescent mode, and **cleveland** cannot be the value of its DBSERVERNAME configuration parameter.

The statement in the next example is equivalent to the **onmode -M 4096** command:

```
EXECUTE FUNCTION ADMIN("ONMODE", "M", "4096KB");
```

If this function call is successful, it dynamically resets DS_TOTAL_MEMORY, the maximum amount of memory available for parallel queries, to 4096 Kb. (This example explicitly specifies kilobytes ("KB") in the last argument, although kilobytes are the default units for all SQL Admin API arguments that specify memory or storage sizes without providing units.)

Arguments to the **ADMIN** or **TASK** function that follow the initial "ONMODE" argument follow the same rules as for any other **ADMIN** or **TASK** *command*. Only the "ONMODE" command, however, has the same name as the command-line utility that it emulates, and only the "ONMODE" command uses the same case-sensitive keyletters as the corresponding command-line utility options to define the requested operation.

Options to the **onmode** utility that are not emulated by SQL Admin API commands that use "ONMODE" as the initial argument to **ADMIN** or to **TASK** are documented in the *IBM Informix Dynamic Server Administrator's Guide* and in the *IBM Informix Administrator's Reference*.

The following table lists SQL Admin API "ONMODE" commands that emulate operations of the **onmode** utility. See the *IBM Informix Administrator's Reference* for additional information about the keyletters and parameters of the **onmode** utility.

Table 6-2. ONMODE SQL Admin API Commands

Command	Effect of Successful Operation	Corresponding onmode Command Line
ONMODE	Add memory	"onmode a: Add a shared-memory segment" on page 6-30
ONMODE	Do a checkpoint of a specified type	"onmode c: Force a checkpoint" on page 6-30
ONMODE	Tune B-tree scanner	"onmode C: Control the B-tree scanner" on page 6-31
ONMODE	Configure a high availability environment	"onmode d: Set data-replication types" on page 6-32
ONMODE	Set PDQ	"onmode D: Set PDQ" on page 6-33
ONMODE	Configure or flush shared statement cache	"onmode e: Change usage of the SQL statement cache" on page 6-33
ONMODE	Free unused memory	"onmode F: Free unused memory segments" on page 6-34
ONMODE	Switch to administration mode	"onmode j: Switch database server to administration mode" on page 6-34
ONMODE	Switch to next logical log	"onmode l: Switch to next logical log" on page 6-35
ONMODE	Switch to multi-user mode	"onmode m: Switch to multi-user mode" on page 6-35
ONMODE	Resize decision support memory	"onmode M: Change decision-support memory" on page 6-35
ONMODE	Unlock resident portion of shared memory	"onmode n: Unlock resident memory" on page 6-36
ONMODE	Override the WAIT mode of the ONSPACESDOWN parameter	"onmode O: Override a space down that blocks a checkpoint" on page 6-36
ONMODE	Add or remove a virtual processor	"onmode p: Add or remove virtual processors" on page 6-36
ONMODE	Reset the limit on decision-support queries	"onmode Q: Set maximum number for decision-support queries" on page 6-39
ONMODE	Lock the resident portion of shared memory	"onmode r: Lock memory resident" on page 6-39
ONMODE	Reset the limit on decision support scans	"onmode S: Specify maximum number of decision support scans" on page 6-39
ONMODE	Reset statement cache attributes	"onmode W: Reset statement cache attributes" on page 6-40
ONMODE	Dynamically update certain configuration parameters	"onmode {wf wm}: Dynamically update a configuration parameter" on page 6-40
ONMODE	Dynamically update how bufferpools are flushed	"onmode wm: Dynamically change LRU tuning status" on page 6-41
ONMODE	Switch the SET EXPLAIN feature on or off, or save query optimizer plans without executing the queries	"onmode Y: Dynamically change SET EXPLAIN for a user session" on page 6-43
ONMODE	Terminate a user session	"onmode z: Terminate a user session" on page 6-44
ONMODE	Heuristically complete a transaction	"onmode Z: Terminate a distributed transaction ungracefully" on page 6-44

Calling Syntax of the "ONMODE" Commands

This section shows the syntax for invoking the commands of the SQL Admin API function calls whose initial command is "ONMODE" and whose argument lists begin with each of the keyletters that Table 6-2 identifies as corresponding to a supported command-line option of the **onmode** utility. These commands are listed in the order of their first keyletter argument.

onmode a: Add a shared-memory segment

Syntax:

```

▶▶ ADMIN (—"ONMODE"—,—"a"—,—"size"—)
└TASK┘

```

Argument	Effect	Key Considerations
<i>size</i>	Specifies the size (in kilobytes) of a new virtual shared-memory segment.	The value of <i>size</i> must be a positive integer that does not exceed the operating-system limit on the size of shared-memory segments.

Ordinarily, you do not need to add segments to the virtual portion of shared memory because the database server automatically adds segments as they are needed. As segments are added, however, the database server might reach the operating-system limit for the maximum number of segments before it acquires the memory that it needs. This situation typically occurs when SHMADD is set to a value so small that the database server exhausts the number of available segments before it acquires sufficient memory for some operation.

If you use this command to add a segment that is larger than the size specified by SHMADD, you can avoid exhausting the operating-system limit for segments but still meet the need that the database server has for additional memory.

onmode c: Force a checkpoint

Syntax:

```

▶▶ ADMIN (—"ONMODE"—,—"c"—,—"
└TASK┘
┌hard┐
│block│
│norm │
└unlock┘
")

```

This command forces a checkpoint that flushes the buffers to disk. You can use the **c** option to force a checkpoint if the most recent checkpoint record in the logical log was preventing the logical-log file from being freed (status U-B-L).

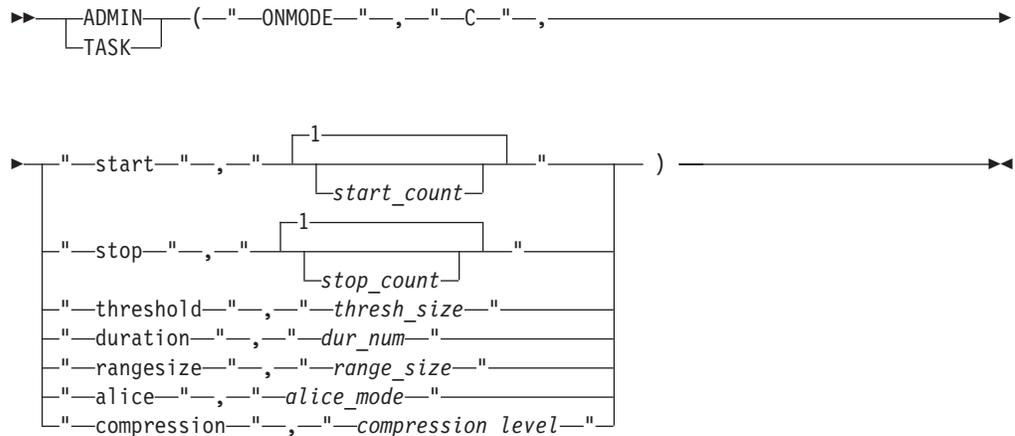
The last argument specifies the type of checkpoint:

- **block** prevents the database server from processing any transactions. While the database server is blocked, users cannot access it, except in read-only mode. Use this option to perform an external backup on Dynamic Server.
- **hard** forces a blocking checkpoint. If you specify no argument after “**c**”, this is the default.
- **norm** forces a nonblocking checkpoint.
- **unblock** unblocks the database server. When the database server is unblocked, data transactions and normal database server operations can resume. Use this option after you complete an external backup on Dynamic Server.

For more information, see the *IBM Informix Backup and Restore Guide*.

onmode C: Control the B-tree scanner

Syntax:



Argument	Effect	Key Considerations
<i>start_count</i>	Starts additional B-tree scanner threads.	If <i>start_count</i> is not specified, 1 additional thread is started. There is no limit on the number of scanner threads that can be specified.
<i>stop_count</i>	Stops B-tree scanner threads.	If <i>stop_count</i> is not specified, a single thread is stopped. Stopping all index scanners prevents all index cleaning. If you specify a larger <i>stop_count</i> value than the number of threads than are running, no error is issued, but all scanner threads are stopped.
<i>thresh_size</i>	Sets the minimum number of deleted items an index must encounter before an index is placed on the hot list.	Once all indexes above the threshold have been cleaned and there is no other work for the B-tree scanner to do, the indexes below the threshold are added to the hot list.
<i>dur_num</i>	The number of seconds that the hot list is valid.	After this number of seconds expires, the hot list will be rebuilt by the next available B-tree scanner thread, even if unprocessed items are on the list. Scanners currently processing requests are not interrupted.
<i>range_size</i>	Determines the size of an index before index range cleaning is enabled.	A size of -1 can be used to disable range scanning.
<i>alice_mode</i>	Sets the system's alice mode.	Valid integer values range from 0 (OFF) to 12.
<i>compression_level</i>	For a database server instance, modifies the level at which two partially used index pages are merged. The pages are merged if the data on those pages totals a set level.	Valid values for the level are low, med (medium), high, and default. The system default value is med.

This command controls the B-tree scanner for cleaning indexes of deleted items. There is no limit to the number of threads that can run at one time. There is a limit of 32 threads, however, that can be started at one time. If, for example, you wanted 60 B-tree scanner threads to run, you could execute two commands:

```
EXECUTE FUNCTION ADMIN ("ONMODE", "C", "start" "30");
EXECUTE FUNCTION ADMIN ("ONMODE", "C", "start" "30");
```

The following command stops all of these threads

```
EXECUTE FUNCTION ADMIN ("ONMODE", "C", "start" "30000");
```

No error is issued when the *stop_count* value is greater than the number of running threads.

The B-tree scanner has statistical information which tracks index efficiency and how much extra work the index currently places on the server. Based on the amount of extra work the index has accomplished because of committed deleted index items, the B-tree scanner develops an ordered list of indexes that have caused the server to do extra work, called the hot list. The index causing the highest amount of extra work is cleaned first and the rest of the indexes are cleaned in descending order. The DBA can allocate cleaning threads dynamically, thus allowing for configurable workloads.

onmode d: Set data-replication types

Syntax:

```
ADMIN TASK (—"ONMODE"—,—"d"—,—"standard"—,—"primary"—,—"dbserver"—,—"secondary"—,—"dbserver"—)
```

Argument	Effect	Key Considerations
<i>dbserver</i>	Identifies the database server name of the primary or secondary database server	The <i>dbserver</i> argument must correspond to the DBSERVERNAME parameter in the ONCONFIG file of the intended secondary database server. It should <i>not</i> correspond to one of the database servers that the DBSERVERALIASES parameter specifies.

Use this command to set the High-Availability Data Replication type as standard, primary, or secondary. You can use the **primary** and **secondary** options only when the database server is in quiescent mode. You can use the **standard** option when the database server is in quiescent, online, or read-only mode.

The *dbserver* argument of the other database server in the data-replication pair and the type of a database server (standard, primary, or secondary) is preserved after reinitialization of shared memory.

Using the "standard" Option

The **standard** option drops the connection between database servers in a data replication pair (if one exists) and sets the database server type of the current database server to standard. This option does not change the mode or type of the other database server in the pair.

Using the "primary" and "dbserver" Options

The **primary** and *dbserver* arguments set the database server type to primary and attempts to connect with the database server that *dbserver* specifies. If the connection is successful, data replication is turned on. The primary database server goes into online mode, and the secondary database server goes into read-only mode. If the connection is not successful, the database server comes to online mode, but data replication is not turned on.

Using the "secondary" and "dbserver" Option

The **secondary** and *dbserver* arguments set the database server type to secondary and attempts to connect with the database server that *dbserver* specifies. If the connection is successful, data replication is turned on. The primary database server goes online, and the secondary database server goes into read-only mode. If the connection is not successful, the database server comes to read-only mode, but data replication is not turned on.

onmode D: Set PDQ

Syntax:

```

>> ADMIN (—"ONMODE"—,—"D"—,—"max_priority"—) <<<
    |
    | TASK
  
```

Argument	Effect	Key Considerations
<i>max_priority</i>	Updated value of MAX_PDQPRIORITY in memory	This value must be an unsigned integer between 0 and 100.

Specify *max_priority* as a factor to temper user requests for PDQ resources. This command can override the MAX_PDQPRIORITY configuration parameter setting while the database server is online. The new value affects only the current database server session, but is not recorded in the ONCONFIG file. If you shut down and restart the database server, the MAX_PDQPRIORITY value reverts to its setting in the ONCONFIG file.

onmode e: Change usage of the SQL statement cache

Syntax:

```

>> ADMIN (—"ONMODE"—,—"e"—,—"ENABLE"—,—"FLUSH"—,—"OFF"—,—"ON"—) <<<
    |
    | TASK
  
```

This command operates on the SQL statement cache. It can specify any one of four modes for user sessions of the current database server session:

- **ENABLE**

This enables the SQL statement cache, if it was disabled before the command successfully executed. Individual user sessions can use the statement cache only after they perform either of the following actions:

- Set the environment variable **STMT_CACHE** to 1
- Execute the SQL statement **SET STATEMENT CACHE ON**.

For more information, see the material on improving query performance in the *IBM Informix Performance Guide*

- **FLUSH**

This mode flushes the statements that are not in use from the SQL statement cache, which remains enabled. After the cache is flushed, the **onstat -g ssc ref_cnt** field shows 0.

- **OFF**

This mode turns off the SQL statement cache, so that no statements are cached.

- **ON**

In this mode, all statements are cached unless the user turns it off by one of the following actions:

- Use this command to specify the OFF mode
- Set the environment variable **STMT_CACHE** to 0
- Execute the SQL statement SET STATEMENT CACHE OFF statement of SQL.

This command cannot modify the STMT_CACHE setting in the configuration file, but the last argument overrides that setting (or the default value, if STMT_CACHE is not set). Any changes to the statement cache behavior that you make with this command are in effect for the current database server session only. When you restart the database server, it uses the setting of the STMT_CACHE parameter in the ONCONFIG file, or else uses no statement cache, if STMT_CACHE is undefined in that file.

onmode F: Free unused memory segments

Syntax:

```
►► [ADMIN TASK] ( ["-ONMODE-", "-F-"] )
```

When you execute this command, the memory manager examines each memory pool for unused memory. The memory manager immediately frees unused blocks of memory that it locates. After the memory manager checks each memory pool, it begins checking memory segments and frees any that the database server no longer needs.

Recommendation: Run this command from an operating-system scheduling facility regularly and after the database server performs any function that creates additional memory segments, including large index builds, sorts, or backups.

Running this command causes a significant degradation of performance for any users that are active when you execute the utility. Although the execution time is brief (1 to 2 seconds), degradation for a single-user database server can reach 100 percent. Systems with multiple CPU virtual processors experience proportionately less degradation.

To confirm that **onmode** freed unused memory, check your message log. If the memory manager frees one or more segments, it displays a message that indicates how many segments and bytes of memory were freed.

onmode j: Switch database server to administration mode

Syntax:

```
►► [ADMIN TASK] ( ["-ONMODE-", "-j-"] )
```

This command uses the **j** option as its last argument to put the database server into the administration mode. Only user **informix** can connect to the database in

this mode, but the configuration parameter `ADMIN_USER_MODE_WITH_DBSA` can be set to enable members of the **DBSA** group, in addition to user **informix**, to connect to the server in this mode.

When the server is changed to administration mode, all sessions lose their connection to the database server except for sessions of the following users:

- user **informix**
- users in the **DBSA** group
- and users who are identified in `ADMIN_MODE_USERS` settings.

onmode I: Switch to next logical log

Syntax:

```

>> ADMIN (—"ONMODE"—,—"I"—)
   TASK
  
```

This command switches the current logical-log file to the next logical-log file.

The following example moves the logical log out of the **root** chunk

```

SELECT task("onmode", "I") FROM sysmaster:syslogfil
WHERE chunk = 1 AND sysmaster:bitval(flags,"0x02")>0;
  
```

For information on switching to the next logical-log file, see the section on managing logical-log files in the *IBM Informix Administrator's Guide*.

onmode m: Switch to multi-user mode

Syntax:

```

>> ADMIN (—"ONMODE"—,—"m"—)
   TASK
  
```

This command brings the database server online from quiescent mode or from administration mode.

onmode M: Change decision-support memory

Syntax:

```

>> ADMIN (—"ONMODE"—,—"M"—,—"size"—)
   TASK
  
```

Argument	Effect	Key Considerations
<i>size</i>	The new size limit (in kilobytes) that overrides the current value of <code>DS_TOTAL_MEMORY</code>	This <i>size</i> specifies the maximum amount of memory available for parallel queries.

These options allow you to override the limit set by the `DS_TOTAL_MEMORY` configuration parameter while the database server is online. The new values affect

only the current instance of the database server; the values are not recorded in the ONCONFIG file. If you shut down and restart the database server, the values of the parameter reverts to the values in the ONCONFIG file. For more information about these configuration parameters, see *IBM Informix Administrator's Reference*.

The *size* that you specify has a platform-dependent upper limit. The value for 32-bit systems must be an unsigned integer between $128 \times \text{DS_MAX_QUERIES}$ and 1,048,576. On 64-bit systems, the limit is generally higher and varies with the operating system. On HP 9000 platforms, for example, the maximum value is 4,294,967,296.

For more information, see the *IBM Informix Performance Guide*.

onmode n: Unlock resident memory

Syntax:

```

▶▶ [ADMIN] (—"ONMODE"—,—"n"—)
   [TASK]

```

This command ends forced residency of the resident portion of shared memory. This command does not affect the value of RESIDENT, the forced-residency parameter in the ONCONFIG file.

onmode O: Override a space down that blocks a checkpoint

Syntax:

```

▶▶ [ADMIN] (—"ONMODE"—,—"O"—)
   [TASK]

```

This command overrides the WAIT mode of the ONDBSPACEDOWN configuration parameter. Use this command only in the following circumstances:

- ONDBSPACEDOWN is set to WAIT.
- A disabling I/O error occurs that causes the database server to block all updating threads.
- You cannot or do not want to correct the problem that caused the disabling I/O error.
- You want the database server to mark the disabled dbspace as down and continue processing.

When you execute this command, the database server marks the dbspace that is responsible for the disabling I/O error as down, completes a checkpoint, and releases any blocked threads.

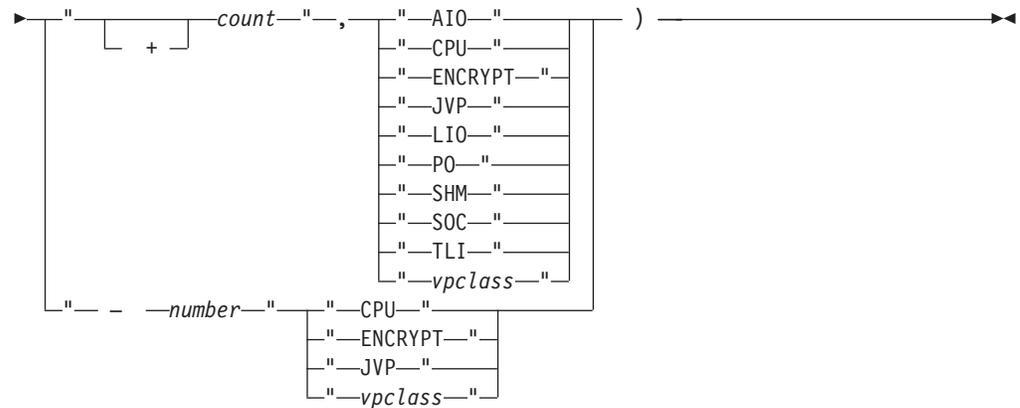
onmode p: Add or remove virtual processors

Syntax:

```

▶▶ [ADMIN] (—"ONMODE"—,—"p"—,
   [TASK]

```



Argument	Effect	Key Considerations
<i>count</i>	This integer specifies how many virtual processors of the specified class to add or to remove	This command adds VPs if <i>count</i> is positive or unsigned, or removes VPs if <i>count</i> is negative. For more details, see “Adding and Dropping Virtual Processors.”
AIO	Performs nonlogging disk I/O to cooked disk spaces	Also performs nonlogging I/O to raw disk spaces if kernel asynchronous I/O (KAIO) is not used.
CPU	Runs all session threads and some system threads	Recommendation: The number of CPU VPs not be greater than the number of physical processors.
ENCRYPT	Executes column-level encryption and decryption routines	Specify more ENCRYPT virtual processors if you have multiple encrypted columns.
JVP	Executes Java user-defined routines in the Java Virtual Machine (JVM)	Specify more JVPs if you are running many Java UDRs.
LIO	Writes to the logical-log files if they are in cooked disk space	Use two LIO virtual processors only if the logical logs are in mirrored dbspaces. The database server allows a maximum of two LIO virtual processors.
PIO	Writes to the physical log if it is in cooked disk space	Use two PIO virtual processors only if the physical log is in a mirrored dbspace. The database server allows a maximum of two PIO virtual processors.
SHM	Performs shared-memory communication	You can use the SHM virtual processor even if the database server is not configured for shared-memory communication.
SOC	Uses sockets to perform network communications	You can use the SOC virtual processor only if the database server is configured for network connections through sockets.
STR	Performs stream pipe connections	
TLI	Uses the Transport Layer Interface (TLI) to perform network communication	You can use the TLI virtual processor only if the database server is configured for network connections through TLI.
<i>vpclass</i>	Specifies a user-defined virtual processor class	The VPCLASS parameter in the ONCONFIG file defines the user-defined virtual-processor class.

You can use this command only when the database server is in online mode, and you can change the number of virtual processors in only one VP class at a time.

Adding and Dropping Virtual Processors

If you are removing virtual processors, the absolute value of *count* cannot exceed the actual number of processors of the specified type. If you are adding virtual processors, the maximum number depends on the operating system.

The following rules about adding or dropping virtual processors apply:

- You can add but not drop virtual processors of the AIO, PIO, LIO, TLI, SHM, SOC, and STR classes.
- You cannot add or drop virtual processors of the OPT, ADM, ADT, and MSC classes. The database server adds these automatically.
- You can add or drop virtual processors of the CPU, ENCRYPT, JVP, and user-defined (*vpclass*) classes.
- On Windows, you can add a virtual processor of any class, but you cannot drop virtual processors.

CPU Virtual Processors

The number of CPU VPs should not exceed the number of physical processors on your system, but no error is issued if the *count* ignores this recommendation. If KAIO is used, CPU VPs perform I/O to raw disk spaces, including I/O to physical and logical logs, and run threads for KAIO (where available) or a single poll thread. The database server uses the number of CPU VPs to allocate resources for parallel database queries (PDQ). If you drop CPU VPs, your queries run significantly slower. After you use this ONMODE command to change the number of CPU VPs, the **Reinit** field in output from the **onstat -g mgm** can show how many queries are waiting for other queries to complete. See also the *IBM Informix Performance Guide* for additional information about performance implications of the CPU VP class.

User-Defined Virtual Processors

Use the VPCLASS parameter in the ONCONFIG file to define a user-defined virtual-processor class. You can include several VPCLASS parameter definitions in your ONCONFIG file. Each VPCLASS parameter defines one class of virtual processors.

If you are running many UDRs on UNIX, you can in some cases improve performance by calling the TASK("ONMODE", "p", "+count", "vpclass") function to add one or more user-defined virtual processors.

On Windows, the value of *count* must be 1 if you specify a *vpclass* with this command, because you can have no more than one user-defined virtual processor class at a time.

Dropping Virtual Processors Automatically

Table 6-3 shows the virtual processors that the database server starts automatically. You cannot add or drop these virtual processors with the **onmode -p** command. To drop these virtual processors, shut down and restart the database server.

Table 6-3. Virtual-Processor Classes That the Database Server Starts Automatically

VP Class	Description
ADM	Performs administrative functions
ADT	Runs auditing processes. The database server starts one virtual processor in the audit class when you turn on audit mode by setting the ADTMODE parameter in the ONCONFIG file.
MSC	Services requests for system calls that require a large stack. The database server starts this virtual processor automatically.
OPT	Performs I/O to the optical disk. The database server starts one OPT virtual processor when you use the Optical Subsystem.

For more information, see the section on using virtual processors in the *IBM Informix Administrator's Guide*.

onmode Q: Set maximum number for decision-support queries

Syntax:

```

▶▶ [ADMIN] (—"ONMODE"—,—"Q"—,—"queries"— )
   [TASK]
  
```

Argument	Effect	Key Considerations
<i>queries</i>	Changes the value of DS_MAX_QUERIES	Specifies the maximum number of concurrently executing parallel queries. This value must be an unsigned integer between 1 and 8,388,608.

For information on parameters used for controlling PDQ, see the *IBM Informix Performance Guide*.

onmode r: Lock memory resident

Syntax:

```

▶▶ [ADMIN] (—"ONMODE"—,—"r"— )
   [TASK]
  
```

This command starts forced residency of the resident portion of shared memory. It does not affect the value of RESIDENT, the forced-memory parameter in the ONCONFIG file.

Important: Set the RESIDENT parameter to 1 before you invoke this command.

For information on using the forced-residency parameter to turn residency on or off for the next time that you restart the database server, see the section on managing shared memory in the *IBM Informix Administrator's Guide*.

onmode S: Specify maximum number of decision support scans

Syntax:

```

▶▶ [ADMIN] (—"ONMODE"—,—"S"—,—"scans"— )
   [TASK]
  
```

Argument	Effect	Key Considerations
<i>scans</i>	Specifies the new value of DS_MAX_SCANS	This value must be an unsigned integer between 10 and 1,048,576.

Use this command to reset the maximum number of concurrently executing parallel scans.

For information on parameters used for controlling PDQ, see the *IBM Informix Performance Guide*.

onmode W: Reset statement cache attributes

Syntax:

```

▶▶ [ADMIN] ( ["ONMODE" , ["W" ,
└─TASK─┘
└─"STMT_CACHE_HITS" , ["hits"
└─"STMT_CACHE_NOLIMIT" , ["value"
)

```

Argument	Effect	Key Considerations
<i>hits</i>	Specifies the number of hits (references) to a statement before it is fully inserted in the SQL statement cache. Set <i>hits</i> to 1 or more to exclude ad hoc queries from entering the cache.	You can only increase or reset the value of STMT_CACHE_HITS. The new value displays in the #hits field of the onstat-g ssc output.
<i>value</i>	Controls whether statements are inserted in the SQL statement cache.	This can reset STMT_CACHE_NOLIMIT to either 0 or 1. If <i>value</i> = 0, the database server inserts no statements into the cache. If <i>value</i> = 1, the database server always inserts statements in the cache.

This command can dynamically reset the STMT_CACHE_HITS or STMT_CACHE_NOLIMIT configuration parameters of the statement cache.

If *hits* is equal to 0, the database server inserts all qualified statements and their memory structures in the cache. If *hits* is greater than 0 and the number of times the SQL statement has been executed is less than STMT_CACHE_HITS, the database server inserts *key-only* entries in the cache. The database server inserts qualified statements in the cache after the specified number of hits has occurred for the statement.

If none of the queries are shared, set the STMT_CACHE_NOLIMIT *value* to 0 to prevent the database server from allocating a large amount of memory for the statement cache.

onmode {wf | wm}: Dynamically update a configuration parameter

Use this command to dynamically change the setting of certain configuration parameters that can take a single value, or that can have multiple entries, each specifying a single value, in the ONCONFIG file.

- If the first argument after the ONMODE command is **wf**, then the setting of the specified configuration parameter is changed in the configuration file.
- If the first argument after the ONMODE command is **wm**, then the setting of the specified configuration parameter is changed in memory.

Syntax:

```

▶▶ [ADMIN] ( ["ONMODE" , [" wf
└─TASK─┘ └─wm─┘
, ["config_param=value"
)

```

Argument	Effect	Key Considerations
<i>config_param=value</i>	<p>This specifies a configuration parameter and its new value, separated by the = symbol, as a single argument. The following configuration parameters can be specified:</p> <ul style="list-style-type: none"> • ADMIN_MODE_USERS • AUTO_AIOVPS • AUTO_CKPTS • DS_MAX_QUERIES • DS_MAX_SCANS • DS_NONPDQ_QUERY_MEM • DS_TOTAL_MEMORY • EXPLAIN_STAT • IFX_EXTEND_ROLE • LISTEN_TIMEOUT • LOG_INDEX_BUILDS • MAX_INCOMPLETE_CONNECTIONS • MAX_PDQPRIORITY • ONLIDX_MAXMEM • RESIDENT • RTO_SERVER_RESTART • SDS_ENABLE • SDS_TIMEOUT • TEMPTAB_NOLOG • USELASTCOMMITTED • VP_MEMORY_CACHE_KB 	The <i>value</i> must be valid for the specified parameter.

For information about the effects of these configuration parameters and the set or range of values can follow the = symbol for each configuration parameter, see the entry for that parameter in *IBM Informix Administrator's Reference*. The section "onmode wm: Dynamically change LRU tuning status" illustrates the syntax of an ONMODE command that uses "WF" or "WM" as its first argument after the ONMODE command, but accepts additional arguments after the "parameter=value" specification to reset the configuration parameter.

onmode wm: Dynamically change LRU tuning status

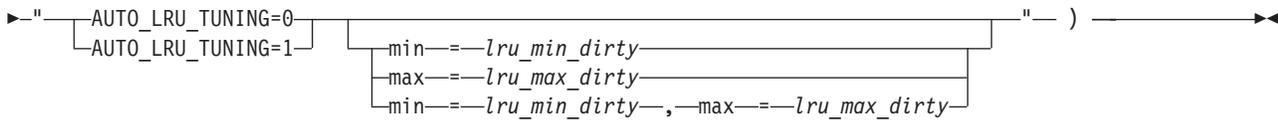
The following section provides the syntax of the ONMODE command for the AUTO_LRU_TUNING configuration parameter, which specifies whether automatic LRU tuning (for flushing each buffer pool between checkpoints) is enabled when the database server is restarted. This command requires **wm** as the argument that immediately follows the ONMODE command string, but the third argument can include additional specifications after the AUTO_LRU_TUNING=*value* specification. Specify the current AUTO_LRU_TUNING value if you intend to reset only the **lru_min_dirty** or **lru_max_dirty** flushing attributes of the BUFFERPOOL configuration parameter,

Syntax:

```

▶▶ [ADMIN | TASK] (—"ONMODE"—,—"wm"—, —————▶

```



Argument	Effect	Key Considerations
<i>lru_max_dirty</i>	Specifies the percentage of modified pages in the LRU queues at which the queue is cleaned	If a value is specified out of the range of 0 through 100 (inclusive), then the default of 60.00 percent is set. Fractional values are allowed.
<i>lru_min_dirty</i>	Specifies the percentage of modified pages in the LRU queues at which page cleaning is no longer mandatory. Page cleaners might continue cleaning beyond this point under some circumstances.	If a value is specified out of the range of 0 through 100 (inclusive), then the default of 80.00 percent is set. Fractional values are allowed.

Use this command to change the LRU tuning status and also specify flushing behavior. The arguments to this command can dynamically change the setting in memory of the `AUTO_LRU_TUNING` configuration parameter for the current database server session, and can also update the setting for this parameter in the configuration file that will take effect when the database server is restarted. The optional **min** and **max** arguments can reset the **lru_max_dirty** and **lru_min_dirty** values of the `BUFFERPOOL` configuration parameter.

- The "`AUTO_LRU_TUNING =1`" argument enables automatic LRU tuning for the current session.
- The "`AUTO_LRU_TUNING=0`" argument disables automatic LRU tuning for the current session.
- If the `AUTO_LRU_TUNING` argument is followed by both **min** and **max** arguments, the command sets the **lru_max_dirty** and **lru_min_dirty** flushing attributes of the `BUFFERPOOL` parameter.
- If the `AUTO_LRU_TUNING` argument is followed by only the **max** argument, the command sets the **lru_max_dirty** flushing attribute of the `BUFFERPOOL` parameter.
- If the `AUTO_LRU_TUNING` argument is followed by only the **min** argument, the command sets the **lru_min_dirty** flushing attribute of the `BUFFERPOOL` parameter.
- If the `AUTO_LRU_TUNING` argument is the last argument, the command has no effect on the LRU flushing attributes of the `BUFFERPOOL` parameter.

LRU flushing attributes can be set when `AUTO_LRU_TUNING` is either on or off. If a **min** value is greater to or equal to the *max* value for a specific buffer pool, the **min** value is ignored and no error message is returned.

For information about the effects of these configuration parameters and the range or set of values that is valid for each parameter, see *IBM Informix Administrator's Reference*.

onmode Y: Dynamically change SET EXPLAIN for a user session

Syntax:

```

ADMIN TASK (—"ONMODE—" ,—"Y—" ,—"sessionid—" ,—" 2 " )

```

Argument	Effect	Key Considerations
<i>sessionid</i>	Identifies the session for which to change output to the sqexplain.out file.	See Table 2-3 on page 2-627 for a summary of what the sqexplain.out file can display.

This command dynamically changes the output of query plan measurements for an individual session. You can use this command to emulate the SET EXPLAIN statement. The last argument determines whether or not to record query measurements, including the plan of the query optimizer, an estimate of the number of rows returned, and the relative cost of the query.

The following options are valid as the last argument to this command:

Argument	Effect
2	Equivalent to SET EXPLAIN ON AVOID_EXECUTE, this enables output to sqexplain.out.sessionid of subsequent query plans, but without executing the query.
1	Equivalent to SET EXPLAIN ON for <i>sessionid</i> , this enables output to sqexplain.out.sessionid
0	Equivalent to SET EXPLAIN OFF, this disables output of query measurements to sqexplain.out for the current session.

If neither this command nor the SET EXPLAIN statement has been issued for the current session, the default behavior for recording query measurement is equivalent to the effect of this command with the 0 option.

When you specify the 1 or 2 option to this command to enable query measurements, the output is written to the **sqexplain.out.sessionid** file. If an **sqexplain.out** file already exists for the session, the database server writes to that file until an administrator turns off the recording of query measurements for the session.

Specifying the 2 option causes the database server to record query plan measurements, but prevents the query from being executed. When this option is in effect for a session, you can use the SQL Admin API to enable query execution by issuing this command with the same *sessionid* argument, but with 1 or 0 as the last argument.

onmode z: Terminate a user session

Syntax:

```

ADMIN (—"ONMODE"—,—"z"—,—"sessionid"—)
TASK
  
```

Argument	Effect	Key Considerations
<i>sessionid</i>	Terminates the specified user session	This value must be an unsigned integer greater than 0, and must be the session identification number of a currently running session.

onmode Z: Terminate a distributed transaction ungracefully

Syntax:

```

ADMIN (—"ONMODE"—,—"Z"—,—"address"—)
TASK
  
```

Argument	Effect	Key Considerations
<i>address</i>	The shared-memory address associated with a distributed transaction.	This argument must be the address of an ongoing distributed transaction that has exceeded the amount of time that TXTIMEOUT specifies.

This command is not valid until the amount of time that the ONCONFIG parameter TXTIMEOUT specifies has been exceeded.

Cross-server distributed transactions can query data on different database servers. For information on initiating independent actions in a two-phase commit protocol, see the section on multiphase commit protocols in the *IBM Informix Administrator's Guide*.

The *address* argument must conform to the operating-system-specific rules for addressing shared-memory. This address is available from **onstat -x** output.

Warning: If applications are performing distributed transactions, killing one of the distributed transactions can leave your client/server database system in an inconsistent state. This command should rarely be used, and only by an administrator of a database server involved in distributed transactions.

High Availability Initial Commands

As their first argument, these SQL Admin API functions accept the following high availability (HA) *command* specifications to configure a high-availability cluster environment. These HA commands all correspond in their effects to **onmode -d** or **onmode -wf** command-line options.

Among the HA commands that accept additional arguments, most expect the name or alias of a secondary or primary database server, or an Enterprise Replication *group* name, as the required second argument. A few HA commands accept as their third and final argument a password on a Remote Standalone Secondary (RSS) server. For a summary of the calling syntax of the HA commands, see "Calling Syntax of the High Availability Commands" on page 6-46.

The following table shows the high availability initial *command* argument, its effect, and the corresponding **onmode -d** or **onmode -wf** option that produces similar results. See the *IBM Informix Dynamic Server Administrator's Guide* and the *IBM Informix Administrator's Reference* for information about the corresponding parameters and command-line flags of the **onmode** utility. For additional information about high availability environments, see the *IBM Informix Dynamic Server Enterprise Replication Guide*.

Table 6-4. High Availability SQL Admin API Commands

Initial Command	Effect of Successful Operation	Corresponding onmode Command Line
HA MAKE PRIMARY	Change the specified secondary server to a primary server. This command can be used on Remote Standalone Secondary (RSS), Shared Disk Secondary (SDS), and HDR secondary servers.	onmode -d make primary <i>alias</i>
HA MAKE PRIMARY FORCE	Change the specified (inactive) secondary server to a primary server, whether or not a secondary server is connected to it. If sessions are active, the call succeeds, but the shared disk subsystem can become corrupted. This command can be used on Remote Standalone Secondary (RSS), Shared Disk Secondary (SDS), and HDR secondary servers.	onmode -d make primary <i>alias</i> force
HA RSS	Set data-replication type of the specified database server to Remote Standalone Secondary (RSS) server. This command can also be used to demote an HDR secondary server to a Remote Standalone Secondary (RSS) server. In this case, <i>dbserver</i> specifies the name of the primary server.	onmode -d RSS <i>dbserver</i> [<i>password</i>]
HA RSS ADD	Register the specified RSS server name in the sysha database	onmode -d add RSS <i>dbserver</i> [<i>password</i>]
HA RSS CHANGE	Change the password on the specified RSS server	onmode -d change RSS <i>dbserver</i> <i>password</i>
HA RSS DELETE	Remove an RSS server definition from the sysha database	onmode -d delete RSS <i>dbserver</i>
HA SDS CLEAR	The server specified by <i>alias</i> no longer acts as a Shared Disk primary server	onmode -d clear SDS primary <i>alias</i>
HA SDS SET	Define the server specified by <i>alias</i> as a Shared Disk primary server	onmode -d set SDS primary <i>alias</i>
HA SDS SET FORCE	Define the (inactive) server specified by <i>alias</i> as a Shared Disk primary server, whether or not a secondary server is connected to it. If sessions are active, the call succeeds, but the shared disk subsystem can become corrupted.	onmode -d set SDS primary <i>alias</i> force
HA SDS PRIMARY	Change the specified secondary server to a Shared Disk primary server	onmode -d make primary <i>alias</i>

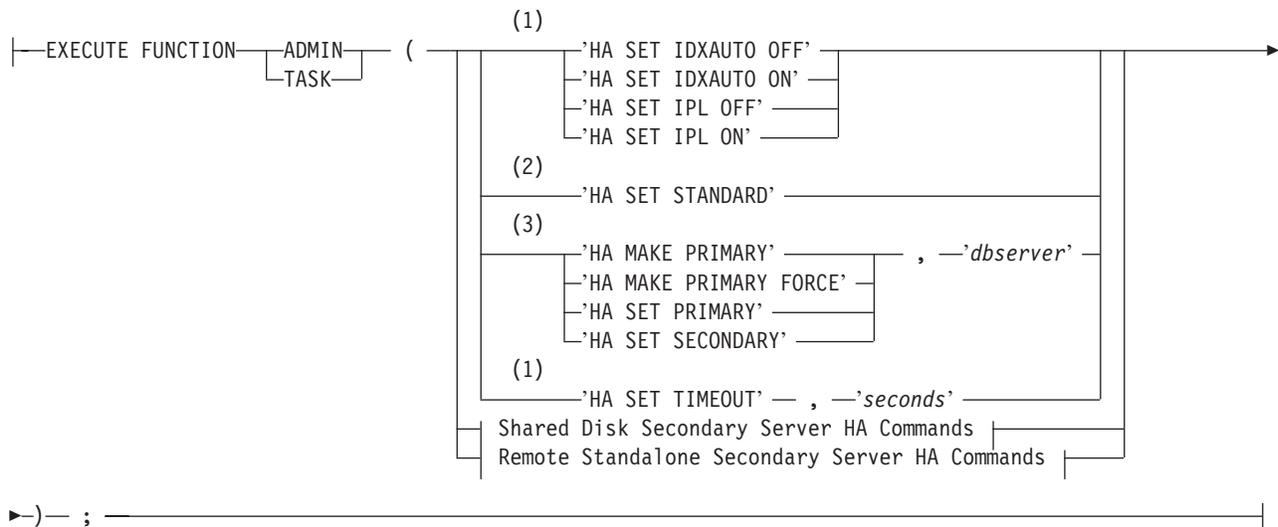
Table 6-4. High Availability SQL Admin API Commands (continued)

Initial Command	Effect of Successful Operation	Corresponding onmode Command Line
HA SDS PRIMARY FORCE	Change the specified (inactive) secondary server to a Shared Disk primary server, whether or not a secondary server is connected to it. If sessions are active, the call succeeds, but the shared disk subsystem can become corrupted.	onmode -d make primary <i>alias</i> force
HA SET IDXAUTO OFF	Disable automatic index replication to the secondary server	onmode -d idxauto off
HA SET IDXAUTO ON	Enable automatic index replication to the secondary server	onmode -d idxauto on
HA SET IPL OFF	Disable logging index builds on the primary server	onmode -wf LOG_INDEX_BUILDS =0
HA SET IPL ON	Enable logging index builds on the primary server	onmode -wf LOG_INDEX_BUILDS =1
HA SET PRIMARY	Set data-replication type of the current database server to primary, and put <i>dbserver</i> into read-only mode.	onmode -d primary <i>dbserver</i>
HA SET SECONDARY	Set data-replication type of the current database server to secondary, and put <i>dbserver</i> into online mode. This command can also be used to promote a Remote Standalone Secondary (RSS) server to an HDR secondary server. In this case, <i>dbserver</i> specifies the name of the HDR primary server.	onmode -d secondary <i>dbserver</i>
HA SET STANDARD	Set data-replication server type of the current database server to standard.	onmode -d standard
HA SET TIMEOUT	Reset SDS_TIMEOUT value	onmode -wf SDS_TIMEOUT= <i>seconds</i>

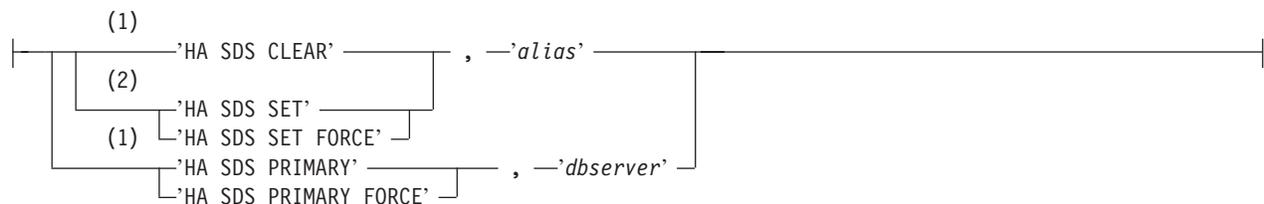
Calling Syntax of the High Availability Commands

The SQL Admin API commands that support High-Availability Data Replication operations have the following calling syntax

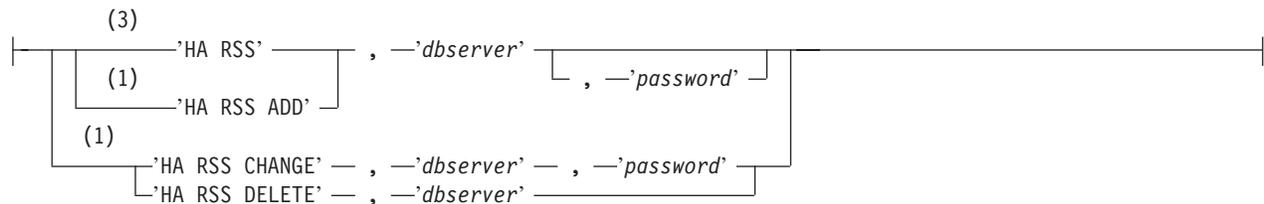
High-Availability Commands:



Shared Disk Secondary Server HA Commands:



Remote Standalone Secondary Server HA Commands:



Notes:

- 1 Run on primary server only.
- 2 Run on primary server or secondary server.
- 3 Run on secondary server only.

Element	Description	Restrictions	Syntax
<i>alias</i>	The alias of a database server that is being enabled or disabled as an SDS primary server	Must be defined in the DBSERVERALIASES configuration setting	“Quoted String” on page 4-189
<i>dbserver</i>	The alias of a database server that is being enabled or disabled as an SDS primary server	Must be defined in the DBSERVERNAME or DBSERVERALIASES configuration setting, or as an Enterprise Replication group name	“Quoted String” on page 4-189

Element	Description	Restrictions	Syntax
<i>password</i>	A password to set or to change on a remote secondary server (RSS)	The <i>password</i> is used only during the first connection attempt. After the primary and secondary server have connected, <i>password</i> cannot be changed.	“Quoted String” on page 4-189
<i>seconds</i>	Number of seconds to which the SDS_TIMEOUT configuration parameter is reset	Must be a positive integer in the range from 2 to 2147483647	“Literal Number ” on page 4-185

For more information about the effects of these operations, see the descriptions of using the **onmode -d** and **onmode -wf** commands for setting secondary server characteristics, replicating indexes, specifying data-replication server types, and resetting LOG_INDEX_BUILDS and SDS_TIMEOUT configuration parameters in the section of *IBM Informix Administrator's Reference* that describes the **onmode** utility.

Enterprise Replication Initial Commands

Use the **cdr** SQL Admin API commands to administer and monitor Enterprise Replication domains.

As their first argument, these SQL Admin API commands accept the Enterprise Replication **cdr** command-line utility specifications to configure an Enterprise Replication environment. The **cdr** commands correspond in their effects to the **cdr** command-line options as described in the appendix for the **cdr** utility commands in the *IBM Informix Dynamic Server Enterprise Replication Guide*.

The SQL Admin API supports most **cdr** commands used to administer and monitor Enterprise Replication. The following commands are not supported:

- **cdr view**
- **cdr list server**
- **cdr list replicate**
- **cdr list replicateset**
- **cdr list template**
- **cdr list repair**
- **cdr stats rqm**
- **cdr stats recv**
- **cdr -V**

For a summary of the calling syntax of the **cdr** commands, see “Calling Syntax of the **cdr** Enterprise Replication Commands.”

Calling Syntax of the **cdr** Enterprise Replication Commands

The SQL Admin API commands that support Enterprise Replication operations have the following calling syntax:

Enterprise Replication Commands

```
EXECUTE FUNCTION ADMIN TASK ( "cdr command_name" (1) "option" );
```

Notes:

- 1 Repeat a maximum of six times.

Argument	Effect	Considerations
<i>command_name</i>	Specifies the name of a cdr command	Cannot include any hyphens, flags, nor other constraining options to <i>command_name</i> that the cdr command-line utility requires. Abbreviations are not supported.
<i>option</i>	Specifies one or more elements of the cdr command-line options to <i>command_name</i>	Must be delimited by quotation marks, and include (in their correct order) any hyphens, flags, or other elements of cdr command-line options that <i>command_name</i> requires. Abbreviations can be used.

Use the **cdr** commands to produce the same effect as with the **cdr** command-line utility to manage Enterprise Replication.

The first argument must include only the **cdr** command names exactly as specified in the appendix for the **cdr** utility in the *IBM Informix Dynamic Server Enterprise Replication Guide*, such as **cdr define server**. Command names are case-sensitive and abbreviations (such as `cdr sto replset` instead of `cdr stop replicaset`) are **not** supported. The SQL Admin API does not perform any validation before passing the parameters to the **cdr** utility.

The second and any following arguments include the command options. The options can be specified in one or more arguments.

The following example illustrates the use of the SQL Admin API to define an Enterprise Replication server:

```
EXECUTE FUNCTION task ( 'cdr define server', '--connect=g_amsterdam
--ats=/local0/er/ats --ris=/local0/er/ris --init g_amsterdam' );
```

The following example shows how the parameters can be spread over several arguments; the above statement can also be written as:

```
EXECUTE FUNCTION task ( 'cdr define server',
'--connect=g_amsterdam',
'--ats=/local0/er/ats',
'--ris=/local0/er/ris',
'--init g_amsterdam' );
```

Arguments with quoted strings must be escaped with the backslash (\). The following example shows an escaped quoted string:

```
EXECUTE FUNCTION task ( 'cdr change replicate',
'-d repl_1 -\'db1@server1:antonio.table1\' \'db2@server2:carlo.table2\'');
```

Restricting the Size of the `command_history` Table

SQL administration API functions that are invoked by the Scheduler, by SQL scripts, or directly by user **informix** can cause the number of rows in the

command_history table to increase over time. A schedule task can purge old data from this table, but you should be aware that any rows that are no longer needed in the **command_history** table should be deleted or truncated periodically.

That table resides in the **sysadmin** database, which Dynamic Server creates, by default, in the **root** dbspace. To avoid the risk of default sensors and the **command_history** table occupying all of the available storage in the **root** dbspace, the **ADMIN** or **TASK** function can move the **sysadmin** database to some other dbspace that the **RESET SYSADMIN** command specifies. If you specify no dbspace, this command drops the **sysadmin** database, and then recreates it in the same dbspace.

Related Information

The administrative utility commands of Dynamic Server that correspond to SQL administration API function calls are identified in sections that describe the syntax of those utilities in the *IBM Informix Administrator's Reference*, which also describes the schema of the **sysadmin** database.

For additional information about using these SQL Admin API functions to support automatic monitoring of the database server and for taking corrective actions, see the *IBM Informix Dynamic Server Administrator's Guide*.

Session Configuration Procedures

These built-in SPL procedures enable the Database Administrator to execute SQL and SPL statements automatically when a user connects to or disconnects from the database.

These routines are called "built-in" procedures in this document because the database server recognizes their names and treats them differently from how it treats other routines, but the database server does not create these routines automatically. To use their features, the DBA must issue the **CREATE PROCEDURE** statement or the **CREATE PROCEDURE FROM** statement to define the actions of these routines and to register them in the database. Only the DBA or user **informix** can create, alter, or drop these routines.

If the DBA specifies the login ID of a user as the *owner* of one of these procedures, the database server executes it when the specified user connects to or disconnects from the database. If the DBA specifies **PUBLIC** as the owner, that routine is automatically executed when a user who is not the owner of any of these built-in session configuration procedures connects to or disconnects from the database. Different databases of the same database server instance can specify either the same or different session configuration procedures for individual users or for **PUBLIC**. These built-in procedures are useful in setting the session environment or activating a role for users of applications whose code cannot easily be modified.

These are the built-in session configuration procedures:

- **sysdbclose**
- **sysdbopen**

Using SYSDBOPEN and SYSDBCLOSE Procedures

To set the initial environment for one or more sessions, create and install the **sysdbopen**(SPL procedure. The typical effect of this procedure is to initialize properties of a session without requiring the properties to be explicitly defined within the session. Such procedures are useful if users access databases through client applications that cannot modify application code or set environment options or environment variables.

The **sysdbopen** procedure is executed whenever users successfully issue the DATABASE or CONNECT statement to explicitly connect to a database where the procedures are installed. (But when a user who is connected to the local database calls a remote UDR or performs a distributed DML operation that references a remote database object by using the *database:object* or *database@server:object* notation, no **sysdbopen** procedure is invoked in the remote database.)

These procedures are exceptions to the general rule that Dynamic Server ignores the name of the owner of a UDR when a routine is invoked in a database that is not ANSI-compliant. For UDRs other than **sysdbopen** and **sysdbclose**, multiple versions of UDRs that have the same SQL identifier but that have different *owner* names cannot be registered in the same database unless the CREATE DATABASE statement that created the database also included the WITH LOG MODE ANSI keywords.

You can also create the **sysdbclose** SPL procedure, which is executed when a user issues the CLOSE DATABASE or DISCONNECT statement to disconnect from the database. If a PUBLIC.**sysdbclose** procedure is registered in the database, and no *user.sysdbclose* procedure is registered for the current user, then the PUBLIC.**sysdbclose** procedure is executed automatically when that *user* disconnects from the database.

You can include valid SQL or SPL language statements that are appropriate when a database is opened or closed. The general restrictions on SQL statements that are valid in SPL procedures also apply to these routines. See the following sections for restrictions on SQL and SPL statements within SPL routines:

- “Subset of SPL Statements Valid in the Statement Block” on page 5-75.
- “SQL Statements Valid in SPL Statement Blocks” on page 5-76.
- “Restrictions on SPL Routines in Data-Manipulation Statements” on page 5-78.

Important: The **sysdbopen** and **sysdbclose** procedures are exceptions to the scope rule for stored procedures. In ordinary UDR procedures, the scope of variables and statements is local. SET PDQPRIORITY and SET ENVIRONMENT statement settings do not persist when these SPL procedures exit. In **sysdbopen** and **sysdbclose** procedures, however, statements that set the session environment remain in effect until another statement resets the options, or the session ends.

For example, the following procedure sets the transaction isolation level to Repeatable Read, and sets the **OPTCOMPIND** environment variable to instruct the query optimizer to prefer nested-loop joins. When a user who owns no *user.sysdbopen* procedure connects to the database, this routine will be executed:

```
CREATE PROCEDURE public.sysdbopen()  
    SET ISOLATION TO REPEATABLE READ;  
    SET ENVIRONMENT OPTCOMPIND '1';  
END PROCEDURE;
```

Procedures do not accept arguments or return values. The **sysdbopen** and **sysdbclose** procedures must be registered in each database in which you want to execute them. (In Extended Parallel Server, they must be executed from the connection coserver.) The DBA can create the following four categories of **sysdbopen** and **sysdbclose** procedures.

Procedure Name	Description
----------------	-------------

<i>user.sysdbopen</i>	
-----------------------	--

	This procedure is executed when the specified <i>user</i> opens the database as the current database.
--	---

public.sysdbopen	
-------------------------	--

	If no <i>user.sysdbopen</i> procedure applies, this procedure is executed when any user opens the database as the current database. To avoid duplicating SPL code, you can call this procedure from a user-specific procedure.
--	--

<i>user.sysdbclose</i>	
------------------------	--

	This procedure is executed when the specified <i>user</i> closes the database, disconnects from the database server, or the user session ends. If <i>user.sysdbclose</i> did not exist when the session opened the database, however, the procedure is not executed when the session closes the database.
--	---

public.sysdbclose	
--------------------------	--

	If no <i>user.sysdbclose</i> procedure applies, this procedure is executed when the user closes or disconnects from the database server, or when the session ends. If public.sysdbopen did not exist when the session opened the database, however, the procedure is not executed when the session closes the database.
--	--

The database server calls *user.sysdbclose* procedure, if it exists in the database, or **public.sysdbclose** if this exists and no version owned by *user* exists, when the CLOSE DATABASE or DISCONNECT statement explicitly terminates the connection. If the application terminates without issuing the CLOSE DATABASE or DISCONNECT statement, the database server forces an implicit close of the database and executes the **sysdbclose** procedure, if a UDR with that name is owned by the user or by PUBLIC.

Make sure that you set file access permissions appropriately to allow intended users to execute the SPL procedure statements. For example, if the SPL procedure executes a command that writes output to a local directory, permissions must be set to allow users to write to this directory. If you want the procedure to continue if permission failures occur, include an ON EXCEPTION error handler for this condition.

For more information about the SQL statements that can appear in SPL routines, and about SPL support for transactions and for roles, see the section “Statement Block ” on page 5-74.

Warning: If a **sysdbclose** procedure fails, the failure is ignored. If a **sysdbopen** procedure fails, however, the database cannot be opened.

To avoid situations in which a database cannot be opened, take the following precaution while you are writing and debugging a **sysdbopen** procedure:

- For Dynamic Server databases, set the **IFX_NODBPROC** environment variable before you connect to the database. When **IFX_NODBPROC** is set, the procedure is not executed, and failures cannot prevent the database from opening.
- For Extended Parallel Server databases, set the **DEBUG** environment variable to 'NODBPROC' before you connect to the database. When 'NODBPROC' is the **DEBUG** setting, the procedure is not executed, and failures cannot prevent the database from opening.

Failures from these procedures can be generated by the system or simulated within the procedures by the **RAISE EXCEPTION** statement of SPL. If the **sysdbopen** routine that is invoked for a user at connection time includes this statement, that user cannot connect to the database. For more information, refer to the description of "RAISE EXCEPTION " on page 3-41.

For security reasons, non-DBAs cannot prevent execution of these procedures. For some applications, however, such as *ad hoc* query applications, users can execute commands and SQL statements that subsequently change the environment.

A default role defined in the **sysdbopen** procedure take precedence over any other role that the user holds when a user establishes a connection to a database in which **sysdbopen** successfully specifies a default role for that user.

Any database objects that are created by DDL statements in a *user.sysdbopen* or *user.sysdbclose* procedure are owned by the connected *user*, and any object created within **PUBLIC.sysdbopen** or within **PUBLIC.sysdbclose** is owned by the **PUBLIC** userid, unless the object name is fully qualified by some other owner name when the object name is declared in the DDL statement.

For ANSI-compliant databases, an explicit **COMMIT WORK** statement is required at the end of the **sysdbopen** or **sysdbclose** definition in the **CREATE PROCEDURE** statement, to prevent any implicit transactions of SQL statements that the **sysdbopen** or **sysdbclose** procedure executes from being rolled back when the procedure terminates. (Omitting the **COMMIT WORK** statement does not cause the connection to fail, but does waste resources in opening and then rolling back the transactions.)

For a list of SQL statements that are not valid in these procedures, see "SQL Statements Valid in SPL Statement Blocks" on page 5-76. For a list of the SPL statements that are valid in these procedures, see "Subset of SPL Statements Valid in the Statement Block" on page 5-75.

For general information about how to write and install SPL procedures, refer to the section about SPL routines in *IBM Informix Guide to SQL: Tutorial*.

The EXPLAIN_SQL Routine

IBM Data Studio consists of a set of tools to use for administration, data modeling, and building queries from data that comes from data servers. The **EXPLAIN_SQL** routine prepares a query and returns a query plan in XML. The IBM Data Studio Administration Edition can use the **EXPLAIN_SQL** routine to obtain a query plan in XML format, interpret the XML, and render the plan visually.

If you plan to use IBM Data Studio to obtain Visual Explain output, you must create and specify a default sbspace name for the SBS spacENAME configuration parameter in your ONCONFIG file. The EXPLAIN_SQL routine creates BLOBs in this sbspace.

For information on using IBM Data Studio, see IBM Data Studio documentation.

UDR Definition Routines

The UDR definition routines are built-in routines that enable users to perform various tasks for developing or modifying external user-defined routines of Dynamic Server, or for enabling Java Common Client (JCC) procedures to access IDS and DB2® databases through the Distributed Relational Database Architecture™ (DRDA) protocol.

If the IFX_EXTEND_ROLE configuration parameter is set to '0n' or 1, authorization to use the built-in routines that manipulate shared objects is available only to the Database Server Administrator, and to users to whom the DBSA has granted the EXTEND role.

These are the built-in UDR definition routines:

- `ifx_replace_module()`
- `ifx_unload_module()`
- `jvpcontrol()`
- `sqlj.alter_java_path()`
- `sqlj.install_jar()`
- `sqlj.remove_jar()`
- `sqlj.replace_jar()`
- `sqlj.alter_java_path()`
- `sqlj.setUDTextName()`
- `sqlj.unsetUDTextName()`
- `sysibm.Metadata()`
- `sysibm.SQLCAMessage()`

IFX_REPLACE_MODULE Function

The IFX_REPLACE_MODULE function replaces a loaded shared-object file of a UDR written in the C language with a new version that has a different name or location. If the IFX_EXTEND_ROLE configuration parameter is set to '0n' or 1, authorization to use this function is available only to the Database Server Administrator (DBSA), and to users whom the DBSA has granted the EXTEND role.

IFX_REPLACE_MODULE Function:

```
|—IFX_REPLACE_MODULE—(—old_module—,—new_module—,—"C" —)———|
```

Argument	Description	Restrictions	Syntax
<i>new_module</i>	Full pathname of the new shared-object file to replace the shared-object file that <i>old_module</i> specifies	The shared-object file must exist with the specified pathname, which can be no more than 255 bytes long	“Quoted String” on page 4-189

Argument	Description	Restrictions	Syntax
<i>old_module</i>	Full pathname of the shared-object file to replace with the shared-object file that <i>new_module</i> specifies	The shared-object file must exist with the specified pathname, which can be no more than 255 bytes long	“Quoted String” on page 4-189

The **IFX_REPLACE_MODULE** function is a DBA-privileged function that returns an integer value to indicate the status of the shared-object-file replacement operation:

- Zero (0) to indicate success
- A negative integer to indicate an error.

Do not use the **IFX_REPLACE_MODULE** function to reload a module of the same name. If the full names of the old and new modules that you send to **IFX_REPLACE_MODULE** are the same, then unpredictable results can occur.

After **IFX_REPLACE_MODULE** completes execution, the database server ages out the *old_module* shared-object file; that is, all statements subsequent to the **IFX_REPLACE_MODULE** function will use UDRs in the *new_module* shared-object file, and the old module will be unloaded when any statements that were using it are complete. Thus, for a brief time, both the *old_module* and the *new_module* shared-object files could be resident in memory. If this aging out behavior is undesirable, use the **IFX_UNLOAD_MODULE** function to unload the shared-object file completely.

On UNIX, for example, suppose you want to replace the **circle.so** shared library, which contains UDRs written in the C language. If the old version of this library resides in the **/usr/apps/opaque_types** directory and the new version in the **/usr/apps/shared_libs** directory, then the following EXECUTE FUNCTION statement executes the **IFX_REPLACE_MODULE** function:

```
EXECUTE FUNCTION ifx_replace_module(
    "/usr/apps/opaque_types/circle.so",
    "/usr/apps/shared_libs/circle.so", "C");
```

On Windows, for another example, suppose you want to replace the **circle.dll** dynamic link library, which contains C UDRs. If the old version of this library resides in the **C:\usr\apps\opaque_types** directory and the new version in the **C:\usr\apps\DLLs** directory, then the following EXECUTE FUNCTION statement executes the **IFX_REPLACE_MODULE** function:

```
EXECUTE FUNCTION ifx_replace_module(
    "C:\usr\apps\opaque_types\circle.dll",
    "C:\usr\apps\DLLs\circle.dll", "C");
```

To execute the **IFX_REPLACE_MODULE** function in an IBM Informix ESQL/C application, you must associate the function with a cursor.

For more information on how to use **IFX_REPLACE_MODULE** to replace a shared-object file, see the chapter on how to design a UDR in *IBM Informix User-Defined Routines and Data Types Developer's Guide*. For information on how to use the **IFX_UNLOAD_MODULE** function, see the section “**IFX_UNLOAD_MODULE** Function.”

IFX_UNLOAD_MODULE Function

The **IFX_UNLOAD_MODULE** function unloads the shared-object file of a UDR written in the C language from shared memory.

IFX_UNLOAD_MODULE Function:

|—IFX_UNLOAD_MODULE—(—*module_name*—,—"C"—)|

Argument	Description	Restrictions	Syntax
<i>module_name</i>	Full pathname of file to unload	Shared-object file must exist and be unused. Pathname can be up to 255 bytes long.	"Quoted String" on page 4-189

The **IFX_UNLOAD_MODULE** function is an owner-privileged function whose owner is user **informix**. It returns an integer value to indicate the status of the shared-object-file unload operation:

- Zero (0) to indicate success
- A negative integer to indicate an error.

The **IFX_UNLOAD_MODULE** function can only unload an unused shared-object file; that is, when no executing SQL statements (in any database) are using any UDRs in the specified shared-object file. If any UDR in the shared-object file is currently in use, then **IFX_UNLOAD_MODULE** raises an error.

On UNIX, for example, suppose you want to unload the **circle.so** shared library, which contains C UDRs. If this library resides in the **/usr/apps/opaque_types** directory, you can use the following EXECUTE FUNCTION statement to execute the **IFX_UNLOAD_MODULE** function:

```
EXECUTE FUNCTION ifx_unload_module  
  ("/usr/apps/opaque_types/circle.so", "C");
```

On Windows, for example, suppose you want to unload the **circle.dll** dynamic link library, which contains C UDRs. If this library is in the **C:\usr\apps\opaque_types** directory, you can use the following EXECUTE FUNCTION statement to execute the **IFX_UNLOAD_MODULE** function:

```
EXECUTE FUNCTION ifx_unload_module  
  ("C:\usr\apps\opaque_types\circle.dll", "C");
```

For more information about using the built-in **IFX_REPLACE_MODULE()** and **IFX_UNLOAD_MODULE()** UDR definition routines, see the *IBM Informix User-Defined Routines and Data Types Developer's Guide* and the *IBM Informix Datablade API Programmer's Guide*.

jvpcontrol Function

The **jvpcontrol()** function is a built-in iterative function that you can use to obtain information about a Java Virtual Processor (JVP) class.

The jvpcontrol Function:

|—informix.jvpcontrol—(—"MEMORY"—*jvp_id*—"")
 └──THREADS──┘

Argument	Description	Restrictions	Syntax
<i>jvp_id</i>	Name of the Java Virtual Processor (JVP) class about which you seek information	The specified Java Virtual Processor class must exist	"Identifier " on page 5-21

You must associate this function with the equivalent of a cursor in the Java language.

Using the MEMORY Keyword

When you specify the MEMORY keyword, the **jvpcontrol** function returns the memory usage on the JVP class that you specify. The following example requests information about the memory usage of the JVP class named 4:

```
EXECUTE FUNCTION INFORMIX.JVPCONTROL ("MEMORY 4");
```

Using the THREADS Keyword

When you specify the THREADS keyword, the **jvpcontrol** function returns a list of the threads running on the JVP class that you specify. The following example requests information about the threads running on the JVP class named 4:

```
EXECUTE FUNCTION INFORMIX.JVPCONTROL ("THREADS 4");
```

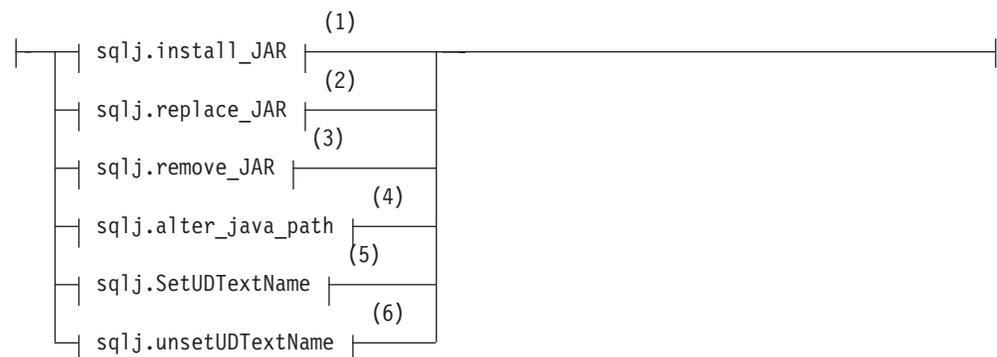
For more information about using **jvpcontrol()** and the built-in **sqlj** routines, see the *J/Foundation Developer's Guide/J/Foundation Developer's Guide*.

SQLJ Driver Built-In Procedures

Use the SQLJ Driver built-in procedures for one of the following tasks:

- To install, replace, or remove a set of Java classes
- To specify a path for Java class resolution for Java classes that are included in a JAR file
- To map or remove the mapping between a user-defined type and the Java type to which it corresponds

SQLJ Driver Built-In Procedures:



Notes:

- 1 See "sqlj.install_jar" on page 6-58
- 2 See "sqlj.replace_jar" on page 6-59
- 3 See "sqlj.remove_jar" on page 6-59
- 4 See "sqlj.alter_java_path" on page 6-60
- 5 See "sqlj.setUDTextName" on page 6-61
- 6 See "sqlj.unsetUDTextName" on page 6-62

A client application must specify the 'sqlj' owner name to invoke these functions from an ANSI-compliant database.

The SQLJ built-in procedures are stored in the **sysprocedures** system catalog table. They are grouped under the **sqlj** schema.

Tip: For any Java static method, the first built-in procedure that you execute must be the **sqlj.install_jar()** procedure. You must install the JAR file before you can create a UDR or map a user-defined data type to a Java type. Similarly, you cannot use any of the other SQLJ built-in procedures until you have used **sqlj.install_jar()**.

sqlj.install_jar

Use the **sqlj.install_jar()** procedure to install a JAR file in the current database and assign to it a JAR identifier.

sqlj.install_jar:

```
sqlj.install_jar(—jar_file—, — Jar Name —, —(1) deploy —)
```

Notes:

- 1 See "Jar Name" on page 5-34

Argument	Description	Restrictions	Syntax
<i>deploy</i>	Integer that causes the procedure to search for deployment descriptor files in the JAR file	None	"Literal Number " on page 4-185
<i>jar_file</i>	URL of the JAR file that contains the Java language UDR	Maximum length of the URL is 255 bytes	"Quoted String" on page 4-189

For example, consider a Java class **Chemistry** that contains the following static method **explosiveReaction()**:

```
public static int explosiveReaction(int ingredient)
```

Here the **Chemistry** class resides in this JAR file on the server computer:

```
/students/data/Courses.jar
```

You can install all classes in the **Courses.jar** file in the current database with the following call to the **sqlj.install_jar()** procedure:

```
EXECUTE PROCEDURE
    sqlj.install_jar("file://students/data/Courses.jar", "course_jar");
```

The **sqlj.install_jar()** procedure assigns the JAR ID, **course_jar**, to the **Courses.jar** file that it has installed in the current database.

After you define a JAR ID in the database, you can use that JAR ID when you create and execute a UDR written in the Java language.

When you specify a nonzero number for the third argument, the database server searches through any included deployment descriptor files. For example, you might want to include descriptor files that include SQL statements to register and grant privileges on UDRs in the JAR file.

File Permissions on Jar Files

After `sqlj.install_jar()` installs a JAR file in the database and declares a JAR ID for the file, Dynamic Server can access that JAR file only if the user who installed the Dynamic Server instance (typically, user 'informix') has permission to read the directory where the JAR file resides. On UNIX systems, for example, this implies that an attempt to read a JAR file that has 600 permissions fails with a `FILENOTFOUND` exception. The same operation can succeed, however, after the `chmod` utility sets the permissions to 660 (`rw-rw----`).

sqlj.replace_jar

Use the `sqlj.replace_jar()` procedure to replace a previously installed JAR file with a new version. When you use this syntax, you provide only the new JAR file and assign to it the JAR ID for which you want to replace the file.

sqlj.replace_jar:

```
|—sqlj.replace_jar—(—jar_file—, — Jar Name —————)—————|  
                                     (1)
```

Notes:

- 1 See "Jar Name" on page 5-34

Argument	Description	Restrictions	Syntax
<i>jar_file</i>	URL of the JAR file that contains the UDR written in Java	The maximum length of the URL is 255 bytes	"Quoted String" on page 4-189

If you attempt to replace a JAR file that is referenced by one or more UDRs, the database server generates an error. You must drop the referencing UDRs before replacing the JAR file.

For example, the following call replaces the **Courses.jar** file, which had previously been installed for the **course_jar** identifier, with the **Subjects.jar** file:

```
EXECUTE PROCEDURE  
  sqlj.replace_jar("file://students/data/Subjects.jar",  
  "course_jar");
```

Before you replace the **Course.jar** file, you must drop the user-defined function `sql_explosive_reaction()` with the `DROP FUNCTION` (or `DROP ROUTINE`) statement.

sqlj.remove_jar

Use the `sqlj.remove_jar()` procedure to remove a previously installed JAR file from the current database.

sqlj.remove_jar:

```
|—sqlj.remove_jar—(— Jar Name —————, — 0 —————)—————|  
                                     (1) deploy
```

Notes:

- 1 See "Jar Name" on page 5-34

Argument	Description	Restrictions	Syntax
<i>deploy</i>	Integer that causes the procedure to search for deployment descriptor files in the JAR file	None	"Literal Number " on page 4-185

If you attempt to remove a JAR file that is referenced by one or more UDRs, the database server generates an error. You must drop the referencing UDRs before you replace the JAR file. For example, the following SQL statements remove the JAR file associated with the **course_jar** JAR ID:

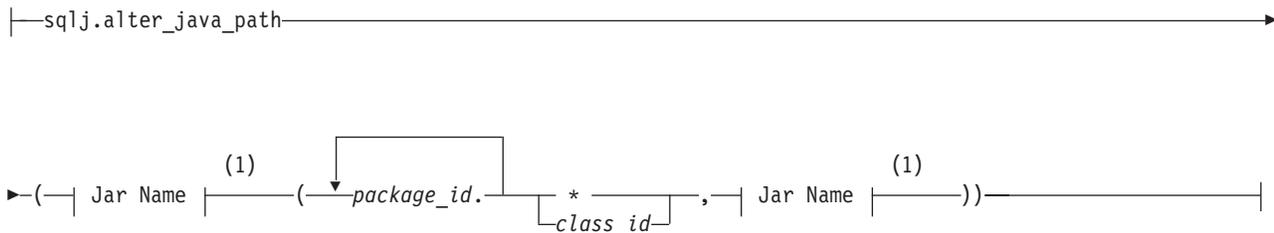
```
DROP FUNCTION sql_explosive_reaction;
EXECUTE PROCEDURE sqlj.remove_jar("course_jar");
```

When you specify a nonzero number for the second argument, the database server searches through any included deployment descriptor files. For example, you might want to include descriptor files that include SQL statements to revoke privileges on UDRs in the associated JAR file and drop them from the database.

sqlj.alter_java_path

Use **sqlj.alter_java_path()** to specify the *jar-file path* to use when the routine manager resolves related Java classes for the JAR file of a UDR written in the Java language.

sqlj.alter_java_path:



Notes:

- 1 See "Jar Name" on page 5-34

Argument	Description	Restrictions	Syntax
<i>class_id</i>	Java class that contains method to implement the UDR	Java class must exist in the JAR file that <i>jar_id</i> specifies. Identifier must not exceed 255 bytes.	Language specific
<i>package_id</i>	Name of the package that contains the Java class	The fully qualified identifier of <i>package_id.class_id</i> must not exceed 255 bytes	Language specific

The JAR IDs that you specify, namely the JAR ID for which you are altering the JAR-file path and the resolution JAR ID, must both have been installed with the **sqlj.install_jar** procedure. When you invoke a UDR written in the Java language, the routine manager attempts to load the Java class in which the UDR resides. At this time, it must resolve the references that this Java class makes to other Java classes.

The three types of such class references are these:

1. References to Java classes that the JVPCLASSPATH configuration parameter specifies (such as Java system classes like **java.util.Vector**)

2. References to classes that are in the same JAR file as the UDR
3. References to classes that are outside the JAR file that contains the UDR.

The routine manager implicitly resolves classes of type 1 and 2 in the preceding list. To resolve type 3 references, it examines all the JAR files in the JAR file path that the latest call to `sqlj.alter_java_path()` specified.

The routine manager issues an exception if it cannot resolve a class reference. The routine manager checks the JAR file path for class references *after* it performs the implicit type 1 and type 2 resolutions.

If you want a Java class to be loaded from the JAR file that the JAR file path specifies, make sure the Java class is *not* present in the JVPCLASSPATH configuration parameter. Otherwise, the system loader picks up that Java class first, which might result in a different class being loaded than what you expect.

Suppose that the `sqlj.install_jar()` procedure and CREATE FUNCTION have been executed as the preceding sections describe. The following SQL statement invokes `sql_explosive_reaction()` function in the `course_jar` JAR file:

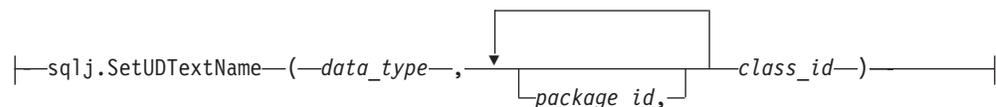
```
EXECUTE PROCEDURE alter_java_path("course_jar",
    "(professor/*, prof_jar)");
EXECUTE FUNCTION sql_explosive_reaction(10000);
```

The routine manager attempts to load the class `Chemistry`. It uses the path that the call to `sqlj.alter_java_path()` specifies to resolve any class references. Therefore, it checks the classes that are in the `professor` package of the JAR file that `prof_jar` identifies.

sqlj.setUDTextName

Use the `sqlj.setUDTextName()` procedure to define the mapping between a user-defined data type and a Java class.

sqlj.SetUDTextName:



Argument	Description	Restrictions	Syntax
<code>class_id</code>	Java class that contains the Java data type	Qualified name <code>package_id.class_id</code> must not exceed 255 bytes	Language-specific rules for Java identifiers
<code>data_type</code>	User-defined type for which to create a mapping	Name must not exceed 255 bytes	"Identifier " on page 5-21
<code>package_id</code>	Name of package that contains the <code>class_id</code> Java class	Same length restrictions as <code>class_id</code>	Language-specific rules for Java identifiers

You must have registered the user-defined data type in the CREATE DISTINCT TYPE, CREATE OPAQUE TYPE, or CREATE ROW TYPE statement.

To look up the Java class for a user-defined data type, the database server searches in the JAR-file path, which the `sqlj.alter_java_path()` procedure has specified. For more information on the JAR-file path, see "sqlj.alter_java_path" on page 6-60.

The SQLJ Driver looks in the path that **CLASSPATH** specifies in the client environment before it asks the database server for the name of the Java class.

The **setUDTextName** routine is an extension to the SQLJ:SQL *Routines Using the Java Programming Language* specification.

sqlj.unsetUDTextName

Use the **sqlj.unsetUDTextName()** routine to remove the mapping from a user-defined data type to a Java class.

sqlj.unsetUDTextName:

|—sqlj.unsetUDTextName—(—data_type—)|

Argument	Description	Restrictions	Syntax
<i>data_type</i>	User-defined data type for which to remove the mapping	Must exist	"Identifier " on page 5-21

This routine removes the SQL-to-Java mapping and consequently removes any cached copy of the Java class from shared memory of the database server.

The **unsetUDTextName** routine is an extension to the SQLJ:SQL *Routines Using the Java Programming Language* specification.

DRDA Support Functions

Dynamic Server provides built-in functions that support the Distributed Relational Database Architecture (DRDA) protocol when the Dynamic Server instance is configured as a DRDA application server. (This is accomplished by setting **DBSERVERALIASES** to **DRDA** in the configuration file.)

- The **sysibm.Metadata** function provides database metadata information to Java Common Client (JCC) client applications.
- The **sysibm.SCLCAMessage** function supports DRDA error handling.

The IDS implementation of these functions conforms to the DR Level 5 SQLAM 7 standard

Metadata Function

The **sysibm.Metadata** function is an SPL routine that can be called by Java Common Client (JCC) applications to dynamically retrieve database metadata. The **Metadata** routine is automatically created in every database of Dynamic Server instances that are configured as DRDA application servers. A JCC client application must specify the '**sysibm**' owner name to invoke this function from an ANSI-compliant database.

It has the following routine definition:

```
create procedure SYSIBM.METADATA() returning
  integer as allProceduresAreCallable,
  integer as allTablesAreSelectable,
  integer as nullsAreSortedHigh,
  integer as nullsAreSortedLow,
  integer as nullsAreSortedAtStart,
```

```

integer as nullsAreSortedAtEnd,
integer as usesLocalFiles,
integer as usesLocalFilePerTable,
integer as storesUpperCaseIdentifiers,
integer as storesLowerCaseIdentifiers,
integer as storesMixedCaseIdentifiers,
integer as storesLowerCaseQuotedIdentifiers,
integer as storesMixedCaseQuotedIdentifiers,
nvarchar(4096) as getSQLKeywords,
varchar(100) as getNumericFunctions,
varchar(100) as getStringFunctions,
varchar(100) as getSystemFunctions,
varchar(100) as getTimeDateFunctions,
varchar(25) as getSearchStringEscape,
varchar(25) as getExtraNameCharacters,
integer as supportsAlterTableWithAddColumn,
integer as supportsAlterTableWithDropColumn,
integer as supportsConvert,
varchar(255) as supportsConvertType,
integer as supportsDifferentTableCorrelationNames,
integer as supportsExpressionsInOrderBy,
integer as supportsOrderByUnrelated,
integer as supportsGroupBy,
integer as supportsGroupByUnrelated,
integer as supportsGroupByBeyondSelect,
integer as supportsMultipleResultSets,
integer as supportsMultipleTransactions,
integer as supportsCoreSQLGrammar,
integer as supportsExtendedSQLGrammar,
integer as supportsANSI92IntermediateSQL,
integer as supportsANSI92FullSQL,
integer as supportsIntegrityEnhancementFacility,
integer as supportsOuterJoins,
integer as supportsFullOuterJoins,
integer as supportsLimitedOuterJoins,
varchar(50) as getSchemaTerm,
varchar(50) as getProcedureTerm,
varchar(50) as getCatalogTerm,
integer as isCatalogAtStart,
varchar(50) as getCatalogSeparator,
integer as supportsSchemasInDataManipulation,
integer as supportsSchemasInProcedureCalls,
integer as supportsSchemasInTableDefinitions,
integer as supportsSchemasInIndexDefinitions,
integer as supportsSchemasInPrivilegeDefinitions,
integer as supportsCatalogsInDataManipulation,
integer as supportsCatalogsInProcedureCalls,
integer as supportsCatalogsInTableDefinitions,
integer as supportsCatalogsInIndexDefinitions,
integer as supportsCatalogsInPrivilegeDefinitions,
integer as supportsPositionedDelete,
integer as supportsPositionedUpdate,
integer as supportsSelectForUpdate,
integer as supportsStoredProcedures,
integer as supportsSubqueriesInComparisons,
integer as supportsUnion,
integer as supportsUnionAll,
integer as supportsOpenCursorsAcrossCommit,
integer as supportsOpenCursorsAcrossRollback,
integer as supportsOpenStatementsAcrossCommit,
integer as supportsOpenStatementsAcrossRollback,
integer as getMaxBinaryLiteralLength,
integer as getMaxCharLiteralLength,
integer as getMaxColumnNameLength,
integer as getMaxColumnsInGroupBy,
integer as getMaxColumnsInIndex,
integer as getMaxColumnsInOrderBy,

```

```

integer as getMaxColumnsInSelect,
integer as getMaxColumnsInTable,
integer as getMaxConnections,
integer as getMaxCursorNameLength,
integer as getMaxIndexLength,
integer as getMaxSchemaNameLength,
integer as getMaxProcedureNameLength,
integer as getMaxCatalogNameLength,
integer as getMaxRowSize,
integer as doesMaxRowSizeIncludeBlobs,
integer as getMaxStatementLength,
integer as getMaxStatements,
integer as getMaxTableNameLength,
integer as getMaxTablesInSelect,
integer as getMaxUserNameLength,
integer as getDefaultTransactionIsolation,
integer as supportsTransactions,
varchar(50) as supportsTransactionIsolationLevel,
integer as supportsDataDefinitionAndDataManipulationTransactions,
integer as supportsDataManipulationTransactionsOnly,
integer as dataDefinitionCausesTransactionCommit,
integer as dataDefinitionIgnoredInTransactions,
varchar(100) as supportsResultSetType,
varchar(100) as supportsResultSetConcurrency,
varchar(100) as ownUpdatesAreVisible,
varchar(100) as ownDeletesAreVisible,
varchar(100) as ownInsertsAreVisible,
varchar(100) as othersUpdatesAreVisible,
varchar(100) as othersDeletesAreVisible,
varchar(100) as othersInsertsAreVisible,
varchar(100) as updatesAreDetected,
varchar(100) as deletesAreDetected,
varchar(100) as insertsAreDetected,
integer as supportsBatchUpdates,
integer as supportsSavepoints,
integer as supportsGetGeneratedKeys

```

sysibm.SQLCAMessage Function

By default, Java Common Client (JCC) for IDS does not return localized error messages. Detailed and localized error messages from the server are expected, however, when property “retrieveMessagesFromServerOnGetMessage=true” is set in the connection URL.

The **SQLCAMessage** function is an SPL routine that supports the retrieval of detailed error message text from remote DB2 or IDS database servers to JCC applications that use the Distributed Relational Database Architecture (DRDA) protocol. The **SQLCAMessage** routine is automatically created in every database of Dynamic Server instances that are configured as DRDA application servers. The JCC client application must specify the ‘**sysibm**’ owner name to invoke this function from an ANSI-compliant database.

The **SQLCAMessage** function retrieves localized error messages, based on the **SQLSTATE** code in the SQL Communications Area (SQLCA).

The definition of this function uses IN, OUT, and INOUT parameters:

```

CREATE function sysibm.SQLCAMessage (
    IN SQLCode      INTEGER,
    IN SQLErrm1    SMALLINT,
    IN SQLErrmc    VARCHAR(70),
    IN SQLErrp     CHAR(8),
    IN SQLErrd0    INTEGER,
    IN SQLErrd1    INTEGER,

```

```

IN SQLErrd2      INTEGER,
IN SQLErrd3      INTEGER,
IN SQLErrd4      INTEGER,
IN SQLErrd5      INTEGER,
IN SQLWarn       CHAR(11),
IN SQLState      CHAR(5),
IN MessageFileName VARCHAR(20),
INOUT Locale     VARCHAR(33),
OUT Message      LVARCHAR(4096),
OUT Rcode        INTEGER)
RETURNING INTEGER
EXTERNAL NAME '(SQLCAMessage)'
LANGUAGE C

```

To invoke the function, you can use this syntax:

sysibm.SQLCAMessage:

```

|—'sysibm'.SQLCAMessage —(—error_code— , —input_locale— , —message_file—)—|

```

Argument	Description	Restrictions	Syntax
<i>error_number</i>	The SQLCODE value of the error,	Must exist	"Expression" on page 4-40
<i>input_locale</i>	Name of the input locale for receiving the message. Default is the U.S. English locale (en_us).	Must exist	"Identifier " on page 5-21
<i>message_file</i>	Name of the message file	Must exist	Pathname

The function retrieves the text from the specified *message_file* for the specified **SQLCODE** and *input_locale*. The return code indicates the success or failure of the call to execute the **SQLCAMessage** routine.

The IDS DRDA application server attempts to retrieve the error message text using the specified input parameters:

- **SQLCODE**
- *input_locale*, and
- *message_file*

The Dynamic Server DRDA application server attempts to retrieve the error message text using the specified input parameters: The default message file (**errmsgtxt**) is used if the **MessageFileName** argument *message_file* is NULL. The default locale (**en_us**) is used to retrieve the error message if retrieval using the specified *input_locale* is unsuccessful. The token array is used to replace the tokens in the retrieved message text, if applicable.

If the retrieval is successful,

- **SQLCODE** is removed from the error message,
- the error message is copied to the OUT parameter '**Message**'
- the locale used for retrieving the message is copied to INOUT parameter '**Locale**'.

If the retrieval is unsuccessful, the error message text "Message not found" is copied to the **Message** parameter.

For both cases, the OUT parameter **Rcod** is set to the return code for executing this SPL routine.

Detailed message for **ISAM** errors are supplied from the **SQLERRD[0]** value. **ISAM** error messages are concatenated to actual error message string and returned to the application.

For the codes of **SQLSTATE** values for which the **SQLCAMEssage** function can return the corresponding error message text, see “ List of **SQLSTATE** Codes” on page 2-387.

Appendix A. Reserved Words for IBM Informix Dynamic Server

The ISO standard SQL language has many keywords. Some are designated as *reserved words* and others as *non-reserved words*. In ISO SQL, reserved words cannot be used as identifiers for database objects (like tables, columns, and so forth). To use such a name in a valid SQL statement requires a delimited identifier (“Delimited Identifiers” on page 5-23) that you enclose between double (" ") quotation marks.

In contrast, the dialect of SQL that IBM Informix database servers implement has very few *reserved words* in the sense of a character string that obeys the rules for identifiers (“Identifier ” on page 5-21) but always produces a compilation error or runtime error when used as an identifier. Your application might encounter restricted functionality, however, or unexpected results, if you define an SPL routine that has the same name as a built-in SQL function, expression, or operator.

This appendix lists the keywords in the IBM Informix implementation of SQL for Dynamic Server. Do not declare any of these keywords as SQL identifiers. If you do, errors or syntactic ambiguities can occur if the identifier appears in a context where the keyword is valid. In addition, your code will be more difficult to read and to maintain. Do not use keywords of C, C++, or other programming language that you will be using in an embedded mode) in your database structures.

If you receive an error message that seems unrelated to the SQL statement that caused the error, you might wish to review this appendix to see if a keyword has been used as an identifier.

To avoid using a keyword as an identifier, you can qualify the identifier with an owner name or modify the identifier. For example, rather than name a database object **CURRENT**, you might name it **o_current** or **juanita.current**. For a discussion of potential problems in using keywords as identifiers, and of additional workarounds for specific keywords, see “Use of Keywords as Identifiers” on page 5-23. See also *IBM Informix Guide to SQL: Tutorial* for more information about using keywords as identifiers in SQL applications.

A

ABS	ALLOCATE	AUDIT
ABSOLUTE	ALTER	AUTHORIZATION
ACCESS	AND	AUTHORIZED
ACCESS_METHOD	ANSI	AUTO
ACOS	ANY	AUTOFREE
ACTIVE	APPEND	AUTO_REPREPARE
ADD	ARRAY	AVG
ADD_MONTHS	AS	AVOID_EXECUTE
ADMIN AFTER	ASC	AVOID_FULL
AGGREGATE	ASCII	AVOID_HASH
ALIGNMENT	ASIN	AVOID_INDEX
ALL	AT	AVOID_INDEX_SJ
ALL_ROWS	ATAN	AVOID_NL
	ATAN2	AVOID_SUBQF
	ATTACH	

B

BEFORE	BITAND	BOOLEAN
BEGIN	BITANDNOT	BOTH
BETWEEN	BITNOT	BUFFERED
BIGINT	BITOR	BUILTIN
BIGSERIAL	BITXOR	BY
BINARY		BYTE

C

CACHE	CLUSTERSIZE	CONST
CALL	COARSE	CONSTRAINT
CANNOHASH	COBOL	CONSTRAINTS
CARDINALITY	CODESET	CONSTRUCTOR
CASCADE	COLLATION	CONTINUE
CASE	COLLECTION	COPY
CAST	COLUMN	COS
CEIL	COMMIT	COSTFUNC
CHAR	COMMITTED	COUNT
CHAR_LENGTH	COMMUTATOR	CRCOLS
CHARACTER_LENGTH	COMPONENT	CREATE
CHARACTER	COMPONENTS	CROSS
CHECK	CONCAT	CURRENT
CLASS	CONCURRENT	CURRENT_ROLE
CLIENT	CONNECT	CURRVAL
CLOSE	CONNECTION	CURSOR
CLUSTER		CYCLE

D

DATABASE	DECODE	DETACH
DATAFILES	DECRYPT_BINARY	DIAGNOSTICS
DATASKIP	DECRYPT_CHAR	DIRECTIVES
DATE	DEFAULT	DIRTY
DATETIME	DEFAULT_ROLE	DISABLED
DAY	DEFERRED	DISCONNECT
DBA	DEFERRED_PREPARE	DISTINCT
DBDATE	DEFINE	DISTRIBUTE_BINARY
DBINFO	DELAY	DISTRIBUTES_REFERENCES
DBPASSWORD	DELETE	DISTRIBUTIONS
DBSERVERNAME	DELETING	DOCUMENT
DBSECADM	DELIMITER	DOMAIN
DEALLOCATE	DELUXE	DONOTDISTRIBUTE
DEBUG	DESC	DORMANT
DEC	DESCRIBE	DOUBLE
DEC_T	DESCRIPTOR	DROP
DECIMAL		DTIME_T
DECLARE		

E

EACH	ENVIRONMENT	EXISTS
ELIF	ERROR	EXIT
ELSE	ESCAPE	EXP
ENABLED	EXCEPTION	EXPLAIN
ENCRYPT_AES	EXCLUSIVE	EXPLICIT
ENCRYPT_TDES	EXEC	EXPRESS
ENCRYPTION	EXECUTE	EXPRESSION
END	EXECUTEANYWHERE	EXTEND
ENUM	EXEMPTION	EXTENT
		EXTERNAL

F

FALSE	FIXCHAR	FORMAT_UNITS
FAR	FIXED	FORTRAN
FETCH	FLOAT	FOUND
FILE	FLOORFLUSH	FRACTION
FILETOBLOB	FOR	FRAGMENT
FILETOCLOB	FOREACH	FREE
FILLFACTOR	FOREIGN	FROM
FILTERING	FORMAT	FULL
FIRST		FUNCTION
FIRST_ROWS		

G

GENERAL	GLOBAL	GRANT
GET	GO	GREATERTHAN
GETHINT	GOTO	GREATERTHANOREQUAL
		GROUP

H

HANDLESNULLS	HEX	HOLD
HASH	HIGH	HOUR
HAVING	HINT	HYBRID

I

IDSLBACREADARRAY	IN	INSTEAD
IDSLBACREADSET	INACTIVE	INT
IDSLBACREADTREE	INCREMENTINDEX	INT8
IDSLBACRULES	INDEXES	INTEG
IDSLBACWRITEARRAY	INDEX_SJ	INTEGER
IDSLBACWRITESET	INDICATOR	INTERNAL
IDSLBACWRITETREE	INFORMIX	INTERNALLENGTH
IDSSECURITYLABEL	INIT	INTERVAL
IF	INITCAP	INTO
IFX_AUTO_REPREPARE	INLINE	INTRVL_T
IFX_INT8_T	INNER	IS
IFX_LO_CREATE_SPEC_T	INOUT	ISCANONICAL
IFX_LO_STAT_T	INSERT	ISOLATION
IMMEDIATE	INSERTING	ITEM
IMPLICIT		ITERATOR

J - K

JAVA	KEEP	KEY
JOIN		

L

LABEL	LENGTH	LOCK
LABELEQ	LESSTHAN	LOCKS
LABELGE	LESSTHANOREQUAL	LOCOPY
LABELGLB	LET	LOG
LABELGT	LEVEL	LOG10
LABELLE	LIKE	LOGN
LABELLT	LIMIT	LONG
LABELLUB	LIST	LOOP
LABELTOSTRING	LISTING	LOTOFILE
LANGUAGE	LOAD	LOW
LAST	LOC_T	LOWER
LAST_DAY	LOCAL	LPAD
LEADING	LOCATOR	LTRIM
LEFT		LVARCHAR

M

MATCHES	MEDIUM	MODERATE
MAX	MEMORY_RESIDENT	MODIFY
MAXERRORS	MIDDLE	MODULE
MAXLEN	MIN	MONEY
MAXVALUE	MINUTE	MONTH
MDY	MINVALUE	MONTHS_BETWEEN
MEDIAN	MOD	MOUNTING
	MODE	MULTISET

N

NAME	NOCACHE	NOT
NCHAR	NOCYCLE	NOTEMPLATEARG
NEGATOR	NOMAXVALUE	NOTEQUAL
NEW	NOMIGRATE	NULL
NEXT	NOMINVALUE	NULLIF
NEXT_DAY	NON_RESIDENT	NUMERIC
NEXTVAL	NONE	NVARCHAR
NO	NOORDER	NVL
	NORMAL	

O

OCTET_LENGTH	OPAQUE	OPTION
OF	OPCLASS	OR
OFF	OPEN	ORDER
OLD	OPERATIONAL	ORDERED
ON	OPTCOMPIND	OUT
ONLINE	OPTICAL	OUTER
ONLY	OPTIMIZATION	OUTPUT
		OVERRIDE

P

PAGE	PERCALL_COST	PREVIOUS
PARALLELIZABLE	PLI	PRIMARY
PARAMETER	PLOAD	PRIOR
PARTITION	POLICY	PRIVATE
PASCAL	POW	PRIVILEGES
PASSEDBYVALUE	POWER	PROCEDURE
PASSWORD	PRECISION	PUBLIC
PDQPRIORITY	PREPARE	PUT

R

RAISE	REPEATABLE	RIGHT
RANGE	REPLACE	ROBIN
RAW	REPLICATION	ROLE
READ	RESERVE	ROLLBACK
REAL	RESOLUTION	ROLLFORWARD
RECORDEND	RESOURCE	ROOT
REFERENCES	RESTART	ROUND
REFERENCING	RESTRICT	ROUTINE
REGISTER	RESUME	ROW
REJECTFILE	RETAIN	ROWID
RELATIVE	RETURN	ROWIDS
RELEASE	RETURNING	ROWS
REMAINDER	RETURNS	RPAD
RENAME	REUSE	RTRIM
REOPTIMIZATION	REVOKE	RULE

S

SAMEAS	SERVERUUID	STABILITY
SAMPLES	SESSION	STACK
SAMPLING	SET	STANDARD
SAVE	SETSESSIONAUTH	START
SAVEPOINT	SHARE	STATEMENT
SCHEDULE	SHORT	STATIC
SCHEMA	SIGNED	STATISTICS
SCRATCH	SIN	STDEV
SCROLL	SITENAME	STEP
SECLABEL_BY_COMP	SIZE	STOP
SECLABEL_BY_NAME	SKALL	STORAGE
SECLABEL_TO_CHAR	SKINHIBIT	STRATEGIES
SECOND	SKIP	STRING
SECONDARY	SKSHOW	STRINGTOLABEL
SECURED	SMALLFLOAT	STRUCT
SECURITY	SMALLINT	STYLE
SECTION	SOME	SUBSTR
SELCONST	SPECIFIC	SUBSTRING
SELECT	SQL	SUM
SELECTING	SQLCODE	SUPPORT
SELFUNC	SQLCONTEXT	SYNC
SELFUNCARGS	SQLERROR	SYNONYM
SEQUENCE	SQLSTATE	SYSDATE
SERIAL	SQLWARNING	SYSDBCLOSE
SERIAL8	SQRT	SYSDBOPEN
SERIALIZABLE		SYSTEM

T

TABLE	TO	TRIGGERS
TAN	TO_CHAR	TRIM
TASK	TO_NUMBER	TRUE
TEMP	TODAY	TRUNC
TEMPLATE	TRACE	TRUNCATE
TEST	TRAILING	TYPE
TEXT	TRANSACTION	TYPEDF
THEN	TREE	TYPEID
TIME	TRIGGER	TYPENAME
TIMEOUT		TYPEOF

U

UNCOMMITTED	UNLOAD	USAGE
UNDER	UNLOCK	USE_HASH
UNION	UNSIGNED	USE_NL
UNIQUE	UPDATE	USE_SUBQF
UNITS	UPDATING	USELASTCOMMITTED
UNKNOWN	UPPER	USER
		USING

V

VALUE
VALUES
VAR
VARCHAR

VARIABLE
VARIANCE
VARIANT
VARYING

VERCOLS
VIEW
VIOLATIONS
VOID
VOLATILE

W - Z

WAIT
WARNING
WEEKDAY
WHEN
WHENEVER
WHERE

WHILE
WITH
WITHOUT
WORK
WRITE
WRITEDOWN

WRITEDUP
XADATASOURCE
XID
XLOAD
XUNLOAD
YEAR

Appendix B. Reserved Words for IBM Informix Extended Parallel Server

The ISO standard SQL language has many keywords. Some are designated as *reserved words* and others as *non-reserved words*. In ISO SQL, reserved words cannot be used as identifiers for database objects (like tables, columns, and so forth). To use such a name in a valid SQL statement requires a delimited identifier (“Delimited Identifiers” on page 5-23) that you enclose between double (" ") quotation marks.

By contrast, the dialect of SQL that IBM Informix database servers implement has very few *reserved words* in the sense of a character string that obeys the rules for identifiers (“Identifier ” on page 5-21) but always produces a compilation error or runtime error when used as an identifier. Your application might encounter restricted functionality, however, or unexpected results, if you define an SPL routine that has the same name as a built-in SQL function, expression, or operator.

This appendix lists the keywords in the IBM Informix implementation of SQL for Extended Parallel Server. Do not declare any of these keywords as SQL identifiers. If you do, errors or syntactic ambiguities can occur if the identifier appears in a context where the keyword is valid. In addition, your code will be more difficult to read and to maintain. Do not use keywords of C or C++ (nor of any other programming language that you will be using in an embedded mode) in your database structures.

If you receive an error message that seems unrelated to the SQL statement that caused the error, you might wish to review this appendix to see if a keyword has been used as an identifier.

To avoid using a keyword as an identifier, you can qualify the identifier with an owner name or modify the identifier. For example, rather than name a database object **CURRENT**, you might name it **o_current** or **'juanita'.current**. For a discussion of potential problems in using keywords as identifiers, and of additional workarounds for specific keywords, see “Use of Keywords as Identifiers” on page 5-23. See also *IBM Informix Guide to SQL: Tutorial* for more information about using keywords as identifiers in SQL applications.

A

ADD	ANSI	ATTACH
AFTER	ANY	AUDIT
ALL	APPEND	AUTHORIZATION
ALL_MUTABLES	AS	AVG
ALTER	ASC	
AND	AT	

B

BEFORE	BINARY	BOUND_IMPL_PDQ
BEGIN	BITMAP	BUFFERED
BETWEEN	BOTH	BY
		BYTE

C

CACHE	CLOSE	CONNECT
CALL	CLUSTER	CONSTRAINT
CASCADE	CLUSTERSIZE	CONSTRAINTS
CASE	COARSE	CONTINUE
CHAR	COBOL	COPY
CHAR_LENGTH	CODESET	COUNT
CHARACTER	COLUMN	CREATE
CHARACTER_LENGTH	COMMIT	CURRENT
CHECK	COMMITTED	CURRENT_ROLE
CLIENT_TZ	COMPUTE_QUOTA	CURSOR

D

DATABASE	DEC	DELUXE
DATAFILES	DECIMAL	DESC
DATASKIP	DECLARE	DETACH
DATE	DECODE	DIRTY
DATETIME	DEFAULT	DISCONNECT
DAY	DEFAULT_ROLE	DISTINCT
DBA	DEFERRED	DISTRIBUTIONS
DBDATE	DEFINE	DOCUMENT
DBSERVERNAME	DELETE	DOUBLE
DEBUG	DELIMITER	DROP

E

EACH	EXCEPTION	EXPLAIN
ELIF	EXCLUSIVE	EXPRESS
ELSE	EXEC	EXPRESSION
END	EXECUTE	EXTEND
ENVIRONMENT	EXISTS	EXTENT
ERROR	EXIT	EXTERNAL
ESCAPE		

F

FETCH	FLOAT	FORTRAN
FILE	FOR	FOUND
FILLFACTOR	FOREACH	FRACTION
FILTERING	FOREIGN	FRAGMENT
FIRST	FORMAT	FROM
		FUNCTION

GGK
GLOBALGO
GOTOGRANT
GROUP

HHASH
HAVINGHIGH
HOLDHOUR
HYBRID

IIF
IMMEDIATE
IMMUTABLE
IMPLICIT_PDQ
IN
INDEX
INDEX_ALLINDICATOR
INIT
INITCAP
INNER
INSERT
INTINT8
INTEGER
INTERVAL
INTO IS
ISOLATION

J-K

JOIN

KEY

LLABELED
LABELGE
LABELGT
LABELLE
LABELLT
LANGUAGELEADING
LEFT
LET
LEVEL
LIKE
LISTING
LOADLOCAL
LOCK
LOCKS
LOG
LOW
LOWER

MMATCHES
MAX
MAXERRORS
MAXSCAN
MEDIUM
MEMORY_RESIDENTMERGE
MIDDLE
MIN
MINUTE
MODE
MODIFYMODULE
MONEY
MONTH
MOUNTING
MOVE
MUTABLE

NNCHAR
NEW
NEXT
NONON_RESIDENT
NONE
NORMAL
NOTNULL
NUMERIC
NVARCHAR
NVL

O

OCTET_LENGTH	ONLY	OPTION
OF	OPEN	OR
OFF	OPERATIONAL	ORDER
OLD	OPTICAL	OUTER
ON	OPTIMIZATION	

P

PAGE	PLOAD	PRIVATE
PASCAL	PRECISION	PRIVILEGES
PDQPRIORITY	PREPARE	PROCEDURE
PLI	PRIMARY	PUBLIC

R

RAISE	REMAINDER	RETURNS
RANGE	RENAME	REVOKE
RAW	REPEATABLE	RIDLIST
READ	RESERVE	ROBIN
REAL	RESOLUTION	ROLLBACK
RECORDEND	RESOURCE	ROLLFORWARD
RECOVER	RESTRICT	ROLE
REFERENCES	RESUME	ROUND
REFERENCING	RETAIN	ROW
REJECTFILE	RETURN	ROWS
RELEASE	RETURNING	

S

SAMEAS	SHARE	STANDARD
SAMPLES	SITENAME	START
SCHEDULE	SIZE	STATIC
SCHEMA	SKALL	STATISTICS
SCRATCH	SKINHIBIT	STDEV
SCROLL	SKSHOW	STEP
SECOND	SMALLFLOAT	STOP
SECTION	SMALLINT	SUBSTR
SELECT	SOME	SUBSTRING
SERIAL	SQL	SUM
SERIAL8	SQLCODE	SYNC
SERIALIZABLE	SQLERROR	SYNONYM
SET	SQLSTATE	SYSTEM
	SQLWARNING	

T

TABLE	TIMEOUT	TRANSACTION
TEMP	TMPSPACE_LIMIT	TRIGGER
TEMP_TABLE_EXT_SIZE	TO	TRIM
TEMP_TABLE_NEXT_SIZE	TODAY	TRUNC
TEXT	TRACE	TRUNCATE
THEN	TRAILING	TYPE

U

UNCOMMITTED	UNITS	UPDATE
UNION	UNLOAD	UPPER
UNIQUE	UNLOCK	USER
		USING

V

VALUES	VARIANT	VIEW
VARCHAR	VARYING	VIOLATIONS
VARIANCE		

W

WAIT	WHENEVER	WITH
WARNING	WHERE	WORK
WHEN	WHILE	WRITE

X-Z

XLOAD	XUNLOAD	YEAR
-------	---------	------

Appendix C. Accessibility

IBM strives to provide products with usable access for everyone, regardless of age or ability.

Accessibility features for IBM Informix Dynamic Server

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility Features

The following list includes the major accessibility features in IBM Informix Dynamic Server. These features support:

- Keyboard-only operation.
- Interfaces that are commonly used by screen readers.
- The attachment of alternative input and output devices.

Tip: The IBM Informix Dynamic Server Information Center and its related publications are accessibility-enabled for the IBM Home Page Reader. You can operate all features using the keyboard instead of the mouse.

Keyboard Navigation

This product uses standard Microsoft Windows navigation keys.

Related Accessibility Information

IBM is committed to making our documentation accessible to persons with disabilities. Our publications are available in HTML format so that they can be accessed with assistive technology such as screen reader software. The syntax diagrams in our publications are available in dotted decimal format. For more information about the dotted decimal format, go to "Dotted Decimal Syntax Diagrams."

You can view the publications for IBM Informix Dynamic Server in Adobe Portable Document Format (PDF) using the Adobe Acrobat Reader.

IBM and Accessibility

See the *IBM Accessibility Center* at <http://www.ibm.com/able> for more information about the commitment that IBM has to accessibility.

Dotted Decimal Syntax Diagrams

The syntax diagrams in our publications are available in dotted decimal format, which is an accessible format that is available only if you are using a screen reader.

In dotted decimal format, each syntax element is written on a separate line. If two or more syntax elements are always present together (or always absent together), the elements can appear on the same line, because they can be considered as a single compound syntax element.

Each line starts with a dotted decimal number; for example, 3 or 3.1 or 3.1.1. To hear these numbers correctly, make sure that your screen reader is set to read punctuation. All syntax elements that have the same dotted decimal number (for example, all syntax elements that have the number 3.1) are mutually exclusive alternatives. If you hear the lines 3.1 USERID and 3.1 SYSTEMID, your syntax can include either USERID or SYSTEMID, but not both.

The dotted decimal numbering level denotes the level of nesting. For example, if a syntax element with dotted decimal number 3 is followed by a series of syntax elements with dotted decimal number 3.1, all the syntax elements numbered 3.1 are subordinate to the syntax element numbered 3.

Certain words and symbols are used next to the dotted decimal numbers to add information about the syntax elements. Occasionally, these words and symbols might occur at the beginning of the element itself. For ease of identification, if the word or symbol is a part of the syntax element, the word or symbol is preceded by the backslash (\) character. The * symbol can be used next to a dotted decimal number to indicate that the syntax element repeats. For example, syntax element *FILE with dotted decimal number 3 is read as 3 * FILE. Format 3* FILE indicates that syntax element FILE repeats. Format 3* * FILE indicates that syntax element * FILE repeats.

Characters such as commas, which are used to separate a string of syntax elements, are shown in the syntax just before the items they separate. These characters can appear on the same line as each item, or on a separate line with the same dotted decimal number as the relevant items. The line can also show another symbol that provides information about the syntax elements. For example, the lines 5.1*, 5.1 LASTRUN, and 5.1 DELETE mean that if you use more than one of the LASTRUN and DELETE syntax elements, the elements must be separated by a comma. If no separator is given, assume that you use a blank to separate each syntax element.

If a syntax element is preceded by the % symbol, this identifies a reference that is defined elsewhere. The string following the % symbol is the name of a syntax fragment rather than a literal. For example, the line 2.1 %OP1 means that you should refer to a separate syntax fragment OP1.

The following words and symbols are used next to the dotted decimal numbers:

- ? Specifies an optional syntax element. A dotted decimal number followed by the ? symbol indicates that all the syntax elements with a corresponding dotted decimal number, and any subordinate syntax elements, are optional. If there is only one syntax element with a dotted decimal number, the ? symbol is displayed on the same line as the syntax element (for example, 5? NOTIFY). If there is more than one syntax element with a dotted decimal number, the ? symbol is displayed on a line by itself, followed by the syntax elements that are optional. For example, if you hear the lines 5 ?, 5 NOTIFY, and 5 UPDATE, you know that syntax elements NOTIFY and UPDATE are optional; that is, you can choose one or none of them. The ? symbol is equivalent to a bypass line in a railroad diagram.
- ! Specifies a default syntax element. A dotted decimal number followed by the ! symbol and a syntax element indicates that the syntax element is the default option for all syntax elements that share the same dotted decimal number. Only one of the syntax elements that share the same dotted decimal number can specify a ! symbol. For example, if you hear the lines

2? FILE, 2.1! (KEEP), and 2.1 (DELETE), you know that (KEEP) is the default option for the FILE keyword. In this example, if you include the FILE keyword but do not specify an option, default option KEEP is applied. A default option also applies to the next higher dotted decimal number. In this example, if the FILE keyword is omitted, default FILE(KEEP) is used. However, if you hear the lines 2? FILE, 2.1, 2.1.1! (KEEP), and 2.1.1 (DELETE), the default option KEEP only applies to the next higher dotted decimal number, 2.1 (which does not have an associated keyword), and does not apply to 2? FILE. Nothing is used if the keyword FILE is omitted.

- * Specifies a syntax element that can be repeated zero or more times. A dotted decimal number followed by the * symbol indicates that this syntax element can be used zero or more times; that is, it is optional and can be repeated. For example, if you hear the line 5.1* data-area, you know that you can include more than one data area or you can include none. If you hear the lines 3*, 3 HOST, and 3 STATE, you know that you can include HOST, STATE, both together, or nothing.

Notes:

1. If a dotted decimal number has an asterisk (*) next to it and there is only one item with that dotted decimal number, you can repeat that same item more than once.
2. If a dotted decimal number has an asterisk next to it and several items have that dotted decimal number, you can use more than one item from the list, but you cannot use the items more than once each. In the previous example, you could write HOST STATE, but you could not write HOST HOST.
3. The * symbol is equivalent to a loop-back line in a railroad syntax diagram.

- + Specifies a syntax element that must be included one or more times. A dotted decimal number followed by the + symbol indicates that this syntax element must be included one or more times. For example, if you hear the line 6.1+ data-area, you must include at least one data area. If you hear the lines 2+, 2 HOST, and 2 STATE, you know that you must include HOST, STATE, or both. As for the * symbol, you can only repeat a particular item if it is the only item with that dotted decimal number. The + symbol, like the * symbol, is equivalent to a loop-back line in a railroad syntax diagram.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licenses of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating

platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>.

Adobe, Acrobat, Portable Document Format (PDF), PostScript, and all Adobe-based trademarks are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Index

Special characters

- jar filename extension 2-120
- [] , brackets
 - array subscripts 2-543
 - range delimiters 4-14
 - substring operator 2-570, 4-64
- @ , at symbol 5-18
- > greater than symbol 4-193
- < less than symbol 4-193
- | , pipe character 2-110, 2-443, 2-577, 2-687
- || , concatenation operator 4-40, 4-56
- { } , braces
 - collection delimiters 4-179
 - comment indicators 1-3, 5-36
 - specifying empty collection 4-84
- _ , underscore 5-23
 - as wildcard character 4-13
- , double hyphen, comment indicator 1-3
- , hyphen symbol
 - DATETIME separator 4-181
 - INTERVAL separator 4-184
- , minus sign
 - binary operator 4-40
 - INTERVAL literals 4-184
 - unary operator 4-184
- :: , cast operator 4-58
- : , colon symbol
 - DATETIME separator 4-181
 - INTERVAL separator 4-184
- ! , exclamation point 4-193
 - in smart-large-object filename 4-126
- ? , question mark
 - as placeholder in PREPARE 2-463, 2-468
 - as wildcard 4-14
 - dynamic parameters 2-317
 - generating unique large-object filename 4-125
 - variables in PUT 2-481
- / , slash symbol
 - arithmetic operator 4-40
 - UNIX path separator 5-71
- /* */ , slash and asterisk
 - comment indicator 1-4, 2-470, 5-36
- . , decimal point
 - DATETIME separator 4-181
 - INTERVAL separator 4-177
 - literal numbers 4-185, 4-192
- . , period symbol
 - DATETIME separator 4-182
 - DECIMAL values 4-186
 - dot notation 4-61
 - INTERVAL separator 4-184
 - MONEY values 4-186
- " , double quotation marks
 - delimiting SQL identifiers 4-190
 - literal in a quoted string 4-192
 - quoted string delimiter 4-190, 4-192
 - with delimited identifiers 5-24, 5-26
- \$, dollar sign 5-23
- * , asterisk
 - all columns of a table 2-532, 2-695
 - all fields of a collection variable 3-25

- * , asterisk (*continued*)
 - all fields of a ROW column 2-540
 - all fields of a ROW value 4-61
 - all fields of a ROW variable 2-550
 - all labels of a security policy 2-513
 - arithmetic operator 4-40
 - in Projection clause 2-531
 - wildcard character 4-14
- \ , backslash
 - as wildcard character 4-13
 - escape character 2-687
- ' , single quotation marks 4-189
 - literal in a quoted string 4-192
- % , percent sign
 - as wildcard 4-13
- + , plus sign
 - binary operator 4-40
 - in optimizer directives 5-36
 - unary operator 4-184, 4-185
- = , equal sign
 - assignment operator 2-695
 - relational operator 4-194, 4-195
- ^ , caret
 - as wildcard character 4-14
- <<label>> statement 3-1

A

- ABS function 4-89, 4-90
- ABSOLUTE keyword, in FETCH statement 2-368
- ACCESS keyword
 - in ALTER TABLE statement 2-62
 - in CREATE TABLE statement 2-222
 - in GRANT statement 2-418
 - in INFO statement 2-428
 - in REVOKE statement 2-513
- Access method
 - attributes 5-51
 - configuring 5-51
 - default operator class, assigning 5-52
 - defined 5-51
 - directives 5-37
 - index 2-125
 - modifying 2-5
 - primary 2-225
 - privileges to alter 2-5
 - privileges to create 2-91
 - privileges to drop 2-322
 - purpose options 5-51
 - registering 2-90, 2-91
 - secondary 2-125
 - specifying for a table 2-224
 - sysams system catalog table settings 5-51
- Access mode for transactions 2-660
- ACCESS_METHOD keyword
 - in ALTER ACCESS_METHOD statement 2-5
 - in CREATE ACCESS_METHOD statement 2-90
 - in DROP ACCESS_METHOD statement 2-323
- accessibility C-1
 - keyboard C-1
 - shortcut keys C-1

- Accessibility C-1
 - dotted decimal format of syntax diagrams C-1
 - syntax diagrams, reading in a screen reader C-1
- ACOS function 4-139, 4-140
- Action clause, in CREATE TRIGGER statement
 - action list 2-257
 - syntax 2-251
- Active connection 2-321, 2-588
- ACTIVE keyword, in SAVE EXTERNAL DIRECTIVES statement 2-523
- Active set
 - constructing with OPEN statement 2-461
 - empty 2-375
 - retrieving with FETCH 2-369
 - sequential cursor 2-294
- ADD BUFFERPOOL command 6-6
- ADD CHUNK command 6-6
- ADD CONSTRAINT keywords
 - in ALTER TABLE statement 2-63
- ADD CRCOLS keywords
 - in ALTER TABLE statement 2-46
- ADD keyword
 - in ALTER ACCESS_METHOD statement 2-5
 - in ALTER FRAGMENT statement 2-23
 - in ALTER FUNCTION statement 2-28
 - in ALTER PROCEDURE statement 2-32
 - in ALTER ROUTINE statement 2-33
 - in ALTER SECURITY LABEL COMPONENT statement 2-36
 - in ALTER TABLE statement 2-46, 2-47, 2-72
- ADD LOG command 6-6
- ADD MEMORY command 6-6
- ADD MIRROR command 6-6
- ADD ROWIDS keywords
 - in ALTER TABLE statement 2-46
- ADD TYPE keywords
 - in ALTER TABLE statement 2-72
- ADD VERCOLS keywords
 - in ALTER TABLE statement 2-47
- ADD_MONTHS function 4-128, 4-129
- ADM VP class 6-38
- ADMIN function 6-2
- ADMIN_MODE_USERS configuration parameter 6-34, 6-41
- ADMIN_USER_MODE_WITH_DBSA configuration parameter 6-34
- ADT VP class 6-38
- ADTMODE configuration parameter 6-38
- Advanced Encryption Standard (AES) 4-114
- AES (Advanced Encryption Standard) 4-114
- AFTER keyword
 - in ALTER FRAGMENT statement 2-9, 2-23
 - in ALTER SECURITY LABEL COMPONENT statement 2-36
 - in CREATE TRIGGER statement 2-251, 2-257
- Aggregate functions
 - arguments 4-167
 - as arguments 5-3
 - AVG 4-165, 4-168
 - COUNT 4-165, 4-168
 - in ESQL 4-174
 - in EXISTS subquery 4-18
 - in expressions 2-539
 - in GROUP BY Clause 2-567
 - in ORDER BY clause 2-570
 - in SELECT statement 2-539
 - MAX 4-165, 4-172
 - MIN 4-165, 4-172
- Aggregate functions (*continued*)
 - RANGE 4-172
 - STDEV 4-173
 - SUM 4-165, 4-172
 - summary 4-174
 - syntax of built-in aggregates 4-165
 - using DISTINCT 4-168
 - VARIANCE 4-173
- AGGREGATE keyword
 - in CREATE AGGREGATE statement 2-92
 - in DROP AGGREGATE statement 2-323
- AIO VP class 6-37
- Algebraic functions
 - ABS 4-90
 - CEIL 4-90
 - FLOOR 4-90
 - MOD 4-91
 - POW 4-91
 - POWER 4-91
 - ROOT 4-91
 - ROUND 4-92
 - SQRT 4-91
 - TRUNC 4-95
- Alias
 - for a collection in SELECT statement 3-26
 - for a table in SELECT statement 2-545, 5-4
 - in optimizer directives 5-37
- ALIGNMENT keyword, in CREATE OPAQUE TYPE statement 2-147
- ALL keyword 2-424, 2-518, 2-528
 - beginning a subquery 2-564
 - in Condition segment 4-18
 - in CREATE INDEX statement 2-141
 - in CREATE VIEW statements 2-276
 - in DISCONNECT statement 2-320
 - in Expression segment 4-165, 4-175
 - in GRANT statement 2-408, 2-416, 2-418
 - in REVOKE statement 2-497, 2-504, 2-511, 2-513
 - in SELECT statement 2-535
 - in SET ENVIRONMENT statement 2-622
 - in SET Transaction Mode statement 2-591, 2-661
 - with UNION operator 2-579
- ALL_MUTABLES keyword, in SET ALL_MUTABLES statement 2-581
- ALL_ROWS keyword
 - in optimizer directives 5-42
 - in SET OPTIMIZATION statement 2-641
- ALLOCATE COLLECTION statement 2-1
- ALLOCATE DESCRIPTOR statement 2-2
- ALLOCATE ROW statement 2-4
- Allocating memory
 - for a collection variable 2-1
 - for system-descriptor area 2-2
- ALLOW_NEWLINE configuration parameter 4-191
- Allowing newline characters in quoted strings 4-191
- ALPHA class 4-156
- ALTER ACCESS_METHOD statement 2-5
- ALTER CHUNK OFFLINE command 6-6
- ALTER CHUNK ONLINE command 6-6
- ALTER FRAGMENT statement
 - am_readwrite purpose flag 5-52
 - hybrid fragmentation 2-17
 - INIT clause 2-18
 - privileges required 2-8
 - restrictions 2-8
 - reverting to nonfragmented table with DETACH 2-17

- ALTER FRAGMENT statement *(continued)*
 - reverting to nonfragmented table *(continued)*
 - with INIT 2-19
 - syntax 2-7
 - with generalized-key index 2-18
- ALTER FUNCTION statement 2-28
- ALTER INDEX statement
 - reclustering a table 2-30
 - specifying lock modes 2-31
- ALTER keyword
 - in GRANT statement 2-401, 2-408
 - in REVOKE statement 2-497, 2-504
- ALTER LOGMODE command 6-6
- ALTER PLOG command 6-6
- ALTER privilege 2-103, 2-401, 2-408, 2-497
- ALTER PROCEDURE statement 2-32
- ALTER ROUTINE statement 2-33
- ALTER SECURITY LABEL COMPONENT statement 2-36
- ALTER SEQUENCE statement 2-40
- ALTER TABLE statement
 - cascading deletes 2-53
 - changing column data type 2-57
 - dropping a column 2-55
 - privileges needed 2-44
 - restrictions
 - ADD clause 2-48
 - DROP clause 2-55
 - general 2-44
 - MODIFY clause 2-57
- am_beginscan purpose task 5-53
- am_check purpose task 5-53
- am_close purpose task 5-53
- am_cluster purpose flag 5-52
- am_create purpose task 5-53
- am_defopclass purpose value 5-52
- am_delete purpose task 5-53
- am_drop purpose task 5-53
- am_endscan purpose task 5-53
- am_getbyid purpose task 5-53
- am_getnext purpose task 5-53
- am_insert purpose task 5-53
- am_keyscan purpose flag 5-52
- am_open purpose task 5-53
- am_parallel purpose flag 5-52
- am_readwrite purpose flag 5-52
- am_rescan purpose task 5-53
- am_rowids purpose flag 5-52
- am_scancost purpose task 5-53
- am_sptype purpose value 5-52
- am_stats purpose task 5-53
- am_truncate access method 2-681
- am_truncate purpose task 5-53
- am_unique purpose flag 5-52
- am_update purpose task 5-53
- Ambiguities in non-unique identifiers 5-26
- AND bitwise logical operation 4-53
- AND keyword
 - in ANSI Joined Tables segment 2-555
 - in Condition segment 4-6, 4-20
 - in CREATE INDEX statement 2-143
 - with BETWEEN keyword 2-562
- ANDNOT bitwise logical operation 4-55
- Angle bracket (<< , , >>) symbols
 - loop label delimiters 3-2
 - statement label delimiters 3-2
- ANSI compliance
 - ansi compilation flag 2-288
- ANSI compliance *(continued)*
 - ansi flag 2-173, 2-193, 2-232
 - comment symbols 1-4
 - creating views 2-275
 - list of SQL statements 1-11
 - privileges on UDRs 2-406
 - renaming a table 2-493
 - table privileges 2-424
 - unbuffered logging 2-640
- ansi flag 2-100
- ANSI keyword
 - in CREATE DATABASE statement 2-97
- ANSI-compliance
 - escape character 4-14
 - implicit transactions 2-652
 - isolation level 2-374, 2-659
 - owner names 5-17
 - SQLSTATE codes 2-387
 - update cursors 2-290, 2-291, 2-307
 - warning after DELETE 2-310
- ANSI-compliant database
 - creating 2-99
 - database object naming 5-48
 - fixed-point DECIMAL values 2-99
 - implicit transactions 2-305
 - opaque-type naming 2-146
 - procedure name 2-715
 - table privileges 2-229
 - upshifting owner names 5-49
 - warning after opening 2-282
 - with BEGIN WORK 2-75
- ANSIOWNER environment variable 1-1, 2-99, 2-518, 5-49
- ANY keyword
 - in Condition segment 4-18
 - in SELECT statement 2-564
- APPEND keyword
 - in SET DEBUG FILE statement 2-600
 - in SET EXPLAIN statement 2-624
 - in SET PLOAD FILE statement 2-646
- Application
 - comment indicators 1-4
 - single-threaded 2-588
 - thread-safe 2-321, 2-588, 2-590
- Application partitioning 2-548
- ARCHIVE FAKE command 6-6
- Argument segment 5-1
- Arguments 5-1
- Arithmetic operators
 - binary 4-51
 - syntax 4-40
 - unary 4-52
- ARRAY keyword
 - in ALTER SECURITY LABEL COMPONENT statement 2-36
 - in CREATE SECURITY LABEL COMPONENT statement 2-177
- Array, with FETCH 2-372
- as DATETIME field qualifier 4-181
- as INTERVAL field qualifier 4-183
- AS keyword 2-531, 5-4, 5-55
 - in ALTER FRAGMENT statement 2-9
 - in CONNECT statement 2-83
 - in CREATE CAST statement 2-95
 - in CREATE DISTINCT TYPE statement 2-100
 - in CREATE INDEX statement 2-142
 - in CREATE TRIGGER statement 2-242
 - Delete triggers 2-254

AS keyword (*continued*)

- in CREATE TRIGGER statement (*continued*)
 - Insert triggers 2-255
 - Select triggers 2-256
 - Update triggers 2-256
 - view column values 2-270
- in CREATE VIEW statement 2-274
- in DROP CAST statement 2-324
- in explicit casts 4-58
- in GRANT FRAGMENT statement 2-423
- in GRANT statement 2-412
- in Iterator segment 2-551
- in MERGE statement 2-452
- in REVOKE FRAGMENT statement 2-517
- in REVOKE statement 2-495, 2-507
- in SELECT statement
 - ANSI table reference 2-555
 - FROM clause 2-545
 - in Iterator segment 2-551
 - with display labels 2-541
 - with table aliases 2-545

AS PARTITION keywords, in ALTER FRAGMENT statement 2-9

AS REMAINDER keywords, in ALTER FRAGMENT statement 2-9

ASC keyword

- in CREATE INDEX statement 2-127
- in SELECT statement 2-568, 2-570
- order with nulls 2-571

Ascending sequence 2-41, 2-185

ASCII code set 4-195

ASCII function 4-146

ASCII keyword

- in CREATE EXTERNAL TABLE statement 2-109
- in SELECT statement 2-577

ASIN function 4-139, 4-141

Assign support function 2-149, 2-437, 2-445, 2-448, 2-698

Associated statement 2-583

Asterisk (*)

- argument to COUNT function 4-169
- arithmetic operator 4-40
- in C-style comment indicators 1-4
- in projection list 3-25
- Projection clause 2-531
- wildcard character 4-14

AT keyword, in INSERT statement 2-430

At symbol (@) 5-16

ATAN function 4-139, 4-141

ATAN2 function 4-139, 4-141

ATTACH keyword, in ALTER FRAGMENT statement 2-9

Attached indexes 2-103, 2-136, 2-327, 2-586

Audit-event mnemonics 4-109

Authorization identifier 2-165, 2-334, 2-651, 4-72, 5-45

AUTHORIZATION keyword

- in CREATE SCHEMA statement 2-172
- in SET SESSION AUTHORIZATION statement 2-651

AUTHORIZED keyword

- in CREATE SECURITY POLICY statement 2-182

Auto Update Statistics (AUS) 2-706

AUTO_AIOVPS configuration parameter 6-41

AUTO_CKPTS configuration parameter 6-41

AUTO_LRU_TUNING configuration parameter 6-41

AUTO_REPREPARE configuration parameter 2-466, 2-478, 2-714

Autofree feature, in SET AUTOFREE 2-582

Automatic recompilation after table schema changes 2-620

AVG function 4-165, 4-168

AVOID_EXECUTE keyword

- in optimizer directives 5-44
- in SET EXPLAIN statement 2-624

AVOID_FULL keyword, in optimizer directives 5-37

AVOID_HASH keyword, in optimizer directives 5-40

AVOID_INDEX keyword, in optimizer directives 5-37

AVOID_INDEX_SJ keyword, in optimizer directives 5-37

AVOID_NL keyword, in optimizer directives 5-40

B

B keyword in SQL Admin API functions 6-2

B-tree cleaner list 2-708

B-tree index 2-51

- btree_ops operator class 2-153
- default operator class 2-153
- detached 2-136
- uses 2-133

B-tree scanner 6-32

B-tree secondary-access method 2-133, 2-151

Background mode 3-45

Backslash (\)

- as escape character 4-14
- as wildcard character 4-13
- escape character 2-687

Base-100 format 4-28

BASE64 encoding of encrypted data 2-612, 4-110

Batch file 3-46

BEFORE keyword

- in ALTER FRAGMENT statement 2-9, 2-23
- in ALTER SECURITY LABEL COMPONENT statement 2-36
- in ALTER TABLE statement 2-47
- in CREATE TRIGGER statement 2-251, 2-257

BEGIN keyword 3-30, 5-75

BEGIN WORK statement 2-74

BETWEEN keyword

- in Condition segment 4-6, 4-9

Big-endian format 2-106

BigDecimal data type of Java 4-29, 5-68

BIGINT data type 4-27

BIGSERIAL data type

- inserting values 2-436
- invalid default 2-195
- last inserted value 4-105
- value range 4-27

BINARY keyword

- in CREATE EXTERNAL TABLE statement 2-105

Binary operator 2-93, 4-51

Binary-format integers 2-106

Bit-hashing function 2-567

BITAND function 4-52, 4-53

BITANDNOT function 4-52, 4-55

BITMAP keyword, in CREATE INDEX statement 2-124

BITNOT function 4-52, 4-55

BITOR function 4-52, 4-54

Bitwise functions 4-52

BITXOR function 4-52, 4-54

Blank characters

- DATETIME separator 4-182
- in index names 2-330
- in literal numbers 4-185
- INTERVAL separator 4-184

BLOB data type 4-29, 4-30

- copying to a file 4-125
- copying to a smart large object 4-126
- creating from a file 4-123

- BLOB data type (*continued*)
 - default value 2-196
 - handle values 4-64
 - size limit 4-31
 - storing 2-62, 2-222
 - unloading 2-685, 2-686
 - unreferenced 2-679
- BLOB keyword
 - in Data Type segment 4-29
 - in DEFINE statement 3-8
- Blobspace 2-17
- BOOLEAN data type
 - defined 4-22
 - literal values 2-196
 - unloading 2-684
- Boolean expression 4-6
- BOTH keyword, in TRIM expressions 4-147
- BOUND_IMPL_PDQ keyword
 - in SET ENVIRONMENT statement 2-616
- Braces ({ })
 - collection delimiters 4-179
 - comment indicator 1-3, 5-36
 - specifying empty collection 4-84
- Brackets ([])
 - range delimiters 4-14
- BROADCAST optimizer directive 5-42
- BUFFERED keyword 2-640
 - in CREATE DATABASE statement 2-97
- BUFFERED LOG keyword
 - in CREATE DATABASE 2-99
- Buffered logging 2-640
- BUFFERPOOL configuration parameter 6-41
- Build table for hash joins 2-628
- Built-in aggregates
 - contrasted with user-defined 2-93
 - defined 4-167
 - extending 2-93
- Built-in data types
 - opaque 4-22
 - owner 2-100
 - privileges on 2-404
 - syntax 4-21
- Built-in roles
 - DBSECADM 2-414, 2-509
 - EXTEND 2-411
- Built-in routines
 - ADMIN 6-2
 - ALTER_JAVA_PATH 6-60
 - EXPLAIN_SQL 6-53
 - IFX_REPLACE_MODULE 6-54
 - IFX_UNLOAD_MODULE 6-55
 - INSTALL_JAR 6-58
 - JVPCONTROL 6-56
 - METADATA 6-62
 - REMOVE_JAR 6-59
 - REPLACE_JAR 6-59
 - SETUDTEXTNAME 6-61
 - SQLCMESSAGE 6-64
 - SYSDBCLOSE 6-50
 - SYSDBOPEN 6-50
 - TASK 6-2
 - UNSETUDTEXTNAME 6-62
- Built-in secondary-access method 2-132
- BY keyword
 - in ALTER FRAGMENT statement 2-20, 2-21
 - in ALTER SEQUENCE statement 2-41
 - in CREATE INDEX statement 2-136
- BY keyword (*continued*)
 - in CREATE SEQUENCE statement 2-185
 - in CREATE TABLE statement 2-212
 - in SELECT statement 2-566, 2-568
- BYTE and TEXT columns, fragment storage 2-13, 2-17
- BYTE column
 - modifying 2-58
- BYTE data
 - effect of isolation on retrieval 2-637, 2-660
 - loading 2-445
 - storage location 4-30
 - unloading 2-685
- BYTE data type
 - declaration syntax 4-30
 - default value 2-196
 - with SET DESCRIPTOR 2-611
 - with SPL routines 3-9, 3-16
- BYTE keyword 5-55
 - in Data Type segment 4-29
 - in DEFINE statement 3-8

C

- C keyword, in External Routine Reference segment 5-19
- C++ API 4-191
- Cache cohesion 2-654
- CACHE keyword
 - in ALTER SEQUENCE statement 2-42
 - in CREATE SEQUENCE statement 2-186
 - in SET STATEMENT CACHE statement 2-652
- Calculated expression
 - restrictions with GROUP BY 2-567
- Calendar 4-71
- CALL keyword, in WHENEVER statement 2-718
- CALL statement 3-3
- Callbacks 5-64
- CANNOTHASH keyword, in CREATE OPAQUE TYPE statement 2-147
- CANNOTHASH modifier 2-567
- CARDINALITY function 4-98
- Caret (^)
 - as wildcard character 4-14
 - use with brackets 4-14
- Cartesian product 5-40
- CASCADE keyword 2-494
 - in ALTER TABLE statement 2-52
 - in CREATE TABLE statement 2-199
 - in DROP SECURITY POLICY statement 2-338
 - in DROP TABLE statement 2-342
 - in DROP VIEW statement 2-346
 - in MOVE TABLE statement 2-455
- Cascading deletes 2-53
 - CREATE TABLE with 2-53
 - locking associated with 2-305
 - logging 2-305
 - multiple child tables 2-305
- Cascading triggers
 - and triggering table 2-264
 - triggered actions 2-253
- Case conversion
 - functions
 - INITCAP 4-156
 - LOWER 4-156
 - UPPER 4-156
- CASE keyword
 - in Expression segment 4-66, 4-67
- CASE statement 3-5

CAST keyword
 in CREATE CAST statement 2-95
 in DROP CAST statement 2-324
 in explicit casts 4-58

Casts
 built-in 2-96, 2-324
 creating 2-95
 dropping 2-324
 explicit 2-96, 4-58
 function for 2-97
 implicit 2-96, 4-58
 operator (::) 2-96, 4-58
 privileges 2-95
 registering 2-95
 symbol 4-58

cdrserver shadow column 2-46
cdrserver, replication column name 2-209, 2-235
cdrtime shadow column 2-46
cdrtime, replication column name 2-209, 2-235
CEIL function 4-89, 4-90
Chaining synonyms 2-190
CHAR data type
 defined 4-22
 in INSERT 4-192
 syntax 4-23
CHAR_LENGTH function 4-118, 4-119
CHARACTER data type
 syntax 4-23
Character data types
 fixed and varying length 4-22, 4-24
 multibyte characters 4-24
 syntax 4-23
CHARACTER VARYING data type
 syntax 4-23
CHARACTER_LENGTH function 4-118, 4-119
Check constraints
 defining 2-202
 moving 2-457
 reject files 2-111
 unloading data 2-111
CHECK DATA command 6-6
CHECK EXTENTS command 6-6
CHECK keyword
 in ALTER TABLE statement 2-54
 in CREATE EXTERNAL TABLE statement 2-107
 in CREATE TABLE statement 2-202
 in CREATE VIEW statement 2-274
CHECK PARTITION command 6-6
CHECKPOINT command 6-6
Child table 2-457
chmod utility 6-59
CLASS keyword, in Routine Modifier segment 5-60
CLASS_ORIGIN keyword, in GET DIAGNOSTICS
 statement 2-391
CLASSPATH environment variable 6-62
CLEAN SBSPACE command 6-6
Client APIs of Dynamic Server 4-191
CLIENT_LOCALE environment variables 4-25
CLIENT_TZ keyword
 in SET ENVIRONMENT statement 2-617, 4-75, 4-77
CLIENTBINVAL data type 2-536, 4-22
CLOB data type
 copying to a file 4-125
 copying to a smart large object 4-126
 creating from a file 4-123
 default value 2-196
 handle values 4-64
 CLOB data type (*continued*)
 size limit 4-31
 unloading 2-685, 2-686
 unreferenced 2-679
CLOB keyword
 in Data Type segment 4-29
 in DEFINE statement 3-8
CLOSE DATABASE statement 2-80
CLOSE statement
 closing a collection cursor 2-78
 closing a select cursor 2-77
 closing an insert cursor 2-78
 cursors affected by transaction end 2-79
 syntax 2-77
CLUSTER keyword
 in ALTER INDEX statement 2-30
 in CREATE INDEX statement 2-125
Clustered index 2-30, 2-126
Clustering, specifying support for 5-52
Co-locality 2-138
COARSE keyword
 in ALTER INDEX statement 2-30
 in CREATE INDEX statement 2-141
Code points, ASCII 4-195
CODESET keyword
 in CREATE EXTERNAL TABLE statement 2-109
 in SELECT statement 2-577
Cogroup name 2-109
Collation
 localized 2-584
 with relational operators 4-195
COLLATION keyword, in SET COLLATION statement 2-584
Collection constructors
 example 4-84
 restrictions 4-84
Collection cursor
 closing 2-78
 DECLARE for ESQL/C variable 2-298
 declaring 3-25
 defined 2-298
 in SPL 3-25
 inserting into 2-374, 2-483
 INTO clause 2-374
 opening 2-463
Collection data type 4-35
 allocating memory 2-1
 defining a column 4-35
 deleting 2-309
 element, searching for with IN 4-11
 IN operator 4-11
 LIST 4-36
 loading 2-447
 MULTISET 4-36
 returning number of elements 4-98
 selecting from 2-549
 SET 4-36
 unloading 2-684
 updating 5-10
COLLECTION keyword
 in ALLOCATE COLLECTION statement 2-1
 in DEALLOCATE COLLECTION statement 2-283
 in DEFINE statement 3-8
 untyped collection variable 2-550
Collection Subquery segment 4-3
Collection variable
 accessing 5-10
 accessing values 5-10

- Collection variable (*continued*)
 - associating cursor with 2-298
 - cursor for 2-374
 - deallocating memory for 2-283
 - in SELECT statement 2-550
 - manipulating values 5-10
 - opening a cursor 2-463
 - selecting from 2-549
 - selecting, inserting elements 2-298
 - untyped 2-2, 2-283
 - updating 2-375, 5-10
 - with DESCRIBE INPUT statement 2-319
 - with DESCRIBE statement 2-314
- Collection-derived table 5-4
 - collection cursor 2-299, 2-374, 2-483
 - collection variables with 5-10
 - FOREACH statement 3-26
 - in SELECT statement 5-10
 - INSERT statement with 2-300, 2-442
 - row types in 5-7
 - row variables with 5-14
 - SELECT statement with 2-549
 - fields from row variable 2-550
 - TABLE keyword 5-5, 5-10, 5-14
 - UPDATE statement with 2-703, 5-10, 5-14
 - where allowed 5-14
- Collection-Derived Table segment 5-4
- Collections
 - accessing a nested collection 5-14
 - accessing elements 5-5
 - allocating memory 2-1
 - constructors 4-84
 - deleting elements from 2-309
 - example of deleting elements 5-11
 - example of inserting elements 5-13
 - example of updating 5-12
 - generating values for 4-83
 - inserting values into 2-437
 - nested 4-180
 - restrictions when accessing elements 5-6
 - restrictions when defining 4-37
 - restrictions with inserting null values 2-437
 - selecting from 2-550
 - specifying literal values 4-179
 - updating 2-697
- Colon symbol (:)
 - cast operator (::) 4-58
 - DATETIME separator 4-181
 - INTERVAL separator 4-184
 - with database qualifier 5-17
- Column definition clause 2-193
- Column expression 2-575
- COLUMN keyword
 - in ALTER TABLE statement 2-55, 2-60
 - in CREATE TABLE statement 2-193
- COLUMN keyword, in RENAME COLUMN statement 2-487
- Column name
 - dot notation 4-61
 - using functions as names 5-26, 5-27
 - using keywords as names 5-27
- Column substring 4-63
- Column-level encryption 2-613, 4-108
- Columns
 - adding 2-47
 - adding a NOT NULL constraint 2-60
 - changing the data type 2-59
 - check constraints 2-108, 2-202
- Columns (*continued*)
 - defining as primary key 2-199
 - dropping 2-55, 2-56
 - expression 2-539, 4-59
 - inserting into 2-432
 - modifying with ALTER TABLE 2-57
 - number, effect on triggers 2-247
 - order in Projection list 2-532
 - primary or foreign key 2-199
 - privileges 2-401
 - projection 4-61
 - referenced and referencing 2-200
 - removing a NOT NULL constraint 2-60
 - renaming 2-487
 - specifying a subscript 2-570, 4-63
 - virtual 2-276
- COLUMNS keyword, in INFO statement 2-428
- COMBINE keyword
 - in CREATE AGGREGATE statement 2-92
- Comma (,) symbol in pathnames 4-124
- Command file 1-4
- command_history table 6-2, 6-49
- Comment symbol
 - braces ({}) 1-3
 - double hyphen (--) 1-3
 - in application programs 1-4
 - in optimizer directives 5-36
 - in prepared statements 2-470
 - slash and asterisk (/ * */) 1-4
- COMMIT WORK statement
 - in ANSI-compliant databases 2-82
 - in non-ANSI databases 2-82
 - syntax 2-81
- COMMITTED keyword
 - in SET ENVIRONMENT statement 2-622
 - in SET ISOLATION statement 2-631, 2-634
 - in SET TRANSACTION statement 2-658
- Committed Read isolation level 2-622, 2-634
- COMMITTED READ keywords
 - in SET ENVIRONMENT statement 2-622
 - in SET ISOLATION statement 2-631
- COMMITTED READ setting, of USELASTCOMMITTED
 - configuration parameter 2-634
- Compacted index 2-133
- Companion functions 2-357, 5-63
- Compare support function 2-149
- Complete-connection level settings
 - of SET EXPLAIN 2-628
 - of SET ISOLATION 2-633
 - of SET LOCK MODE 2-639
- Complex data type
 - invalid in distributed queries 2-537
 - loading element values 2-447
 - unloading 2-687
- Complex numbers 4-29
- Complex table expression 2-546
- Complex view 2-267
- COMPONENT keyword
 - in ALTER SECURITY LABEL COMPONENT statement 2-36
 - in CREATE SECURITY LABEL COMPONENT statement 2-177
 - in CREATE SECURITY LABEL statement 2-175
 - in DROP SECURITY statement 2-338
 - in RENAME SECURITY statement 2-490
- COMPONENTS keyword
 - in CREATE SECURITY POLICY statement 2-182

- Composite key 2-208
- Compound assignment 3-32
- COMPUTE_QUOTA keyword
 - in SET ENVIRONMENT statement 2-618
- CONCAT function 4-142
- CONCAT() operator function 4-57
- Concatenation operator (||) 4-40, 4-56
- Concurrency
 - with CREATE INDEX 2-143
 - with DROP INDEX 2-330
 - with SET ENVIRONMENT statement 2-622
 - with SET ISOLATION 2-632
 - with SET TRANSACTION 2-660
 - with START VIOLATIONS TABLE 2-665
- CONCURRENT keyword, in CONNECT statement 2-83
- Condition segment
 - DELETING 4-12
 - INSERTING 4-12
 - IS NOT NULL 4-12
 - IS NULL 4-12
 - join conditions 2-564
 - keywords
 - ALL 4-18
 - ANY 4-18
 - BETWEEN 4-9
 - EXISTS 4-18
 - LIKE 4-13
 - MATCHES 4-13
 - NOT 4-13
 - SOME 4-18
 - null values 4-19
 - SELECTING 4-12
 - subquery in SELECT 4-15
 - syntax 4-5
 - UPDATING 4-12
- Conditional expressions
 - CASE 4-65
 - DECODE 4-65
 - NVL 4-65
- Conditions
 - comparison 4-6, 4-8
 - IN operator 4-10
 - NOT IN operator 4-10
- Configuration parameters
 - ADMIN_MODE_USERS 6-34, 6-41
 - ADMIN_USER_MODE_WITH_DBSA 6-34
 - ADTMODE 6-38
 - ALLOW_NEWLINE 4-191
 - AUTO_AIOVPS 6-41
 - AUTO_CKPTS 6-41
 - AUTO_LRU_TUNING 6-41
 - AUTO_REPREPARE 2-466, 2-478, 2-714
 - BUFFERPOOL 6-41
 - DATASKIP 2-598
 - DB_LIBRARY_PATH 5-71
 - DBCREATE_PERMISSION 2-98
 - DBSERVERALIASES 2-310, 2-439, 2-538, 2-699, 6-32, 6-46, 6-62
 - DBSERVERNAME 2-310, 2-439, 2-538, 2-699, 6-32, 6-46
 - DEADLOCK_TIMEOUT 2-330
 - DEF_TABLE_LOCKMODE 2-69, 2-226, 2-452, 2-634
 - DIRECTIVES 5-36
 - DS_ADM_POLICY 2-650
 - DS_MAX_QUERIES 6-36, 6-39, 6-41
 - DS_MAX_SCANS 6-41
 - DS_NONPDQ_QUERY_MEM 2-570, 6-41
 - DS_TOTAL_MEMORY 6-36, 6-41
- Configuration parameters (*continued*)
 - EXPLAIN_STAT 2-628, 6-41
 - EXT_DIRECTIVES 2-523, 2-620, 5-45
 - FILLFACTOR 2-134
 - IFX_EXTEND_ROLE 2-115, 2-123, 2-411, 2-507, 6-41, 6-54
 - JVPCCLASSPATH 6-60
 - LISTEN_TIMEOUT 6-41
 - LOCKS 2-326, 2-451
 - LOG_INDEX_BUILDS 6-46
 - MAX_INCOMPLETE_CONNECTIONS 6-41
 - MAX_PDQPRIORITY 2-643, 6-23, 6-33, 6-41
 - ONDBSPACEDOWN 6-36
 - ONLIDX_MAXMEM 6-41
 - OPT_GOAL 5-43
 - OPTCOMPIND 2-653, 5-42
 - RESIDENT 6-23, 6-39, 6-41
 - RTO_SERVER_RESTART 6-23, 6-41
 - SBSPACENAME 2-222, 4-114, 4-115
 - SDS_TIMEOUT 6-46
 - SEQ_CACHE_SIZE 2-42, 2-187
 - SHMADD 6-30
 - STACKSIZE 5-66
 - STMT_CACHE 2-652, 6-34
 - STMT_CACHE_HITS 2-654, 6-40
 - STMT_CACHE_NOLIMIT 2-654, 6-40
 - STMT_CACHE_SIZE 2-654
 - SYSSBSPACENAME 2-708
 - TEMPTAB_NOLOG 2-232
 - TXTIMEOUT 6-44
 - USELASTCOMMITTED 2-450, 2-622, 2-634, 2-658, 6-23, 6-41
 - USEOSTIME 4-75
 - VP_MEMORY_CACHE_KBS 6-23, 6-41
 - VPCLASS 6-37, 6-38
- Conflict resolution 2-46
- CONNECT keyword 2-397, 2-496
- Connect privilege 2-496
 - granting 2-398
- CONNECT statement 2-83
- CONNECTION keyword, in SET CONNECTION statement 2-587
- Connection URL 6-64
- CONNECTION_NAME keyword, in GET DIAGNOSTICS statement 2-391
- Connections
 - active 2-321, 2-588
 - context 2-320, 2-588
 - current 2-84, 2-321, 2-590
 - default 2-88, 2-321, 2-589
 - dormant 2-84, 2-320, 2-587
 - explicit 2-81, 2-88, 2-326
 - implicit 2-88, 2-98, 2-282, 2-321, 2-589
 - returning connection names 2-393
 - single-threaded applications 2-588
 - specifying connection names 2-86, 2-589
- Constant expression 4-70
 - in SELECT 2-539
 - inserting with PUT 2-480
- CONSTRAINT keyword
 - in ALTER TABLE statement 2-51, 2-63, 2-66
 - in CREATE TABLE statement 2-203
- Constraints
 - adding primary-key 2-65
 - adding referential 2-65
 - adding to a column with data 2-61
 - adding unique 2-65
 - adding with ALTER TABLE 2-63, 2-64

- Constraints (*continued*)
 - affected by dropping a column from table 2-56
 - B-tree indexes 2-126
 - checking 2-268, 2-454, 2-661
 - detached checking 2-209
 - disabled 2-204, 2-595
 - dropping 2-66
 - enabled 2-204
 - encountering violations while adding 2-66
 - filtering 2-204
 - filtering to violations table 2-595
 - foreign key 2-508
 - limit on size 2-206
 - mode 2-204
 - modifying a column that has constraints 2-57
 - moving 2-457
 - multiple-column 2-205
 - name 2-203
 - number of columns allowed 2-206
 - privileges needed to create 2-65
 - referential 2-199, 2-673
 - renaming 2-489
 - single-column 2-197
 - system catalog tables 2-203
 - transaction mode 2-661
- CONSTRAINTS keyword
 - in SET CONSTRAINTS statement 2-591
 - in SET Database Object Mode statement 2-593
 - in SET Transaction Mode statement 2-661
- constrid column 2-19
- Constructors
 - functions
 - collections 4-83
 - row 4-82
- CONTINUE keyword
 - in WHENEVER statement 2-718
- CONTINUE statement 3-7
- Coordinated Universal Time (UTC) 4-106
- Correlated subquery 2-546
 - defined 4-16
- Correlation name 2-545
 - declaring 2-116, 2-157, 2-254
 - in routines 2-265
 - qualifying values 2-262
 - scope of reference 2-261
- COS function 4-139, 4-140
- Coserver number 2-109, 4-106
- COSTFUNC keyword, in Routine Modifier segment 5-60
- COUNT field 2-381
 - in ALLOCATE DESCRIPTOR statement 2-3
 - with DESCRIBE INPUT statement 2-318
 - with DESCRIBE statement 2-313
- COUNT function
 - defined 4-168
 - restriction with CREATE TRIGGER statement 2-261
 - syntax 4-166
- COUNT keyword
 - in SET DESCRIPTOR statement 2-605
- CPU usage cost 5-64
- CPU VP class 2-618, 5-61, 6-38
- CRCOLS keyword
 - in ALTER TABLE statement 2-46, 2-47
 - in CREATE TABLE statement 2-209
 - in CREATE Temporary TABLE statement 2-235
- CREATE ACCESS_METHOD statement 2-91
- CREATE AGGREGATE statement 2-92
- CREATE BLOBSPACE command 6-6
- CREATE CAST statement 2-95
- CREATE CHUNK command 6-6
- CREATE DATABASE statement
 - ANSI compliance 2-99
 - syntax 2-97
 - using with PREPARE 2-98
- CREATE DBSPACE command 6-6
- CREATE DISTINCT TYPE statement 2-100
- CREATE DUPLICATE statement 2-102
- CREATE EXTERNAL TABLE statement 2-104
- CREATE FUNCTION FROM statement 2-122
- CREATE FUNCTION statement 2-113
- CREATE INDEX statement
 - cluster with fragments 2-126
 - composite indexes 2-129
 - disabled indexes 2-140
 - index-type options 2-125
 - sort order 2-129
 - specifying lock modes 2-141
 - specifying object modes 2-138
 - storage options 2-134
- CREATE OPAQUE TYPE statement 2-145
- CREATE OPCLASS statement 2-149
- CREATE PROCEDURE FROM statement 2-163
- CREATE PROCEDURE statement 2-153
- CREATE ROLE statement 2-164
- CREATE ROUTINE FROM statement 2-166
- CREATE ROW TYPE statement 2-168
- CREATE SBSPACE command 6-6
- CREATE SCHEMA statement 2-172
- CREATE SCRATCH TABLE statement 2-174
- CREATE SECURITY LABEL COMPONENT statement 2-177
- CREATE SECURITY LABEL statement 2-175
- CREATE SECURITY POLICY statement 2-181
- CREATE SEQUENCE statement 2-184
- CREATE SYNONYM statement 2-187
- CREATE TABLE statement
 - access-method option 2-224
 - column definition clause 2-193
 - column-level constraints 2-234
 - constraints
 - check 2-202
 - composite keys 2-205, 2-208
 - defining 2-197
 - distinct 2-198
 - example 2-207
 - NOT NULL 2-195, 2-198
 - primary key 2-199
 - referential 2-199
 - restrictions 2-198
 - unique 2-198
 - creating composite columns 2-205
 - defining constraints 2-205
 - fragmenting
 - by expression 2-214
 - by hash 2-216
 - by hybrid 2-217
 - by range 2-218
 - round-robin 2-214
 - locking options 2-225
- ON DELETE CASCADE keywords 2-53
- OPERATIONAL keyword 2-193
- PUT clause 2-222
- RAW keyword 2-193
- specifying cascading deletes 2-201
- specifying column-default values 2-195
- specifying storage location 2-210

CREATE TABLE statement (*continued*)
 STANDARD keyword 2-193
 STATIC keyword 2-193
 syntax 2-191
 WITH ROWIDS keywords 2-214

CREATE TEMP TABLE statement 2-230

CREATE TEMPDBSPACE command 6-6

CREATE Temporary TABLE statement 2-231

CREATE TRIGGER statement 2-239

CREATE VIEW statement 2-274, 2-547

CREATE WITH_CHECK CHUNK command 6-6

CREATE WITH_CHECK DBSPACE command 6-6

CREATE WITH_CHECK SBSPACE command 6-6

CREATE WITH_CHECK TEMPDBSPACE command 6-6

CREATE XADATASOURCE statement 2-279

CREATE XADATASOURCE TYPE statement 2-280

Creation-time settings of environment variables 2-258

Cross joins 2-553

CROSS keyword in SELECT statement 2-557

Cross-database DML operations 2-310, 2-439, 2-538, 2-699

CRPT audit-event mnemonic 4-109

CSN encryption 4-108

CTRL-J
 newline
 preserving in quoted strings 4-191

Current database, specifying with DATABASE 2-281

Current date 4-75, 4-77

CURRENT DORMANT keywords, in SET CONNECTION statement 2-587

CURRENT function
 as an argument 4-76
 as constant expression 4-71
 in ALTER TABLE statement 2-49
 in Condition segment 4-10
 in CREATE TABLE statement 2-195
 in DEFINE statement 3-10
 in INSERT statement 2-439
 in WHERE condition 4-76

CURRENT keyword 2-689
 in DELETE statement 2-303
 in DISCONNECT statement 2-321
 in FETCH statement 2-368
 in SET CONNECTION statement 2-587
 in UPDATE statement 2-702

CURRENT_ROLE operator
 defined 4-73
 syntax 4-71
 usage 4-73

CURRVAL operator 2-185, 4-79, 4-80

Cursor
 activating with OPEN 2-459
 affected by transaction end 2-79
 characteristics 2-293
 closing 2-77
 closing with ROLLBACK WORK 2-521
 declaring 2-286
 direct 2-78
 for update
 restricted statements 2-296
 using in ANSI-mode databases 2-296, 2-297
 using in non-ANSI databases 2-296, 2-297
 freeing automatically with SET AUTOFREE 2-582
 implicit 3-23
 manipulation statements 1-8
 opening 2-460, 2-461
 prepared statement with 2-297

Cursor (*continued*)
 read-only
 restricted statements 2-296
 using in ANSI-mode databases 2-296, 2-297
 using in non-ANSI databases 2-296, 2-297
 where required 2-573
 retrieving values with FETCH 2-367
 select hold examples 2-295
 sequence of program operations 2-289
 stability 2-657
 statement identifier with 2-297
 types of 2-460
 with INTO keyword in SELECT 2-542
 with transactions 2-300

Cursor function 3-24, 5-59

CURSOR keyword
 in DECLARE statement 2-286
 in SET ISOLATION statement 2-631

Cursor Stability isolation level 2-635

CURSOR STABILITY keywords, in SET ISOLATION statement 2-631

CYCLE keyword
 in ALTER SEQUENCE statement 2-42
 in CREATE SEQUENCE statement 2-186

D

Dangling child records 2-306

Data
 access statements 1-9
 confidentiality 2-611
 data definition statements 1-6
 data manipulation statements 1-8
 encryption 4-107
 inserting with LOAD 2-443
 integrity statements 1-8

Data buffering 2-377

data column of sysprocbody table 4-109

Data distributions
 confidence level 2-710
 on temporary tables 2-706

DATA field 2-381
 in SET DESCRIPTOR statement 2-606
 with DESCRIBE INPUT statement 2-318
 with DESCRIBE statement 2-313

Data replication 2-46

Data type segment 4-21

Data types 4-21
 alignment 2-101
 casting 2-95, 4-58
 changing with ALTER TABLE 2-59
 collection 4-35
 complex 4-34
 considerations for INSERT 2-436, 4-192
 distinct 4-32
 opaque 2-145
 representation 2-101
 simple large object 4-30
 specifying with CREATE VIEW 2-275
 varying-length 4-24

Data-integrity violations 2-664, 2-676

Database administrator 2-399

Database administrator (DBA) 2-496
 granting privileges 2-397

DATABASE keyword
 in CLOSE DATABASE statement 2-80
 in CREATE DATABASE statement 2-97

DATABASE keyword (*continued*)
 in DATABASE statement 2-281
 in DROP DATABASE statement 2-325
 in MOVE TABLE statement 2-455
 in RENAME DATABASE statement 2-488
 Database object
 naming 5-16
 owner 5-45
 Database object mode
 for triggers 2-242, 2-595
 privileges required 2-592
 specifying 2-591
 Database Object Name segment 5-16
 Database Server Administrator 5-20
 Database Server Administrator (DBSA) 2-115, 2-166, 2-328, 2-333, 2-336, 2-411, 2-506, 6-2
 Database servers
 bringing online from quiescent mode 6-35
 unblocking 6-30
 DATABASE statement
 determining database type 2-282
 exclusive mode 2-282
 specifying current database 2-281
 SQLWARN after 2-282
 syntax 2-281
 Database statements 2-88
 database-level 2-397, 2-496
 Database-level privilege 2-411
 not available for roles 2-496
 revoking 2-496
 Databases
 ANSI-compliant 2-282
 closing with CLOSE DATABASE 2-80
 data warehousing 2-193
 default isolation levels 2-635, 2-659
 dropping 2-326
 external 5-17
 global variables 3-10
 isolation level 2-657
 lock 2-282
 naming conventions 5-15
 nonlogging 2-660
 nonlogging database 2-267
 OLTP 2-193
 opening in exclusive mode 2-282
 optimizing queries 2-706
 read-only mode 2-573
 remote 5-16
 renaming 2-488
 running in secondary mode 2-282
 DataBlade API 5-64
 DataBlade API (LIBDMI) 4-191
 DataBlade Developers Kit 5-71
 DataBlade modules 2-153
 DATAFILES keyword
 in CREATE EXTERNAL TABLE statement 2-108
 in SELECT statement 2-574
 DATASKIP configuration parameter 2-598
 DATASKIP keyword
 in SET DATASKIP statement 2-598
 DATE data type
 declaration syntax 4-32
 functions 4-128
 literal DATE values 2-196
 DATE function 4-128, 4-130
 DATETIME data type 4-32, 4-38, 4-181
 as quoted string 4-192
 DATETIME data type (*continued*)
 cast to Coordinated Universal Time (UTC) 4-106
 functions 4-128
 in INSERT 4-192
 literal values 4-78
 DATETIME data types 4-181
 DATETIME Field Qualifier segment 4-38
 DATETIME keyword, in Literal DATETIME 4-181
 datetime.h header file 2-609
 DAY function 4-128, 4-130
 DAY keyword 4-38, 4-176, 4-181, 4-183
 in DATETIME Field Qualifier 4-39
 DB_LIBRARY_PATH configuration parameter 5-71
 DB_LOCALE environment variable 2-136, 2-282, 2-585, 2-628, 5-23
 DB_LOCALE environment variables 4-25
 DBA keyword 2-397, 2-496
 in CREATE FUNCTION statement 2-113
 in CREATE PROCEDURE statement 2-153
 DBA privilege
 with CREATE ACCESS_METHOD statement 2-91
 with CREATE SCHEMA 2-173
 with DROP DATABASE 2-325
 with DROP TRIGGER statement 2-344
 with REVOKE statement 2-496
 DBA-privileged UDR 2-115, 2-156
 DBA. 2-399
 DBACCNOIGN environment variable 2-75
 DBANSIWARN environment variable 2-100, 2-173, 2-193, 2-232, 2-275
 DBBLOBBUF environment variable 2-446, 2-686
 DBCENTURY environment variable 2-15, 2-258, 2-444, 4-182
 DBCREATE_PERMISSION configuration parameter 2-98
 DBDATE environment variable 2-196, 2-684, 4-192
 DBDELIMITER environment variable 2-448, 2-687
 dbhostname option of DBINFO 4-104
 DBINFO function 4-99
 DBMONEY environment variable 2-445, 2-685, 4-186
 DBPATH environment variable 2-86, 2-88
 DBSA group 2-411, 6-2
 DBSA. 2-411
 dbschema utility 2-42
 DBSECADM keyword
 in GRANT statement 2-414
 in REVOKE statement 2-509
 DBSECADM role 2-17
 dbsendrecv data type 2-147
 DBSERVERALIASES configuration parameter 2-310, 2-439, 2-538, 2-699, 6-32, 6-46, 6-62
 DBSERVERNAME configuration parameter 2-310, 2-439, 2-538, 2-699, 6-32, 6-46
 DBSERVERNAME function
 constant expression 4-74
 in ALTER TABLE statement 2-49
 in Condition segment 4-10
 in CREATE TABLE statement 2-195
 in DEFINE statement 3-10
 dbspace
 renaming with onspaces 2-8
 dbspaces
 number 4-118
 skipping if unavailable 2-598
 DBSPACETEMP configuration parameter 2-576
 DBSPACETEMP environment variable 2-138, 2-237, 2-576, 2-714, 3-16
 DBTIME environment variable 2-445, 2-685
 DBUPSPACE environment variable 2-712, 2-714

- DDL (Data Definition Language) statements
 - listed 1-6
- Deadlock 2-374, 3-45
- Deadlock detection 2-638
- DEADLOCK_TIMEOUT setting in ONCONFIG 2-638
- DEALLOCATE COLLECTION statement 2-283
- DEALLOCATE DESCRIPTOR statement 2-284
- DEALLOCATE ROW statement 2-285
- DEBUG environment variable 6-52
- DEBUG keyword
 - in SET DEBUG FILE statement 2-599
- Debugging sysdbopen() routines 6-52
- DECIMAL data type 4-27, 4-29
 - literal values 4-186
- Decimal point (.)
 - DATETIME separator 4-181
 - INTERVAL separator 4-177
 - literal numbers 4-185, 4-192
- declaration syntax 4-29
- Declarative statements 3-39
- DECLARE statement
 - cursor characteristics 2-293
 - CURSOR keyword 2-294
 - cursors with prepared statements 2-297
 - cursors with transactions 2-300
 - FOR UPDATE keywords 2-290
 - function cursor 2-289
 - insert cursor 2-292, 2-296
 - restrictions with SELECT with ORDER BY 2-572
 - SCROLL keyword 2-294
 - select cursor 2-289
 - syntax 2-286
 - updating specified columns 2-291
 - WHERE CURRENT OF clause 2-290
 - WITH HOLD keywords 2-294
 - with SELECT statement 2-543
- DECODE function 4-69
- DECRYPT_BINARY function 4-114
- DECRYPT_CHAR function 4-113
- DEF_TABLE_LOCKMODE configuration parameter 2-69, 2-452, 2-634
- Default connection 2-88, 2-321
- Default database server 2-88, 2-321, 2-393
- Default escape character 2-687
- Default isolation level 2-659
- DEFAULT keyword
 - in ALTER TABLE statement 2-49
 - in CONNECT statement 2-83
 - in CREATE EXTERNAL TABLE statement 2-109
 - in CREATE TABLE statement 2-195
 - in DEFINE statement 3-8
 - in DISCONNECT statement 2-321
 - in GRANT statement 2-410
 - in REVOKE statement 2-506
 - in SET CONNECTION statement 2-587, 2-589
 - in SET DATASKIP statement 2-598
 - in SET Default Table Space statement 2-601
 - in SET Default Table Type statement 2-602
 - in SET ENVIRONMENT statement 2-616, 2-620, 2-621
 - in SET PDQPRIORITY statement 2-643
 - in SET ROLE statement 2-649
- Default role 2-166, 2-410, 2-506
- Default user ID 2-88
- DEFAULT_ATTACH environment variable 2-136
- DEFAULT_ROLE operator
 - defined 4-74
 - syntax 4-71
- DEFAULT_ROLE operator (*continued*)
 - usage 4-73
- DEFERRED keyword, in SET Transaction Mode statement 2-661
- DEFERRED_PREPARE keyword
 - in SET DEFERRED_PREPARE statement 2-603
- Deferred-Prepare feature 2-603
- DEFINE keyword, in Statement Block segment 5-75
- DEFINE statement
 - default value clause 3-10
 - syntax 3-8
- Delete join 2-309
- DELETE keyword 2-424, 2-518
 - in CREATE TABLE statement 2-200
 - in CREATE TRIGGER statement 2-246
 - in GRANT statement 2-400
 - in REVOKE statement 2-497
- Delete privilege 2-400, 2-497
- DELETE statements
 - and triggers 2-259
 - cascading 2-305
 - collection columns with 2-309
 - cursor with 2-290
 - distributed 2-309
 - OUT parameter 4-165
 - restrictions when using a join 2-308
 - syntax 2-303
 - using joins 2-308
 - with SELECT... FOR UPDATE 2-572
 - with update cursor 2-307
 - within a transaction 2-304
- Delete trigger 2-244, 2-271
- Deleting from a specific table in a table hierarchy 2-305
- DELETING operator 2-118, 2-159, 4-12
- DELIMIDENT environment variable 2-85, 2-681, 4-190, 4-192, 5-22, 5-23, 5-24
- delimited identifiers 5-23
- Delimited identifiers 5-23
 - in database server names 5-18
 - multibyte characters 5-24
 - non-ASCII characters 5-24
- DELIMITED keyword
 - in CREATE EXTERNAL TABLE statement 2-109
 - in SELECT statement 2-577
- Delimiter
 - for LOAD input file 2-448
 - specifying with UNLOAD 2-687
- DELIMITER keyword
 - in CREATE EXTERNAL TABLE statement 2-109
 - in LOAD statement 2-443
 - in SELECT statement 2-577
 - in UNLOAD statement 2-684
- DELUXE keyword
 - in CREATE EXTERNAL TABLE statement 2-109
- DES3 (Triple Data Encryption Standard) 4-115
- DESC keyword
 - in CREATE INDEX statement 2-127
 - in SELECT statement 2-568, 2-570
 - order with nulls 2-571
- Descending sequence 2-41, 2-185
- DESCRIBE INPUT statement 2-315
- DESCRIBE statement
 - collection variable with 2-314
 - distinct data type with 2-384
 - opaque data type with 2-384
 - relation to GET DESCRIPTOR 2-383
 - syntax 2-311

DESCRIBE statement (*continued*)
 with SET DESCRIPTOR 2-611

DESCRIPTOR keyword
 in ALLOCATE DESCRIPTOR statement 2-2, 2-3
 in DEALLOCATE DESCRIPTOR statement 2-284
 in DESCRIBE INPUT statement 2-317
 in DESCRIBE statement 2-313
 in EXECUTE statement 2-350, 2-353, 2-355
 in FETCH statement 2-368
 in GET DESCRIPTOR statement 2-381
 in OPEN statement 2-460
 in PUT statement 2-479

destroy() support function 2-149, 2-309, 2-343, 2-681

DETACH keyword, in ALTER FRAGMENT statement 2-16

Detached index 2-23, 2-104, 2-136

Detached statement 2-583

dev/null output destination 2-646

Diagnostics area 2-386

DIAGNOSTICS keyword, in GET DIAGNOSTICS statement 2-386

Diagnostics table
 creating 2-664
 examples 2-677
 filtering mode 2-595
 how to stop 2-676
 relationship to target table 2-668
 relationship to violations table 2-668
 restriction on dropping 2-344
 restriction on moving 2-458
 schema 2-672

Direct cursor 2-78

Directive 5-35

DIRECTIVES configuration parameter 5-36

DIRECTIVES keyword, in SAVE EXTERNAL DIRECTIVES statement 2-523

Dirty Read isolation level 2-143, 2-330, 2-622, 2-633

DIRTY READ keywords
 in SET ENVIRONMENT statement 2-622
 in SET ISOLATION statement 2-631

DIRTY READ setting, of USELASTCOMMITTED configuration parameter 2-634

Disabilities, visual
 reading syntax diagrams C-1

disability C-1

DISABLED keyword
 in ALTER TABLE statement 2-51
 in CREATE INDEX statement 2-138
 in CREATE TABLE statement 2-203, 2-205
 in CREATE TRIGGER statement 2-242, 2-243
 in SET AUTOFREE statement 2-582
 in SET Database Object Mode statement 2-594, 2-597
 in SET DEFERRED_PREPARE statement 2-603

DISCONNECT statement 2-320

DISK keyword
 in CREATE EXTERNAL TABLE statement 2-108

Display labels
 in CREATE VIEW statement 2-276
 in Projection clause 2-541
 in SELECT statement 2-541, 2-575

Distinct data types 4-32
 base types 4-32
 casting 2-101
 casts and DROP TYPE 2-345
 creating with CREATE DISTINCT TYPE 2-100
 DESCRIBE with 2-384
 distributed queries 2-536, 2-537
 dropping 2-345

Distinct data types (*continued*)
 dynamic SQL with 2-384
 GET DESCRIPTOR with 2-384
 in dynamic SQL 2-610
 privileges 2-100, 2-404, 2-501
 restrictions on source type 2-100
 source data type 2-384, 2-610
 Usage privilege 2-404
 with SET DESCRIPTOR 2-610

DISTINCT data types
 distributed queries 4-33

DISTINCT keyword
 in ALTER TABLE statement 2-50, 2-64
 in CREATE DISTINCT TYPE statement 2-100
 in CREATE INDEX statement 2-125, 2-141
 in CREATE TABLE statement 2-197, 2-205
 in CREATE Temporary TABLE statement 2-234
 in Expression segment 4-165, 4-175
 in SELECT statement 2-535
 in subquery 4-17

Distributed DML operations 2-309, 2-438, 2-698

Distributed queries 2-536, 2-537
 cross-database queries 2-536
 cross-server queries 2-537
 terminating ungracefully 6-44

Distributed Relational Database Architecture (DRDA) 6-62

Distribution bins 2-712

Distributions 2-711
 dropping 2-709
 privileges required to create 2-710

DISTRIBUTIONS keyword, in UPDATE STATISTICS statement 2-705, 2-711

divide() operator function 4-51

Division (/) symbol, arithmetic operator 4-40

DML (Data Manipulation Language) statements 1-8

DOCUMENT keyword
 in CREATE FUNCTION statement 2-113
 in CREATE PROCEDURE statement 2-153

Dollar (\$) sign 5-23

Dollar (\$) symbol
 in MONEY values 4-192
 in prepared statements 2-471
 in SQL identifiers 5-21
 prefix for C variables 5-71
 prefix in pathnames 5-71

Dominant table 2-553

DORMANT keyword, in SET CONNECTION statement 2-587

Dot notation 4-61

Dotted decimal format of syntax diagrams C-1

Double hyphen (--) comment indicator 1-3, 5-36

DOUBLE PRECISION data type 4-29

Double quotation marks (") 4-189
 delimiting SQL identifiers 4-190
 literal in a quoted string 4-192
 quoted string delimiter 4-192

DRDA application server 6-62, 6-64

DRDA client-server communication protocol 2-387

DROP ACCESS_METHOD statement 2-322

DROP AGGREGATE statement 2-323

DROP BLOBSPACE command 6-6

DROP CAST statement 2-324

DROP CHUNK command 6-6

DROP CONSTRAINT keywords, in ALTER TABLE statement 2-66

DROP CRCOLS keywords
 in ALTER TABLE statement 2-47

- DROP DATABASE statement 2-325
- DROP DISTRIBUTIONS keywords, in UPDATE STATISTICS statement 2-705
- DROP DUPLICATE statement 2-326
- DROP FUNCTION statement 2-327
- DROP INDEX statement 2-329
- DROP keyword 2-705
 - in ALTER ACCESS_METHOD statement 2-5
 - in ALTER FRAGMENT statement 2-25
 - in ALTER FUNCTION statement 2-28
 - in ALTER PROCEDURE statement 2-32
 - in ALTER ROUTINE statement 2-33
 - in ALTER TABLE statement 2-46, 2-47, 2-55, 2-60
 - in TRUNCATE statement 2-678
- DROP LOG command 6-6
- DROP OPCLASS statement 2-331
- DROP PARTITION keywords, in ALTER FRAGMENT statement 2-25
- DROP PROCEDURE statement 2-332
- DROP ROLE statement 2-334
- DROP ROUTINE statement 2-335
- DROP ROW TYPE statement 2-337
- DROP ROWIDS keywords
 - in ALTER TABLE statement 2-46
- DROP SBSPACE command 6-6
- DROP SECURITY LABEL COMPONENT statement 2-338
- DROP SECURITY LABEL statement 2-338
- DROP SECURITY POLICY statement 2-338
- DROP SECURITY statement 2-338
- DROP SEQUENCE statement 2-340
- DROP SYNONYM statement 2-341
- DROP TABLE statement 2-342
- DROP TEMPDBSPACE command 6-6
- DROP TRIGGER statement 2-344
- DROP TYPE statement 2-345
- DROP VERCOLS keywords
 - in ALTER TABLE statement 2-47
- DROP VIEW statement 2-346
- DROP XADATASOURCE statement 2-347
- DROP XADATASOURCE TYPE statement 2-347
- DS_ADM_POLICY configuration parameter 2-650
- DS_MAX_QUERIES configuration parameter 6-36, 6-39, 6-41
- DS_MAX_SCANS configuration parameter 6-41
- DS_NONPDQ_QUERY_MEM configuration parameter 2-570, 6-41
- DS_TOTAL_MEMORY configuration parameter 6-36, 6-41
- DS_TOTAL_TMPSPACE configuration parameter 2-619
- DUPLICATE keyword
 - in CREATE DUPLICATE statement 2-102
 - in DROP DUPLICATE statement 2-327
- Duplicate table
 - CREATE DUPLICATE 2-102
- Duplicate values
 - in a query 2-535
- Dynamic cursor names 2-288
- Dynamic link library 5-20
- Dynamic log feature 3-45
- Dynamic management statement 1-9
- Dynamic parameters 2-317
- Dynamic routine-name specification 2-360
 - of SPL functions 2-360
 - of SPL procedures 2-367

E

- EACH keyword, in CREATE TRIGGER statement 2-251, 2-257
- East Asian locales 4-119
- EBCDIC keyword
 - in CREATE EXTERNAL TABLE statement 2-109
 - in SELECT statement 2-577
- Element variable 3-25
- ELIF keyword, in IF statement 3-28
- ELSE keyword
 - in CASE statement 3-5
 - in Expression segment 4-66, 4-67
 - in IF statement 3-28
- ENABLED keyword
 - in ALTER TABLE statement 2-51
 - in CREATE INDEX statement 2-138
 - in CREATE TABLE statement 2-203, 2-205
 - in CREATE TRIGGER statement 2-242
 - in SET AUTOFREE statement 2-582
 - in SET Database Object Mode statement 2-594, 2-597
 - in SET DEFERRED_PREPARE statement 2-603
- ENCRYPT VP class 6-37
- ENCRYPT_AES function 4-114
- ENCRYPT_TDES function 4-115
- Encrypted data 2-196, 4-107
- Encryption and decryption functions
 - DECRYPT_BINARY 4-114
 - DECRYPT_CHAR 4-113
 - ENCRYPT_AES 4-114
 - ENCRYPT_TDES 4-115
 - GETHINT 4-116
 - syntax 4-107
- Encryption communication support module 2-611, 4-108
- ENCRYPTION keyword
 - in SET ENCRYPTION PASSWORD statement 2-611
- END CASE keywords, in CASE statement 3-5
- END EXCEPTION keywords, in ON EXCEPTION statement 3-38
- END FOR keywords, in FOR statement 3-19
- END FOREACH keywords, in FOREACH statement 3-23
- END FUNCTION keyword
 - in CREATE FUNCTION statement 2-113
- END IF keywords, in IF statement 3-28
- END keyword 5-75
 - in Expression segment 4-66, 4-67
- END PROCEDURE keywords
 - in CREATE PROCEDURE statement 2-153
- END WHILE keywords, in WHILE statement 3-34, 3-49
- Enterprise Replication 6-48
 - creating shadow columns 2-209, 2-235
- ENVIRONMENT keyword
 - in SET ENVIRONMENT statement 2-616
- Environment variables
 - ANSIOWNER 1-1, 2-518, 5-49
 - CLASSPATH 6-62
 - CLIENT_LOCALE 4-25
 - DB_LOCALE 2-136, 2-282, 2-585, 2-628, 4-25, 5-23
 - DBACCNOIGN 2-75
 - DBANSIWARN 2-100, 2-173, 2-275, 2-288
 - DBBLOBBUF 2-446, 2-686
 - DBCENTURY 2-15, 2-258, 2-444, 4-182
 - DBDATE 2-196, 2-444, 2-684, 4-192
 - DBDELIMITER 2-448, 2-687
 - DBMONEY 2-445, 2-685
 - DBPATH 2-88
 - DBSPACETEMP 2-138, 2-714
 - DBTIME 2-445, 2-685
 - DBUPSPACE 2-712, 2-714
 - DEBUG 6-52
 - DEFAULT_ATTACH 2-136

Environment variables (*continued*)

- DELIMIDENT 2-85, 2-681, 4-190, 5-23, 5-24, 5-32
- GL_DATE 2-196, 2-444, 2-684, 4-192
- GL_DATETIME 2-445, 2-685, 4-183
- IBM_XPS_PARAMS 2-618, 4-75, 4-77
- IFX_DEF_TABLE_LOCKMODE 2-69, 2-226, 2-634
- IFX_DIRECTIVES 5-36
- IFX_DIRTY_WAIT 2-450, 2-679
- IFX_EXTDIRECTIVES 2-523, 2-620, 5-45
- IFX_LONGID 5-22
- IFX_MULTIPREPSTMT 2-575
- IFX_NODBPROC 6-52
- IFX_PAD_VARCHAR 4-24
- IFX_TABLE_LOCKMODE 2-452
- IFX_UPDESC 2-313, 2-316
- INFORMIXCONCSMCFG 2-614
- INFORMIXSERVER 2-88, 4-74
- MAXERRORS 2-110
- NODEFDAC 2-116, 2-156, 2-502
- OPT_GOAL 5-43
- OPTCOMPIND 2-653, 6-51
- PDQPRIORITY 2-616, 2-644, 2-653
- RECORDEND 2-110
- SERVER_LOCALE 4-25
 - setting with SYSTEM statement 3-46
- STMT_CACHE 2-653
- USERTABLENAME 2-44, 2-343, 2-679

Equal sign (=)

- assignment operator 2-695
- relational operator 4-194, 4-195

equal() operator function 2-198, 4-195

Error checking

- continuing after error in SPL routine 3-40
- error status with ON EXCEPTION 3-39
- with SYSTEM 3-45
- with WHENEVER 2-720

ERROR keyword

- in ALTER TABLE statement 2-51
- in CREATE INDEX statement 2-138
- in CREATE TABLE statement 2-203
- in SET CONSTRAINTS statement 2-590
- in SET Database Object Mode statement 2-594
- in SET INDEXES statement 2-631
- in WHENEVER statement 2-718
- synonym for SQLERROR 2-720

ESCAPE keyword

- in Condition segment 4-6, 4-13
- in CREATE EXTERNAL TABLE statement 2-109
- in SELECT statement 2-578
- with LIKE keyword 2-562, 4-15
- with MATCHES keyword 2-563, 4-15

esql compiler 5-36

ESQL/C

- collection cursor with FETCH 2-374
- collection cursor with PUT 2-483
- cursor example 2-295
- deallocating collection-variable memory 2-283
- deallocating row-variable memory 2-285
- error checking for aggregate functions 4-174
- inserting collection variables with 2-437
- inserting row variables 2-438, 2-439
- SQL statements valid only in ESQL/C 4-57
- statements valid only in ESQL/C 4-56

ESQL/C API 4-191

Exception handler 5-77

EXCEPTION keyword

- in GET DIAGNOSTICS statement 2-391

EXCEPTION keyword (*continued*)

- in ON EXCEPTION statement 3-38
- in RAISE EXCEPTION statement 3-41

Excess-65 format 4-28

Exclamation point (!) 4-193

- in smart-large-object filename 4-126

EXCLUSIVE keyword

- in DATABASE statement 2-281
- in LOCK TABLE statement 2-449

Exclusive lock mode 2-125, 2-143, 2-330, 2-449, 2-622, 2-634

Executable file location 5-70

Executable statements 3-39

EXECUTE FUNCTION keywords 3-23, 5-75

- in DECLARE statement 2-286
- in INSERT statement 2-441

EXECUTE FUNCTION statement 2-358

- and triggers 2-259
- how it works 2-357
- preparing 2-359
- syntax 2-356

EXECUTE IMMEDIATE statement 2-360

- restricted statement types 2-361

EXECUTE ON keywords

- in GRANT statement 2-405
- in REVOKE statement 2-501

EXECUTE PROCEDURE keywords 3-23, 5-75

- in DECLARE statement 2-286
- in INSERT statement 2-441

EXECUTE PROCEDURE statement 3-23

- in triggered action 2-259
- syntax 2-364

EXECUTE statement

- INTO clause 2-350
- INTO SQL DESCRIPTOR clause 2-352
- parameterizing a statement 2-353
- returned SQLCODE values 2-353
- syntax 2-348
- USING DESCRIPTOR clause 2-355
- with USING keyword 2-353

EXEMPTION keyword

- in GRANT statement 2-416
- in REVOKE statement 2-511

EXISTS keyword

- beginning a subquery 2-563
- in Condition segment 4-18
- in Condition subquery 4-18

EXIT statement 3-17

EXP function 4-117

EXPLAIN keyword

- in optimizer directives 5-44
- SET EXPLAIN statement 2-623

EXPLAIN_SQL routine 6-53

EXPLAIN_STAT configuration parameter 2-628, 6-41

EXPLICIT keyword

- in CREATE CAST statement 2-95

Exponential function 4-117

Exponential number 4-186

Export support function 2-148, 2-685

exportbinary support function 2-685

Exportbinary support function 2-149

EXPRESS keyword

- in CREATE EXTERNAL TABLE statement 2-109

Expression

- Boolean 4-6, 4-8
- casting 4-58
- constant 4-70
- list of 4-42

- Expression (*continued*)
 - ordering by 2-571
 - smart large objects in 4-65
- EXPRESSION keyword
 - in ALTER FRAGMENT statement 2-20, 2-21
 - in CREATE INDEX statement 2-136
 - in CREATE TABLE statement 2-212
- Expression segment
 - aggregate expressions 4-165
 - cast expressions 4-58
 - column expressions 4-59
 - combined expressions 4-51
 - list of expressions 4-42
 - syntax 4-40
- EXT_DIRECTIVES configuration parameter 2-523, 2-620, 5-45
- EXTDIRECTIVES keyword, in SET ENVIRONMENT
 - statement 2-523, 2-620, 5-45
- EXTEND function 4-128, 4-134
- EXTEND keyword
 - in GRANT statement 2-411
 - in REVOKE statement 2-506
- EXTEND role 2-115, 2-123, 2-161, 2-164, 2-165, 2-167, 2-328, 2-333, 2-334, 2-336
- EXTENT keyword
 - in ALTER TABLE statement 2-62, 2-66
- EXTENT SIZE keywords
 - in ALTER TABLE statement 2-62
 - in CREATE TABLE statement 2-222, 2-224
- Extents
 - counting how many in a table 4-169
 - revising the size 2-619
- External function
 - as operator-class support function 2-152
 - CREATE FUNCTION 2-120
 - dropping 2-328
 - executing 2-356, 2-471
 - limits on return values 5-56
 - non-variant 5-20
 - OUT parameter 4-164
 - registering 2-120
 - strategy functions 2-152
 - variant 5-20
- EXTERNAL keyword
 - in CREATE EXTERNAL TABLE statement 2-105
 - in External Routine Reference segment 5-19
 - in SAVE EXTERNAL DIRECTIVES statement 2-523
 - in SELECT statement 2-574
- External language 2-120
- EXTERNAL NAME keywords
 - in ALTER FUNCTION statement 2-28
 - in ALTER PROCEDURE statement 2-32
 - in ALTER ROUTINE statement 2-33
 - in CREATE FUNCTION statement 2-113
 - in CREATE PROCEDURE statement 2-153
- External optimizer directives 2-522
- External procedure
 - creating body of 2-161
 - dropping 2-333
 - executing 2-471
- External Routine Reference segment 5-19
- External routines
 - as triggered action 2-259
 - CREATE PROCEDURE FROM statement in 2-163
 - creating a function in 2-123
 - defined 2-156
 - dropping 2-336
 - EXTEND role 2-506

- External routines (*continued*)
 - pathname syntax 5-70
 - preparing 2-471
 - referencing 5-20
- External synonym 2-343
- External tables
 - creating 2-104
 - integer data types 2-106
 - NULL values 2-106
 - restrictions in joins and subqueries 2-548
 - restrictions on calculating statistics 2-706
 - restrictions on optimizer directives 5-37
 - with SELECT statement 2-577
- EXTYPEID field 2-381
 - in SET DESCRIPTOR statement 2-606
 - with DESCRIBE INPUT statement 2-318
 - with DESCRIBE statement 2-313
- EXTYPELENGTH field 2-381
 - in SET DESCRIPTOR statement 2-606
 - with DESCRIBE INPUT statement 2-318
 - with DESCRIBE statement 2-313
- EXTYPENAME field 2-381
 - in SET DESCRIPTOR statement 2-606
 - with DESCRIBE INPUT statement 2-318
 - with DESCRIBE statement 2-313
- EXTYPEOWNERLENGTH field 2-381
 - in SET DESCRIPTOR statement 2-606
 - with DESCRIBE INPUT statement 2-318
 - with DESCRIBE statement 2-313
- EXTYPEOWNERNAME field 2-381
 - in SET DESCRIPTOR statement 2-606
 - with DESCRIBE INPUT statement 2-318
 - with DESCRIBE statement 2-313

F

- FETCH statement 2-367
 - relation to GET DESCRIPTOR 2-382
- Field projection 2-697, 4-61
- Field qualifier 4-38, 4-176, 4-183
- field qualifiers 4-38
- FILE TO keywords
 - in SET DEBUG FILE statement 2-600
 - in SET EXPLAIN statement 2-624
 - in SET PLOAD FILE statement 2-645, 2-646
- Files
 - for LOAD input 2-684
 - in FIXED format 2-107
 - saving output from UNLOAD 2-684
 - saving output from UPDATE STATISTICS 2-712
 - saving query plans in sqexplain.out 2-626
 - sending output with the OUTPUT statement 2-467
- FILETOBLOB function 4-122, 4-123
- FILETOCLOB function 4-122, 4-123
- FILLFACTOR configuration parameter 2-134
- FILLFACTOR keyword
 - in CREATE INDEX statement 2-133
- FILTERING keyword
 - in ALTER TABLE statement 2-51
 - in CREATE INDEX statement 2-138
 - in CREATE TABLE statement 2-203, 2-205
 - in SET Database Object Mode statement 2-594
 - with diagnostics tables 2-675
- FINAL keyword
 - in CREATE AGGREGATE statement 2-92
- FIRST keyword
 - in FETCH statement 2-368

- FIRST keyword (*continued*)
 - in SELECT statement 2-533, 2-534
 - invalid in INSERT 2-441
- FIRST_ROWS keyword
 - in optimizer directives 5-42
 - in SET OPTIMIZATION statement 2-641
- Fixed format
 - files 2-106
- FIXED keyword
 - in CREATE EXTERNAL TABLE statement 2-109
- Fixed-length opaque data type 2-146
- Fixed-point numbers 4-186
- FLOAT data type 4-29
 - literal values 4-186
 - systems not supporting 2-282
- Floating-point numbers 4-186
- FLOOR function 4-89, 4-90
- FLUSH statement 2-377
- Flushing an insert buffer 2-484
- for DATETIME 4-38, 4-181
- FOR EACH ROW keywords, in CREATE TRIGGER statement 2-251, 2-257
- for INTERVAL 4-176, 4-183
- FOR keyword 2-705, 3-23
 - in CONTINUE statement 3-7
 - in CREATE FUNCTION statement 2-113
 - in CREATE OPCLASS statement 2-149
 - in CREATE PROCEDURE statement 2-153
 - in CREATE SYNONYM statement 2-187
 - in CREATE TRIGGER statement 2-251, 2-257
 - in DECLARE statement 2-286
 - in EXIT statement 3-17
 - in GRANT statement 2-416, 2-418
 - in INFO statement 2-428
 - in REVOKE statement 2-511, 2-513
 - in SAVE EXTERNAL DIRECTIVES statement 2-523
 - in SELECT statement 2-572, 2-573
 - in SET AUTOFREE statement 2-582
 - in SET CONSTRAINTS statement 2-590
 - in SET Database Object Mode statement 2-593
 - in SET INDEXES statement 2-631
 - in SET TRIGGERS statement 2-663
 - in START VIOLATIONS TABLE statement 2-664
 - in STOP VIOLATIONS TABLE statement 2-676
 - in UPDATE STATISTICS statement 2-714
- FOR READ ONLY keywords
 - in DECLARE statement 2-286
 - in SELECT statement 2-573
- FOR statement 3-19
- FOR TABLE keywords in UPDATE STATISTICS statement 2-705
- FOR UPDATE keywords 2-528
 - in DECLARE statement 2-286
 - in SELECT statement 2-572
 - with column list 2-291
- FOREACH keyword
 - in CONTINUE statement 3-7
 - in EXIT statement 3-17
- FOREACH statement
 - syntax 3-22
- Foreign key
 - dropping 2-66
 - establishing 2-52, 2-200
 - examples 2-53, 2-208
 - multiple columns 2-206
- Foreign key constraint 2-207, 2-508
- FOREIGN KEY keywords
 - in ALTER TABLE statement 2-64
 - in CREATE TABLE statement 2-205, 2-227
- FORMAT keyword
 - in CREATE EXTERNAL TABLE statement 2-109
 - in SELECT statement 2-574, 2-578
- FORMAT_UNITS function 4-158
- FRACTION keyword 4-176, 4-181, 4-183
 - in DATETIME Field Qualifier segment 4-38
- FRAGMENT BY keywords
 - in ALTER FRAGMENT statement 2-20, 2-21
 - in CREATE INDEX statement 2-136
 - in CREATE TABLE statement 2-212
- FRAGMENT keyword
 - in ALTER FRAGMENT statement 2-20
 - in GRANT FRAGMENT statement 2-423
 - in REVOKE FRAGMENT statement 2-517
- Fragment-level privilege
 - granting 2-423
 - revoking 2-517
- Fragmentation
 - adding a fragment 2-23
 - adding rowids 2-46
 - altering fragments 2-7
 - arbitrary rule 2-215
 - built-in hash distribution scheme 2-216
 - by expression 2-15
 - combining tables 2-9
 - Dataskip feature 2-598
 - defining a new strategy 2-136
 - detaching a table fragment 2-16
 - dropping an existing fragment 2-25
 - dropping rowids 2-46
 - hybrid 2-15
 - insufficient log space or disk space 2-9
 - list of dbspaces 2-598
 - modifying an existing fragment expression 2-26
 - nonremainder fragment 2-26
 - number of rows in fragment 2-9
 - of indexes 2-136
 - of tables 2-212
 - of temporary tables 2-237
 - reading data from local fragments 2-548
 - reinitializing strategy 2-21
 - remainder 2-25
 - reverting to nonfragmented 2-19
 - round-robin 2-14
 - rowid 2-19
 - rowid columns with 2-214
 - strategy
 - by expression 2-214, 2-599
 - by round-robin 2-214, 2-599
 - range rule 2-215
 - TEXT and BYTE data types 2-13
- Fragmentation strategy, modifying 2-18
- FRAGMENTS keyword, in INFO statement 2-428
- FREE statement 2-379
- FROM keyword 2-494
 - in CREATE INDEX statement 2-142
 - in DELETE statement 2-303
 - in LOAD statement 2-443
 - in PREPARE statement 2-468
 - in PUT statement 2-479
 - in REVOKE FRAGMENT statement 2-517
 - in REVOKE statement 2-509, 2-511, 2-513, 2-515
 - in SELECT statement 2-544
 - in TRIM expressions 4-147

FROM keyword (*continued*)
 in UPDATE statement 2-699

FULL keyword
 in optimizer directives 5-37
 in SELECT statement 2-557

Full outer joins 2-553

Function
 companion 2-357
 negator 2-357
 selectivity 5-65

Function cursor
 defined 3-24
 opening 2-461
 reopening 2-461

Function expressions 4-87

FUNCTION keyword
 in ALTER PROCEDURE statement 2-28
 in DECLARE statement 2-286
 in DROP FUNCTION statement 2-327
 in EXECUTE FUNCTION statement 2-356
 in GRANT statement 2-405
 in REVOKE statement 2-501
 in SELECT statement 2-551
 in UPDATE STATISTICS statement 2-714

Functional index 2-127, 2-128, 2-129, 2-143, 2-614

Functions
 altering with ALTER FUNCTION 2-28
 casting 2-97
 collection manipulation 5-64
 creating indirectly from a stored file 2-123
 creating with CREATE FUNCTION 2-113
 creating with CREATE FUNCTION FROM 2-122
 cursor 3-24
 distributed transactions 3-43, 5-57
 dropping with DROP FUNCTION 2-327
 dropping with DROP ROUTINE 2-335
 modifying
 path to executable file 2-29
 modifying routine modifiers 2-29
 noncursor 3-24
 nonvariant 5-20
 of SQL administration API 6-3
 smart large object 4-122
 specific name 5-73
 system catalog tables 2-119
 thread-safe 5-64
 trigger 2-116
 user-defined
 defined 2-155
 variant 5-20

Functions,
 security label support 4-119

Functions, SQL 4-128

ABS 4-90

ACOS 4-140

ADD_MONTHS 4-129

ADMIN 6-2

ASCII 4-146

ASIN 4-141

ATAN 4-141

ATAN2 4-141

AVG 4-168

BITAND 4-52

BITANDNOT 4-52

BITNOT 4-52

BITOR 4-52

BITXOR 4-52

Functions, SQL (*continued*)

CARDINALITY 4-98

CASE 4-66

CAST 4-58

CEIL 4-90

CHAR_LENGTH 4-119

CHARACTER_LENGTH 4-119

CONCAT 4-142

COS 4-140

COUNT 4-168

CURRENT 4-75

CURRENT_ROLE 4-73

CURRVAL 4-80

DATE 4-130

DAY 4-130

DBINFO 4-99

DECODE 4-69

DECRYPT_BINARY 4-114

DECRYPT_CHAR 4-113

DEFAULT_ROLE 4-73

ENCRYPT_AES 4-114

ENCRYPT_TDES 4-115

Encryption and decryption 4-107

EXP 4-117

Exponential 4-116

EXTEND 4-134

FILETOBLOB 4-123

FILETOCLOB 4-123

FLOOR 4-90

FORMAT_UNITS 4-158

GETHINT 4-116

HEX 4-117

IFX_REPLACE_MODULE 6-54

INITCAP 4-156

LAST_DAY 4-132

LENGTH 4-119

LOCOPY 4-126

LOG10 4-117

Logarithmic 4-116

LOGN 4-117

LOTOFILE 4-125

LOWER 4-156

LPAD 4-154

LTRIM 4-149

MAX 4-172

MDY 4-135

MIN 4-172

MOD 4-91

MONTH 4-131

MONTHS_BETWEEN 4-131

NEXT_DAY 4-133

NEXTVAL 4-79

NULLIF 4-68

NVL 4-68

OCTET_LENGTH 4-119

POW 4-91

POWER 4-91

RANGE 4-172

REPLACE 4-153

ROOT 4-91

ROUND 4-91

RPAD 4-155

RTRIM 4-149

SECLABEL_BY_COMP 4-120

SECLABEL_BY_NAME 4-120

SECLABEL_TO_CHAR 4-121

SIN 4-140

Functions, SQL (*continued*)

- SQLCODE 4-98
- SQRT 4-91
- STDEV 4-173
- SUBSTR 4-152
- SUBSTRING 4-150
- SUM 4-172
- SYSDATE 4-76
- TAN 4-140
- TASK 6-2
- TO_CHAR 4-135
- TO_DATE 4-138
- TO_NUMBER 4-139
- TODAY 4-75
- TRIM 4-147
- TRUNC 4-95
- UPPER 4-156
- USER 4-72
- VARIANCE 4-173
- WEEKDAY 4-131
- YEAR 4-131

Fuzzy index 2-152

G

G keyword in SQL Admin API functions 6-2

Generalized-key index

- affected by dropping a column 2-57
- affected by modifying a column 2-61
- defined 2-141
- no renamed table 2-493
- no table expressions 2-547
- SELECT clause 2-141
- WHERE clause 2-143

Generic B-tree index 2-132

Generic CASE expressions 4-66

GET DESCRIPTOR statement

- syntax 2-381
- use with FETCH statement 2-372

GET DIAGNOSTICS statement

- SQLSTATE codes 2-387
- syntax 2-386

GET keyword

- in GET DESCRIPTOR statement 2-381
- in GET DIAGNOSTICS statement 2-386

GET_TZ option of DBINFO 4-105

GETHINT function 4-116

GK 2-143

GK index 2-104, 2-125

GK INDEX keywords, in CREATE INDEX statement 2-124

GL_DATE environment variable 2-196, 2-444, 2-684, 4-192

GL_DATETIME environment variable 2-445, 2-685, 4-183

Global environment 3-10

GLOBAL keyword, in DEFINE statement 3-8

Global variables 3-10

GO TO keywords, in WHENEVER statement 2-718

GOTO keyword, in WHENEVER statement 2-721

GOTO statement 3-27

GRANT FRAGMENT statement 2-423

GRANT keyword 2-396

- in GRANT FRAGMENT statement 2-423

GRANT statement 2-395

Greater than (>) symbol 4-193

greaterthan() operator function 4-195

greaterthanorequal() operator function 4-195

GROUP BY keywords, in SELECT statement 2-566

H

Handle value 4-64, 4-123

HANDLESNULLS keyword

- in CREATE AGGREGATE statement 2-92
- in Routine Modifier segment 5-60

Hash join 2-621, 2-628, 5-41

HASH keyword

- in ALTER FRAGMENT statement 2-20
- in CREATE INDEX statement 2-136
- in CREATE TABLE statement 2-212

HAVING keyword

- in SELECT statement 2-567

HEADINGS keyword, in OUTPUT statement 2-467

HEX function 4-64, 4-117

HEX keyword

- in CREATE EXTERNAL TABLE statement 2-105

Hexadecimal digits 2-687

Hexadecimal dump format 2-685

Hexadecimal smart-large-object identifier 2-686, 4-125

High availability 6-44

HIGH INTEG keywords

- in ALTER TABLE statement 2-62
- in CREATE TABLE statement 2-222

HIGH keyword 2-705

- in SET OPTIMIZATION statement 2-641

- in SET PDQPRIORITY statement 2-643

High-Performance Loader 2-148, 2-573

HINT keyword

- in SET ENCRYPTION PASSWORD statement 2-611

Hold cursor

- defined 2-293
- insert cursor with hold 2-296
- update cursor with hold 2-292

HOLD keyword 3-23

- in DECLARE statement 2-286

hosts.equiv file 2-88

Hot list 6-32

HOUR keyword 4-38, 4-176, 4-181, 4-183

HYBRID keyword

- in ALTER FRAGMENT statement 2-20
- in CREATE INDEX statement 2-136
- in CREATE TABLE statement 2-212, 2-218

Hyphen symbol (-)

- DATETIME separator 4-181
- INTERVAL separator 4-184

I

IBM

- format binary
- in external tables 2-106

IBM Data Studio 6-53

IBM_XPS_PARAMS environment variable 2-618

IDATA field 2-381

- in SET DESCRIPTOR statement 2-606
- with X/Open programs 2-383

Identifier 5-23

- column names 5-29
- connection name 2-86
- cursor name 5-30, 5-33
- defined 5-21
- delimited identifiers 5-24
- multibyte characters 5-23
- non-ASCII characters 5-23
- non-unique 5-26
- routines 5-30

Identifier (*continued*)

- storage objects 5-24
- syntax 5-21
- table names 5-28, 5-29, 5-30
- uppercase characters 5-22
- using keywords 5-23
- using keywords as column names 5-27
- variable name 5-32

IDSLBACREADARRAY keyword

- in GRANT statement 2-416
- in REVOKE statement 2-511

IDSLBACREADSET keyword

- in GRANT statement 2-416
- in REVOKE statement 2-511

IDSLBACREADTREE keyword

- in GRANT statement 2-416
- in REVOKE statement 2-511

IDSLBACRULES

- granting exemptions 2-416, 2-511
- in distributed queries 2-538

IDSLBACRULES keywords

- in CREATE SECURITY POLICY statement 2-182

IDSLBACWRITEARRAY keyword

- in GRANT statement 2-416
- in REVOKE statement 2-511

IDSLBACWRITESSET keyword

- in GRANT statement 2-416
- in REVOKE statement 2-511

IDSLBACWRITETREE keyword

- in GRANT statement 2-416
- in REVOKE statement 2-511

IDSSECURITYLABEL data type

- in distributed queries 2-416, 2-511, 2-538

IF statement 3-28

IFX_ALLOW_NEWLINE function

- effect on quoted strings 4-191
- syntax 4-160

IFX_AUTO_REPREPARE keyword, in SET ENVIRONMENT statement 2-620

IFX_AUTO_REPREPARE session environment variable 2-466, 2-478, 2-714

IFX_DEF_TABLE_LOCKMODE environment variable 2-69, 2-226, 2-634

IFX_DIRECTIVES environment variable 5-36

IFX_DIRTY_WAIT environment variable 2-450, 2-679

IFX_EXTDIRECTIVES environment variable 2-523, 2-620, 5-45

IFX_EXTEND_ROLE configuration parameter 2-29, 2-33, 2-35, 2-115, 2-123, 2-161, 2-164, 2-167, 2-411, 2-507, 5-71, 6-41, 6-54

ifx_insert_checksum shadow column 2-47

IFX_LO_SPEC data type 2-536, 4-22

IFX_LO_STAT data type 2-536, 4-22

IFX_LONGID environment variable 5-22

IFX_MULTIPREPSTMT environment variable 2-575

IFX_NODBPROC environment variable 6-52

IFX_PAD_VARCHAR environment variables 4-24

IFX_REPLACE_MODULE function 6-54

ifx_row_version shadow column 2-47

IFX_TABLE_LOCKMODE environment variable 2-452

IFX_UPDDESC environment variable 2-313, 2-316

ILENGTH field 2-381

- in SET DESCRIPTOR statement 2-606
- with X/Open programs 2-383

Ill-behaved C UDR 5-62

Imaginary numbers 4-29

IMMEDIATE keyword

- in EXECUTE IMMEDIATE statement 2-360
- in SET Transaction Mode statement 2-661

IMMUTABLE keyword

- in SET ALL_MUTABLES statement 2-581
- in SET DEFAULT TABLE_SPACE statement 2-601
- in SET DEFAULT TABLE_TYPE statement 2-602
- in SET PDQPRIORITY statement 2-644

IMPEX data type 2-537, 4-23

IMPEXBIN data type 2-537, 4-23

Implicit connection 2-88

Implicit cursor 3-23

Implicit inner join 2-557

IMPLICIT keyword

- in CREATE CAST statement 2-95

Implicit transactions 2-652

IMPLICIT_PDQ keyword

- in SET ENVIRONMENT statement 2-618

Import support function 2-148, 2-445

Importbinary support function 2-148, 2-445

in != relational operator 4-193

- in Collection Subquery segment 4-3
- in Collection-Derived Table segment 5-4
- in Collection-Subquery segment 4-3
- in DATETIME data type 4-38
- in DATETIME field qualifier 4-181
- in DATETIME Field Qualifier 4-38, 4-181
- in DATETIME Field Qualifier segment 4-38
- in EXECUTE FUNCTION statement 2-358
- in FOR statement 3-19
- in FOREACH 3-23
- in FOREACH statement 3-23
- in GET DESCRIPTOR statement 2-381
- in GRANT FRAGMENT statement 2-424
- in GRANT statement 2-396, 2-397
- in IF statement 3-30
- in INSERT 2-430
- in INSERT statement 2-430
- in INTERVAL data type 4-183
- in INTERVAL field qualifier 4-176
- in INTERVAL Field Qualifier 4-176, 4-183

IN keyword 3-19

- as a condition 4-17
- in ALTER FRAGMENT statement 2-20, 2-21
- in ALTER TABLE statement 2-62
- in Condition segment 4-10, 4-17
- in CREATE DATABASE statement 2-97
- in CREATE DUPLICATE statement 2-102
- in CREATE FUNCTION statement 2-113
- in CREATE INDEX statement 2-134, 2-136
- in CREATE PROCEDURE statement 2-153
- in CREATE TABLE statement 2-210, 2-212, 2-218, 2-219, 2-222
- in CREATE Temporary TABLE statement 2-175, 2-236
- in Data Type segment 4-29
- in LOCK TABLE statement 2-449
- in ON EXCEPTION statement 3-38
- in SELECT statement 2-561

- in Literal Collection 4-178
- in Literal DATETIME 4-181
- in Literal Row segment 4-186
- in relational operators 4-193
- in Return Clause segment 5-55
- in Return Clause Segment 5-55
- in REVOKE FRAGMENT statement 2-518
- in REVOKE statement 2-494, 2-496
- in SELECT 2-528
- in SELECT statement 2-528, 2-531
- in SET LOG statement 2-640
- in SQL identifiers 5-23

- in Statement Block segment 5-75
- in UPDATE statement 2-689
- in UPDATE STATISTICS statement 2-705
- INACTIVE keyword, in SAVE EXTERNAL DIRECTIVES statement 2-523
- INCREMENT keyword
 - in ALTER SEQUENCE statement 2-41
 - in CREATE SEQUENCE statement 2-185
- Index
 - access method 2-149
 - altering table fragmentation 2-13
 - attached 2-27, 2-136
 - B-tree 2-128, 2-136
 - bidirectional traversal 2-130
 - clustered fragments 2-126
 - compacted 2-133
 - composite 2-129
 - converting during upgrade 2-705, 2-717
 - creating 2-124
 - delete flag 2-708
 - detached 2-23, 2-136
 - disabled 2-140, 2-595
 - displaying information for 2-429
 - Dropping with DROP INDEX 2-329
 - filtering to violations table 2-595
 - fragmented 2-8, 2-136, 2-208
 - functional 2-128, 2-136, 2-143
 - fuzzy 2-152
 - generalized-key 2-141
 - internal 2-208
 - key filter 2-628
 - maximum key size 2-127, 2-129
 - memory resident 2-630
 - multilingual index 2-586
 - nonfragmented 2-19
 - on ORDER BY columns 2-572
 - on temporary tables 2-576
 - online 2-144
 - privilege 2-8
 - provide for expansion 2-133
 - R-Tree 2-128, 2-136
 - renaming 2-489
 - residency status 2-630, 2-646
 - ROOT argument 4-91
 - self-join keys 2-628
 - self-join path 5-37
 - shared 2-51, 2-53, 2-65
 - side-effect 2-152
 - system-generated 2-208, 2-456, 2-457, 5-37
 - unique 2-139
 - restrictions 2-22
 - unique keys 5-52
 - virtual 2-143, 2-331
- INDEX keyword
 - in ALTER FRAGMENT statement 2-7
 - in ALTER INDEX statement 2-30
 - in CREATE INDEX statement 2-124
 - in DROP INDEX statement 2-330
 - in GRANT statement 2-401
 - in optimizer directives 5-37
 - in RENAME INDEX statement 2-489
 - in REVOKE statement 2-497
 - in SET INDEX statement 2-630
 - in SET Residency statement 2-646
- INDEX optimizer directive 5-38
- Index privilege 2-401, 2-497
- INDEX_SJ keyword
 - in optimizer directives 5-37
- INDEXES keyword
 - in INFO statement 2-428
 - in SET Database Object Mode statement 2-593
 - in SET INDEXES statement 2-631
- INDEXKEYARRAY data type 2-536, 4-22
- INDICATOR field 2-381
 - in SET DESCRIPTOR statement 2-606
- INDICATOR keyword 2-358
 - in EXECUTE statement 2-350, 2-353
 - in FETCH statement 2-368
 - in PUT statement 2-479
 - in SELECT statement 2-542
- Indicator variable
 - in expression 4-174
- Indirect typing 3-16
- INFO statement 2-428
- Informix 4GL 4-191
- INFORMIX keyword
 - in CREATE EXTERNAL TABLE statement 2-109
 - in External Routine Reference segment 5-19
 - in SELECT statement 2-577
- INFORMIX parameter style 5-20
- informix user name 2-325, 2-399, 5-17
- Informix-Admin group 2-414
- INFORMIX.JVPCONTROL function 6-56
- INFORMIXCONCSMCFG environment variable 2-614
- INFORMIXSERVER environment variable 2-86, 2-88, 4-74
- Inheritance hierarchy
 - dropping tables 2-343
 - named ROW types 2-169, 2-337
- INIT keyword
 - in ALTER FRAGMENT statement 2-18
 - in CREATE AGGREGATE statement 2-92
- INITCAP function 4-156
- Initial-cap characters, converting to 4-156
- Inner joins 2-553
- INNER keyword in SELECT statement 2-557
- INOUT parameter 4-162
- INOUT parameters 5-20, 5-70
- Input support function 2-148
- Insert buffer 2-484
 - counting inserted rows 2-378, 2-485
 - filling with constant values 2-480
 - inserting rows
 - with a cursor 2-433
 - storing rows with PUT 2-479
 - triggering flushing 2-484
- Insert cursor 2-293
 - benefits 2-293
 - closing 2-78
 - declaring 2-288
 - in INSERT 2-433
 - in PUT 2-480
 - opening 2-462
 - reopening 2-463
 - result of CLOSE in SQLCA 2-78
 - with hold 2-296
- INSERT INTO keywords 2-430
 - in LOAD 2-448
- INSERT keyword 2-424, 2-518
 - in CREATE TRIGGER statement 2-246
 - in DECLARE statement 2-286
 - in GRANT statement 2-400
 - in LOAD statement 2-443
 - in MERGE statement 2-452

INSERT keyword (*continued*)
 in REVOKE statement 2-497
 Insert privilege 2-400, 2-497
 INSERT statements 2-430
 and triggers 2-259
 AT clause 2-432
 collection-column values 2-437
 collection-derived table, with 2-442
 effect of transactions 2-434
 ESQL/C 2-437, 2-438, 2-439
 filling insert buffer with PUT 2-479
 flex inserts 2-620
 in dynamic SQL 2-443
 insert cursor compared with 2-293
 insert triggers 2-244
 inserting
 rows through a view 2-432
 rows with a cursor 2-433
 into collection cursor 2-483
 nulls 2-439
 OUT parameter and SLVs 4-165
 protected tables 2-440
 row type field values 2-437
 row variables 2-442
 SERIAL and SERIAL8 columns 2-436
 smart large objects with 4-64
 specifying values to insert 2-434
 syntax 2-430
 using functions 2-439
 VALUES clause, expressions with 2-439
 with DECLARE statement 2-286
 with insert cursor 2-292
 with SELECT statement 2-441
 Insert trigger 2-244, 2-271
 INSERTING operator 2-118, 2-159, 4-12
 install_jar() procedure 2-120, 2-161
 INSTEAD OF keywords, in CREATE TRIGGER
 statement 2-239
 INSTEAD OF trigger 2-270
 INT8 data type 4-27
 INTEG keyword
 in ALTER TABLE statement 2-62
 in CREATE TABLE statement 2-222
 INTEGER data type 4-27
 in external tables 2-106
 literal values 4-185
 INTERNAL keyword, in Routine Modifier segment 5-60
 INTERNLENGTH keyword, in CREATE OPAQUE TYPE
 statement 2-145
 INTERVAL data type 4-183
 as quoted string 4-192
 field qualifier 4-176
 in expression 4-78
 in INSERT 4-192
 literal 4-184
 loading 2-445
 precision 4-176
 syntax 4-32
 INTERVAL keyword, in Literal INTERVAL 4-183
 INTO clause 2-358
 INTO DESCRIPTOR keywords, in EXECUTE 2-352
 INTO EXTERNAL keywords
 in SELECT statement 2-577
 INTO keyword 2-358, 2-430, 3-23
 in DESCRIBE INPUT statement 2-315
 in DESCRIBE statement 2-311
 in EXECUTE FUNCTION statement 2-356
 INTO keyword (*continued*)
 in EXECUTE PROCEDURE statement 2-364
 in EXECUTE statement 2-350
 in FETCH statement 2-368
 in LOAD statement 2-443
 in MERGE statement 2-452
 in SELECT statement 2-541
 INTO SCRATCH keywords, in SELECT statement 2-578
 INTO SQL DESCRIPTOR keywords, in EXECUTE
 statement 2-352
 INTO TEMP clause
 in SELECT statement 2-575
 invalid in INSERT 2-441
 with UNION operator 2-579
 IS keyword
 in Condition segment 4-6
 in WHERE clause 2-562
 IS NOT NULL keywords
 Condition segment 4-12
 in WHERE clause of a query 2-562
 IS NULL keywords
 Condition segment 4-12
 in WHERE clause of a query 2-562
 ISAM error code 3-38, 3-41, 3-44
 ISOLATION keyword
 in SET ISOLATION statement 2-632
 in SET TRANSACTION statement 2-656
 Isolation level
 defined 2-633, 2-658
 with FETCH statement 2-373
 Item descriptor 2-3
 ITEM keyword 4-3
 ITER keyword
 in CREATE AGGREGATE 2-92
 Iterator functions 2-551, 2-628, 3-44, 5-63
 ITERATOR keyword, in Routine Modifier segment 5-60
 ITYPE field 2-381
 in SET DESCRIPTOR statement 2-606
 with X/Open programs 2-383

J

Jagged rows 2-227
 Jar files
 installing in the database 6-59
 name of a jar ID 5-34
 read permissions 6-59
 renaming 2-489
 Java class
 installing 6-58
 jar file where defined 5-34
 mapping to a UDT 6-61
 package where defined 5-72
 removing a jar file 6-59
 replacing a jar file 6-59
 Java Common Client (JCC) 6-62, 6-64
 JAVA keyword, in External Routine Reference segment 5-19
 Java UDRs 6-58, 6-59, 6-60
 CLASS routine modifier 5-60
 creating functions 2-120
 creating procedures 2-161
 getting JVP memory information 6-57
 Getting JVP thread information 6-57
 installing a Jar file 6-58
 Java signature 5-73
 jvpcontrol function 6-56
 shared-object file 5-72

Java UDRs (*continued*)
 sqlj.alter_java_path procedure 6-60
 sqlj.replace_jar procedure 6-59
 sqlj.setUDTextName procedure 6-61
 sqlj.unsetUDTextName procedure 6-62
 static method 5-72
 unmapping a user-defined type 6-62

Java Virtual Processor Class
 CLASS modifier 5-62
 getting memory information 6-57
 getting thread information 6-57

Java Virtual-Table Interface 2-91

JDBC API 4-191

JDBC connection 5-64

JDBC Driver built-in function 6-56

Join
 condition 2-552
 hash join 2-621, 2-628, 5-41
 in Condition segment 2-564
 in DELETE statement 2-308
 in UPDATE statement 2-699
 index self-join path 5-37
 multiple-table join 2-565
 nested loop join 5-41
 nested-loop join 2-621, 2-628
 outer 2-560
 outer, Informix extension syntax 2-566
 self-join 2-565
 sort-merge join 2-621
 transitively joined on key 2-142

Join filter 2-557

JOIN keyword
 in SELECT statement 2-553, 2-555

Join on key query 2-142

Join-method directive 2-557, 5-40

Join-order directive 5-40

JVP VP class 6-37

JVPCONFIGURATION configuration parameter 6-60

jpvccontrol function 6-56

K

K keyword in SQL Admin API functions 6-2

KEEP ACCESS TIME keywords
 in ALTER TABLE statement 2-62
 in CREATE TABLE statement 2-222

KEEP keyword
 in ALTER TABLE statement 2-62

keep option of esqlc 5-36

Kernel asynchronous I/O (KAIO) 6-37

KEY keyword
 CREATE TABLE statement 2-227
 in CREATE TABLE statement 2-206
 in CREATE Temporary TABLE statement 2-235

Key management for encrypted data 4-109

Key-only index scan 2-141

Keywords
 as identifiers 5-23
 list for IDS A-1
 list for XPS B-1

Killing a user session 6-44

L

Label
 data security label 2-175, 2-338

Label (*continued*)
 loop label 3-2
 statement label 3-2, 3-27
 user security label 2-175, 2-338, 2-418, 2-513

LABEL COMPONENT keywords
 in DROP SECURITY statement 2-338
 in RENAME SECURITY statement 2-490

LABEL keyword
 in ALTER SECURITY LABEL COMPONENT statement 2-36
 in CREATE SECURITY LABEL COMPONENT statement 2-177
 in CREATE SECURITY LABEL statement 2-175
 in CREATE SECURITY POLICY statement 2-182
 in DROP SECURITY statement 2-338
 in GRANT statement 2-418
 in RENAME SECURITY statement 2-490
 in REVOKE statement 2-513

Language
 external 2-120
 privileges on 2-406, 2-503

LANGUAGE keyword
 in External Routine Reference segment 5-19
 in GRANT statement 2-406
 in REVOKE statement 2-503

Large objects 4-29
 constraints 2-198, 2-206
 pointer structure 2-149

LAST COMMITTED keywords, in SET ISOLATION statement 2-631

LAST keyword
 in FETCH statement 2-368
 in SET ISOLATION statement 2-634

LAST_DAY function 4-128, 4-132

LEADING keyword, in TRIM expressions 4-147

Leap second 4-106

LEFT keyword in SELECT statement 2-557

Left outer joins 2-553

LENGTH field 2-381
 in SET DESCRIPTOR statement 2-606
 with DATETIME and INTERVAL types 2-609
 with DECIMAL and MONEY types 2-609
 with DESCRIBE INPUT statement 2-318
 with DESCRIBE statement 2-313

Length function 4-118

LENGTH function 2-539, 4-118, 4-119

Less than (<) symbol 4-193

lessthan() operator function 4-195

lessthanorequal() operator function 4-195

LET statement 3-32

Lettercase conversion 4-156

LEVEL keyword
 in SET SCHEDULE LEVEL statement 2-649
 in SET TRANSACTION statement 2-657

Level-0 backup 2-193, 2-574

Light appends 2-70, 2-193, 2-242

Light scan 4-25

LIKE keyword
 in Condition segment 4-6, 4-13
 in DEFINE statement 3-8
 in Routine Parameter List segment 5-68
 in SELECT statement 2-562
 wildcard characters 2-562

like() operator function 4-14

LIO VP class 6-37

LIST data type
 columns, generating values for 4-83

LIST data type (*continued*)
 defined 4-84
 deleting elements from 2-309
 unloading 2-684
 updating elements 2-703
 LIST keyword 4-178
 in DEFINE statement 3-13
 in Expression segment 4-83
 LISTEN_TIMEOUT configuration parameter 6-41
 LISTING keyword
 in CREATE FUNCTION statement 2-113
 in CREATE PROCEDURE statement 2-153
 Literal
 BOOLEAN 2-196
 DATE 2-196
 DATETIME 4-181
 in ALTER TABLE statement 2-49
 in INSERT statement 2-434
 with IN keyword 2-561
 INTERVAL 4-183
 in expression 4-78
 in INSERT statement 2-434
 nested row 4-189
 number 4-72
 Number 4-185
 in INSERT 2-434
 with IN keyword 4-11
 Literal collection
 nested 4-180
 syntax 4-178
 Literal number, exponential notation 4-186
 Literal Row segment 4-186
 literal values 4-181
 Literal values, specifying as default values 2-196
 Little-endian format 2-106
 LOAD statement 2-443
 LOCAL keyword
 in SELECT statement 2-548
 local option to esqlc command 2-288, 2-349
 Local variable 3-12
 Localized collation order 2-570, 2-585, 4-14, 4-196
 LOCK keyword
 in ALTER INDEX statement 2-30
 in ALTER TABLE statement 2-68
 in CREATE Temporary TABLE statement 2-175
 in SET LOCK MODE statement 2-637
 LOCK MODE keywords
 in ALTER INDEX statement 2-30
 in ALTER TABLE statement 2-68
 in CREATE INDEX statement 2-141
 in CREATE TABLE statement 2-225
 Lock table overflow 2-292
 LOCK TABLE statement
 syntax 2-449
 use in transactions 2-74
 Locking
 blobspaces 2-17
 during
 inserts 2-434
 updates 2-292, 2-692
 exclusive locks 2-292, 2-449, 2-622, 2-634
 granularity 2-68, 2-452
 in transactions 2-74
 overriding row-level 2-451
 promotable lock 2-292
 releasing with COMMIT WORK statement 2-82, 2-292
 releasing with ROLLBACK WORK statement 2-520
 Locking (*continued*)
 shared locks 2-449
 types of locks 2-225
 update cursors effect on 2-292
 update locks 2-637, 2-692
 waiting period 2-638
 when creating a referential constraint 2-54, 2-201
 with
 SET ISOLATION statement 2-632
 SET LOCK MODE statement 2-637
 UNLOCK TABLE statement 2-688
 with FETCH statement 2-373
 with SET TRANSACTION statement 2-656
 write lock 2-292
 Locking granularity 2-141, 2-452, 2-634
 LOCKS configuration parameter 2-326, 2-451
 LOCKS keyword, in SET ISOLATION statement 2-636
 LOCOPY function 4-122, 4-126
 Log files, for load and unload jobs 2-645
 LOG keyword 2-640
 ALTER TABLE statement 2-62
 in CREATE DATABASE statement 2-97
 in CREATE TABLE statement 2-222
 in CREATE Temporary TABLE statement 2-231
 LOG_INDEX_BUILDS configuration parameter 6-46
 LOG10 function 4-116
 Logarithmic functions
 LOG10 function 4-116
 LOGN function 4-117
 Logging
 buffered versus unbuffered 2-640
 cascading deletes 2-305
 changing mode with SET LOG 2-639
 in CREATE DATABASE statement 2-98
 log space requirements 2-9
 table type options 2-193
 temporary tables 2-238
 with triggers 2-269
 Logical operator, in Condition segment 4-20
 LOGN function 4-116
 Lohandles support function 2-149
 LOLIST data type 2-537, 4-23
 Long transaction rollback 3-45
 Loop
 controlled 3-19
 indefinite with WHILE 3-34, 3-49
 LOOP keyword 3-19
 in CONTINUE statement 3-7
 in EXIT statement 3-17
 LOOP statement 3-34
 LOTOFILE function 4-122, 4-125
 LOW keyword 2-705
 in SET OPTIMIZATION statement 2-641
 in SET PDQPRIORITY statement 2-643
 LOWER function 4-156
 Lower index filter 2-628
 Lowercase characters, converting to 4-156
 LPAD function 4-154
 LTRIM function 4-149
 LVARCHAR data type 4-22, 4-24
 syntax 4-23

M

M keyword in SQL Admin API functions 6-2
 Mail, sending from SPL routines 3-45
 Mantissa 4-185

Master replicate 2-681
 MATCHED keyword in MERGE statement 2-452
 MATCHES keyword
 in Condition segment 4-6, 4-13
 in SELECT statement 2-562
 wildcard characters 2-563
 matches() operator function 4-14
 Materialized table expression 2-546
 Materialized view 2-275
 MAX function 4-165, 4-172
 MAX keyword
 in ALLOCATE DESCRIPTOR statement 2-3
 in CREATE TABLE statement 2-219
 in START VIOLATIONS TABLE statement 2-664
 MAX ROWS keywords, in START VIOLATIONS TABLE statement 2-664
 MAX VIOLATIONS keywords, in START VIOLATIONS TABLE statement 2-664
 MAX_INCOMPLETE_CONNECTIONS configuration parameter 6-41
 MAX_PDQPRIORITY configuration parameter 2-643, 6-23, 6-33, 6-41
 MAXERRORS environment variable 2-110
 MAXERRORS keyword
 in CREATE EXTERNAL TABLE statement 2-109
 MAXLEN keyword, in CREATE OPAQUE TYPE statement 2-147
 MAXSCAN keyword
 in SET ENVIRONMENT statement 2-618
 MAXVALUE keyword
 in ALTER SEQUENCE statement 2-42
 in CREATE SEQUENCE statement 2-186
 MDY function 4-128, 4-135
 MEDIUM keyword, in UPDATE STATISTICS statement 2-705
 MEDIUM mode 2-711
 Membership (.) operator 4-61
 Memory
 allocating for collection variable 2-1
 allocating for query 2-616
 allocating for ROW variable 2-4
 deallocating cursors 2-380
 deallocating for collection variable 2-283
 deallocating for cursors 2-582
 deallocating for row variable 2-285
 deallocating prepared objects 2-380, 2-584
 MEMORY keyword, in EXECUTE FUNCTION statement 6-56
 MEMORY_RESIDENT keyword
 in SET INDEX statement 2-630
 in SET Residency statement 2-631, 2-646
 in SET TABLE statement 2-656
 MERGE statement 2-452
 MESSAGE_LENGTH keyword, in GET DIAGNOSTICS statement 2-391
 MESSAGE_TEXT keyword, in GET DIAGNOSTICS statement 2-391
 Metadata function 6-62
 mi_collection* functions 5-64
 mi_trigger*() functions 2-259
 MIN function 4-165, 4-172
 MIN keyword, in CREATE TABLE statement 2-219
 Minus (-) sign
 binary operator 4-40
 Minus sign (-)
 INTERVAL literals 4-184
 unary operator 4-184
 minus() operator function 4-51
 MINUTE keyword 4-38, 4-176
 MINVALUE keyword
 in ALTER SEQUENCE statement 2-42
 in CREATE SEQUENCE statement 2-186
 Missing arguments 5-2
 Mixed-case characters, converting to 4-156
 MOD function 4-89, 4-91
 MODE keyword
 in ALTER INDEX statement 2-30
 in ALTER TABLE statement 2-68
 in CREATE DATABASE statement 2-97
 in CREATE INDEX statement 2-141
 in CREATE TABLE statement 2-225
 in CREATE Temporary TABLE statement 2-175
 in LOCK TABLE statement 2-449
 in SET LOCK MODE statement 2-637
 MODIFY EXTENT SIZE keyword
 in ALTER TABLE statement 2-66
 MODIFY EXTERNAL NAME keywords
 in ALTER FUNCTION statement 2-29
 in ALTER PROCEDURE statement 2-33
 in ALTER ROUTINE statement 2-35
 MODIFY keyword
 in ALTER ACCESS_METHOD statement 2-5
 in ALTER FRAGMENT statement 2-26
 in ALTER FUNCTION statement 2-28
 in ALTER PROCEDURE statement 2-32
 in ALTER ROUTINE statement 2-33
 in ALTER TABLE statement 2-57
 MODIFY NEXT SIZE keyword
 in ALTER TABLE statement 2-67
 Modifying routine modifiers
 with ALTER FUNCTION statement 2-29
 with ALTER PROCEDURE statement 2-32
 with ALTER ROUTINE statement 2-34
 Modulus 4-91
 MONEY data type
 literal values 4-186
 loading 2-445
 syntax 4-26
 MONTH function 4-128, 4-131
 MONTH keyword 4-38, 4-176
 MONTHS_BETWEEN function 4-128, 4-131
 MORE keyword, in GET DIAGNOSTICS statement 2-390
 MOVE TABLE statement 2-455
 MQ DataBlade module 5-54
 MSC VP class 6-38
 Multi-index scan 5-38
 Multibyte characters 4-24
 Multibyte locales 4-119
 Multilingual index 2-586
 Multiple triggers
 example 2-247
 preventing overriding 2-269
 Multiple-column constraints
 in ALTER TABLE statement 2-64
 in CREATE TABLE statement 2-205
 Multiplication sign (*), arithmetic operator 4-40
 Multirepresentational data 2-343, 2-437, 2-698
 Multirow query 2-370
 MULTISET columns, generating values for 4-83
 MULTISET data type
 collection subqueries 4-4
 defined 4-84
 deleting elements from 2-309
 unloading 2-684
 updating elements 2-703

MULTISET keyword 4-3, 4-178
in DEFINE statement 3-13
in Expression segment 4-83
in FROM clause of SELECT statement 4-5

Mutability
property 2-581
setting for all variables 2-581

MUTABLE keyword
in SET ALL_MUTABLES statement 2-581
in SET DEFAULT TABLE_SPACE statement 2-601
in SET DEFAULT TABLE_TYPE statement 2-602
in SET ENVIRONMENT statement 2-616
in SET PDQPRIORITY statement 2-644

N

NAME field 2-381
in SET DESCRIPTOR statement 2-606
with DESCRIBE INPUT statement 2-318
with DESCRIBE statement 2-313

NAME keyword
External Routine Reference segment 5-19
in ALTER FUNCTION statement 2-28
in ALTER PROCEDURE statement 2-32
in ALTER ROUTINE statement 2-33
in CREATE FUNCTION statement 2-113
in CREATE PROCEDURE statement 2-153

Named row type
assigning with ALTER TABLE 2-73
associating with a column 4-34
creating with CREATE ROW TYPE 2-168
dropping with DROP ROW TYPE 2-336
inheritance 2-169
privileges on 2-404
Under privilege 2-508
unloading 2-685, 2-687
updating fields 2-703

Naming convention
database 5-15
database objects 5-17

Natural logarithms 4-117

NCHAR data type 4-25
syntax 4-23

negate() operator function 4-52

Negator functions 2-357, 2-406, 2-502, 5-63

NEGATOR keyword
Routine Modifier segment 5-60, 5-63
NESTED keyword in optimizer directives 5-44

Nested loop join 2-621, 2-628, 5-41

Nested ordering
in SELECT statement 2-571

NET API 4-191

NEW keyword
in CREATE FUNCTION statement 2-116
in CREATE PROCEDURE statement 2-157
in CREATE TRIGGER statement 2-256
Insert triggers 2-255
Update triggers 2-256

NEW keyword, in CREATE TRIGGER statement 2-254

Newline characters in quoted strings 4-191

NEXT keyword
in ALTER TABLE statement 2-67
in CREATE TABLE statement 2-224
in FETCH statement 2-368

NEXT SIZE keywords
Extents
revising the size 2-67

NEXT SIZE keywords (*continued*)
in ALTER TABLE statement 2-67
in CREATE TABLE statement 2-224

NEXT_DAY function 4-128, 4-133

NEXTVAL operator 2-185, 4-79

NO KEEP ACCESS TIME keywords
in ALTER TABLE statement 2-62
in CREATE TABLE statement 2-222

NO keyword
in CREATE TEMP TABLE statement 2-230
in SET COLLATION statement 2-584

NO LOG keywords
in ALTER TABLE statement 2-62
in CREATE TABLE statement 2-222
in CREATE TEMP TABLE statement 2-230
in SELECT statement 2-574

NOCACHE keyword
in ALTER SEQUENCE statement 2-42
in CREATE SEQUENCE statement 2-186

NOCYCLE keyword
in ALTER SEQUENCE statement 2-42
in CREATE SEQUENCE statement 2-186

NODBPROC setting of DEBUG environment variable 6-52
NODEFDAC environment variable 2-116, 2-156, 2-402, 2-502
effects on new routine 2-116, 2-156
effects on new table 2-229
GRANT statement with 2-406

NOMAXVALUE keyword
in ALTER SEQUENCE statement 2-42
in CREATE SEQUENCE statement 2-186

NOMINVALUE keyword
in ALTER SEQUENCE statement 2-42
in CREATE SEQUENCE statement 2-186

NON_RESIDENT keyword
in SET INDEX statement 2-630
in SET Residency statement 2-631, 2-646
in SET TABLE statement 2-656

Noncursor function 5-59

Nondefault code sets 4-196

NONE keyword
in SET ENVIRONMENT statement 2-622
in SET ROLE statement 2-648

NONE role 2-165, 2-334

Nonlogging temporary tables
creating 2-233
duration 2-238

Nonvariant functions 5-20

NOORDER keyword
in ALTER SEQUENCE statement 2-43
in CREATE SEQUENCE statement 2-187

NORMAL keyword
in ALTER INDEX statement 2-30, 2-31
in CREATE INDEX statement 2-141

NOT AUTHORIZED keywords
in CREATE SECURITY POLICY statement 2-182

NOT bitwise logical operation 4-55

NOT FOUND keywords, in WHENEVER statement 2-718

NOT keyword
in ALTER INDEX statement 2-31
in Condition segment 4-6, 4-10, 4-13, 4-18
in MERGE statement 2-452
in SELECT statement 2-561, 2-562
in SET LOCK MODE statement 2-638
Routine Modifier segment 5-60
with BETWEEN keyword 2-562
with IN keyword 2-563

- NOT NULL keywords
 - in ALTER TABLE statement 2-50
 - in collection data type declarations 4-37
 - in CREATE EXTERNAL TABLE statement 2-107
 - in CREATE ROW TYPE statement 2-170
 - in CREATE TABLE statement 2-197
 - in CREATE Temporary TABLE statement 2-234
 - in DEFINE statement 3-13
 - in SELECT statement 2-561
- NOT VARIANT keywords, in External Routine Reference segment 5-19
- NOT WAIT keywords in SET LOCK MODE 2-638
- notequal() operator function 4-195
- NULL keyword 5-1
 - ambiguous as a routine variable 5-31
 - in ALTER TABLE statement 2-49
 - in Condition segment 4-9
 - in CREATE EXTERNAL TABLE statement 2-105
 - in CREATE ROW TYPE statement 2-171
 - in CREATE TABLE statement 2-195
 - in CREATE Temporary TABLE statement 2-234
 - in DEFINE statement 3-8
 - in Expression segment 4-66, 4-67, 4-69
 - in INSERT statement 2-434
 - in SELECT statement 2-561
 - in SET ROLE statement 2-648
 - in UPDATE statement 2-693, 2-694
- Null values
 - checking for in SELECT statement 2-350, 2-354
 - in IF statement 3-30
 - inserting with the VALUES clause 2-439
 - invalid for collection types 4-37
 - loading 2-445
 - returned implicitly by SPL function 3-43
 - updating a column 2-694
 - used in Condition with NOT operator 4-19
 - used in the ORDER BY clause 2-571
 - WHILE statement 3-49
 - with AND and OR keywords 4-20
 - with NULLIF function 4-68
 - with NVL function 4-68
- NULLABLE field 2-381
 - in SET DESCRIPTOR statement 2-606
 - with DESCRIBE INPUT statement 2-318
 - with DESCRIBE statement 2-313
- NULLIF function 4-68
- NUMBER keyword, in GET DIAGNOSTICS statement 2-390
- Numeric data types 4-26
- NVARCHAR data type 4-25
 - syntax 4-23
- NVL function 4-68

O

- Object-List format, in SET Database Object Mode statement 2-593
- Octal numbers 2-110, 2-578
- OCTET_LENGTH function 4-118, 4-119
- ODBC API 4-191
- OF keyword 2-689
 - in CREATE DUPLICATE statement 2-102
 - in CREATE TRIGGER statement 2-239, 2-247, 2-248, 2-271
 - in CREATE VIEW statement 2-274
 - in DECLARE statement 2-286
 - in DELETE statement 2-303
 - in DROP DUPLICATE statement 2-326
 - in SELECT statement 2-572

- OF TYPE keywords
 - in CREATE TABLE statement 2-227
 - in CREATE VIEW statement 2-274
- OFF keyword
 - in SET DATASKIP statement 2-598
 - in SET ENVIRONMENT statement 2-616, 2-620
 - in SET EXPLAIN statement 2-624
 - in SET PDQPRIORITY statement 2-643
 - in SET STATEMENT CACHE statement 2-652
 - in TRACE statement 3-47
- OLD keyword
 - in CREATE FUNCTION statement 2-116
 - in CREATE PROCEDURE statement 2-157
 - in CREATE TRIGGER statement
 - Delete triggers 2-254
 - Select triggers 2-256
 - Update triggers 2-256
- OLEDB API 4-191
- OLTP 2-330
- OLTP (on-line transaction processing) 2-193
- ON DELETE CASCADE keywords
 - in ALTER TABLE statement 2-52
 - in CREATE TABLE statement 2-199
 - restrictions with triggers 2-244
- ON EXCEPTION keyword 5-75
- ON EXCEPTION statement 3-38
- ON keyword
 - in ALTER FRAGMENT statement 2-7
 - in CREATE INDEX statement 2-124
 - in CREATE TABLE statement 2-201
 - in CREATE TRIGGER statement 2-242, 2-246, 2-248, 2-271
 - in GRANT FRAGMENT statement 2-423
 - in GRANT statement 2-403, 2-406, 2-407, 2-416, 2-421
 - in MERGE statement 2-452
 - in REVOKE FRAGMENT statement 2-517
 - in REVOKE statement 2-497, 2-500, 2-501, 2-503, 2-504, 2-511, 2-515
 - in SELECT statement 2-555
 - in SET DATASKIP statement 2-598
 - in SET ENVIRONMENT statement 2-616, 2-620
 - in SET EXPLAIN statement 2-624
 - in SET STATEMENT CACHE statement 2-652
 - in TRACE statement 3-47
- on triggering view 2-271
- oncheck utility 2-149, 4-118, 6-3
- ONCONFIG parameters
 - AUTO_REPREPARE 2-466, 2-478, 2-714
 - DATASKIP 2-598
 - DBC_CREATE_PERMISSION 2-98
 - DBSERVERALIASES 2-310, 2-439, 2-538, 2-699, 6-46, 6-62
 - DBSERVERNAME 2-310, 2-439, 2-538, 2-699, 6-46
 - DEADLOCK_TIMEOUT 2-638
 - DEF_TABLE_LOCKMODE 2-69, 2-226, 2-452, 2-634
 - DIRECTIVES 5-36
 - DS_ADM_POLICY 2-650
 - DS_NONPDQ_QUERY_MEM 2-570
 - DS_TOTAL_TMPSPACE 2-619
 - EXPLAIN_STAT 2-628
 - EXT_DIRECTIVES 2-523, 5-45
 - FILLFACTOR configuration parameter 2-134
 - IFX_EXTEND_ROLE 2-411, 2-507, 6-54
 - JVPCLASSPATH 6-60
 - LOCKS 2-326, 2-451
 - LOG_INDEX_BUILDS 6-46
 - MAX_PDQPRIORITY 2-643
 - OPT_GOAL 5-43
 - OPTCOMPIND 2-653, 5-42

- ONCONFIG parameters (*continued*)
 - SBS spacename 2-222, 4-114, 4-115
 - SDS TIMEOUT 6-46
 - STACKSIZE 5-66
 - STMT_CACHE 2-652
 - STMT_CACHE_HITS 2-654
 - STMT_CACHE_NOLIMIT 2-654
 - STMT_CACHE_NUMPOOL 2-655
 - STMT_CACHE_SIZE 2-654
 - SYSSBS spacename 2-708
 - TEMPTAB_NOLOG 2-232
 - USELASTCOMMITTED 2-450, 2-622, 2-634, 2-658
 - USEOSTIME 4-75
- ondblog utility 2-640
- ONDBSPACEDOWN configuration parameter 6-36
- oninit utility 4-104
- ONLIDX_MAXMEM configuration parameter 6-41
- ONLINE keyword
 - in CREATE INDEX statement 2-143
 - in DROP INDEX statement 2-330
- Online transaction processing 2-330
- ONLY keyword 2-689, 2-705
 - in DECLARE statement 2-286
 - in DELETE statement 2-303, 2-305
 - in SAVE EXTERNAL DIRECTIVES statement 2-523
 - in SELECT statement 2-544, 2-549
 - in SET TRANSACTION statement 2-656
 - in TRUNCATE (XPS) statement 2-682
 - in UPDATE statement 2-691
 - in UPDATE STATISTICS statement 2-709, 2-711
- onmode utility 2-652, 6-27, 6-30
- onspaces utility 2-8, 2-135, 2-426, 2-518, 6-3
- onstat utility 2-598, 2-618
- onutil utility 2-109, 2-113, 4-118
- Opaque data types
 - alignment of 2-147
 - as argument 2-147
 - associating with a column 4-33
 - creating 2-145
 - DESCRIBE with 2-384
 - dropping 2-345
 - extended identifier 2-384, 2-610
 - GET DESCRIPTOR with 2-384
 - in DELETE 2-309
 - in DROP TABLE 2-343
 - in dynamic SQL 2-610
 - in INSERT 2-436
 - in LOAD 2-448
 - in UPDATE 2-698
 - loading 2-445, 2-448
 - modifiers 2-147
 - name of 2-384, 2-610
 - naming 2-146
 - owner name 2-384, 2-610
 - support functions 2-147
 - unloading 2-685
 - with SET DESCRIPTOR 2-610
- OPEN statement 2-459
- Open-Fetch-Close Optimization 2-604
- OpenAdmin Tool for IDS (OAT) 2-706
- OPERATIONAL keyword
 - in ALTER TABLE statement 2-69
 - in CREATE TABLE 2-193
 - in CREATE TABLE statement 2-191
 - in SELECT statement 2-575
 - in SET Default Table Type statement 2-602
- Operator class
 - btree_ops 2-153
 - creating 2-149
 - default 2-153, 5-52
 - default for B-Tree 2-153
 - defined 2-132, 2-149
 - dropping with DROP OPCLASS 2-331
 - rtree_ops 2-153
 - specifying with CREATE INDEX 2-127, 2-132
- Operator function
 - concat() 4-56
 - divide() 4-51
 - equal() 4-195
 - greaterthan() 4-195
 - greaterthanorequal() 4-195
 - lessthan() 4-195
 - lessthanorequal() 4-195
 - like() 4-14
 - matches() 4-14
 - minus() 4-51
 - negate() 4-52
 - notequal() 4-195
 - plus() 4-51
 - positive() 4-52
 - times() 4-51
- OPT VP class 6-38
- OPT_GOAL configuration parameter 5-43
- OPT_GOAL environment variable 5-43
- OPTCOMPIND configuration parameter 5-42
- OPTCOMPIND environment variable 2-621, 2-653, 6-51
- OPTCOMPIND keyword, in SET ENVIRONMENT statement 2-621
- Optical Subsystem
 - list of statements 1-10
- Optimization
 - specifying a high or low level 2-641
- OPTIMIZATION keyword
 - in SET OPTIMIZATION statement 2-641
- Optimizer
 - and SAVE EXTERNAL DIRECTIVES statement 2-522
 - and SET OPTIMIZATION statement 2-641
 - Optimizer Directives segment 5-35
 - strategy functions 2-151
 - with UPDATE STATISTICS 2-716
- Optimizer directives
 - /BROADCAST 5-40
 - /BUILD 5-40
 - /PROBE 5-40
 - access-method 5-37
 - ALL_ROWS 5-42
 - AVOID_EXECUTE 5-43
 - AVOID_FULL 5-38
 - AVOID_HASH 5-40
 - AVOID_INDEX 5-38
 - AVOID_INDEX_SJ 5-38
 - AVOID_NL 5-40
 - comment symbols 5-36
 - EXPLAIN 5-43
 - explain-mode 5-43
 - external 2-523
 - FIRST_ROWS 5-42
 - FULL 5-38
 - INDEX 5-38
 - INDEX_ALL 5-38
 - INDEX_SJ 5-38
 - inline 2-523
 - join-method 5-40

Optimizer directives (*continued*)

- join-order 5-40
- NESTED 5-44
- not followed 2-557
- optimization-goal 5-42
- ORDERED 5-40
- restrictions 5-37, 5-38
- rewrite method 5-44
- segment 5-35
- USE_HASH 5-40
- USE_NL 5-40

Optimizing

- a database server 2-641
- a query 2-522, 2-624
- across a network 2-642

OPTION keyword 2-396

- in CREATE TRIGGER statement 2-274
- in CREATE VIEW statement 2-277
- in GRANT FRAGMENT statement 2-423

OR bitwise logical operation 4-54

OR keyword

- defined 4-20
- in ANSI Joined Tables segment 2-555
- in Condition segment 4-6

ORDER BY keywords

- in SELECT statement 2-568
- restricted in INSERT 2-441

ORDER keyword

- in ALTER SEQUENCE statement 2-43
- in CREATE SEQUENCE statement 2-187

ORDERED keyword, in optimizer directives 5-40

OUT keyword 5-67

OUT parameter

- default parameter style 5-20
- user-defined function 5-69
- with a statement-local variable 4-162, 4-164

OUTER keyword in SELECT statement 2-555

OUTPUT statement 2-466

Output support function 2-148

Overflow bin 2-712

Overloaded routine 5-2, 5-19

OVERRIDE keyword

- in CREATE SECURITY POLICY statement 2-182

Owner 5-16

- ANSI-compliant database 5-48
- case-sensitivity 2-426, 2-495, 2-506, 2-518, 5-46
- Database Object Name segment 5-45
- in ANSI-compliant database 2-426, 2-495, 2-506, 2-518
- in CREATE SYNONYM 2-189
- in DROP SEQUENCE 2-340
- in RENAME TABLE statement 2-493
- in system catalog table 2-66
- Owner Name segment 5-45

Owner Name segment 5-45

Owner-privileged UDR 2-115, 2-156

P

P keyword in SQL Admin API functions 6-2

Package, jar name component 5-34

Packed decimal 2-106

PACKED keyword

- in CREATE EXTERNAL TABLE statement 2-105

PAGE keyword

- DEF_TABLE_LOCKMODE setting 2-634
- IFX_DEF_TABLE_LOCKMODE setting 2-634
- in ALTER TABLE statement 2-68

PAGE keyword (*continued*)

- in CREATE TABLE statement 2-225
- in CREATE Temporary TABLE statement 2-175

Page number 4-64

Page-level locking, in CREATE TABLE 2-226

Parallel distributed queries

- SET PDQPRIORITY statement 2-643

Parallelizable data query 5-63

PARALLELIZABLE keyword, in Routine Modifier segment 5-60, 5-62

Parameter

- BYTE or TEXT in SPL 3-16
- dynamic 2-317
- Java method 5-73
- UDRs 5-2

PARAMETER keyword

- External Routine Reference segment 5-19

Parameterizing

- prepared statements 2-353

Parameterizing a statement, with SQL identifiers 2-474

Parent table 2-457

Parent-child relationship 2-200

PARTITION BY keywords

- in ALTER FRAGMENT statement 2-20

PARTITION keyword

- in ALTER FRAGMENT statement 2-9, 2-16, 2-20, 2-21, 2-22, 2-23, 2-26
- in CREATE INDEX statement 2-136
- in CREATE TABLE statement 2-212

Partition number 4-101

Partitions 2-426, 2-518

PASSEDBYVALUE keyword, in CREATE OPAQUE TYPE statement 2-147

passwd file 2-87

PASSWORD keyword

- in SET ENCRYPTION PASSWORD statement 2-611

Pathnames with commas 4-124

PB keyword in SQL Admin API functions 6-2

PDQ

- CPU VPs 6-38
- SET ENVIRONMENT statement 2-615
- SET PDQPRIORITY statement 2-643

PDQ thread safe functions 5-64

PDQPRIORITY environment variable 2-616, 2-643, 2-644, 2-653

PDQPRIORITY keyword

- in SET PDQPRIORITY statement 2-643

PERCALL_COST keyword, Routine Modifier segment 5-60

Percent (%) sign

- as wildcard 4-13

Period symbol (.)

- DATETIME separator 4-182
- DECIMAL values 4-186
- INTERVAL separator 4-184
- membership operator 4-61
- MONEY values 4-186

Phantom row 2-633, 2-659

PIO VP class 6-37

Pipe character (|) 2-110, 2-443, 2-577, 2-687

PIPE keyword

- in CREATE EXTERNAL TABLE statement 2-108
- in OUTPUT statement 2-467

PLOAD keyword

- in SET PLOAD FILE statement 2-646

Plus (+) sign

- arithmetic operator 4-40

- Plus operator (+)
 - unary 4-184, 4-185
- Plus sign (+)
 - in optimizer directives 5-36
 - unary operator 4-184, 4-185
- plus() operator function 4-51
- POINTER data type 2-536, 4-22
- Polar coordinates 4-141
- POLICY keyword
 - in ALTER TABLE statement 2-70
 - in CREATE SECURITY POLICY statement 2-182
 - in CREATE TABLE statement 2-229
 - in DROP SECURITY statement 2-338
 - in RENAME SECURITY statement 2-490
- positive() operator function 4-52
- POW function 4-89, 4-91
- POWER function 4-89, 4-91
- Precedence in dot notation 4-62
- PRECISION field 2-381
 - in SET DESCRIPTOR statement 2-606
 - with DESCRIBE INPUT statement 2-318
 - with DESCRIBE statement 2-313
- PREPARE statement
 - deferring 2-603
 - for collection variables 2-471
 - increasing efficiency 2-477
 - multistatement text 2-362, 2-476
 - parameterizing a statement 2-473
 - parameterizing for SQL identifiers 2-474
 - question (?) mark as placeholder 2-468
 - releasing resources with FREE 2-380
 - restrictions with SELECT 2-471
 - statement identifier 2-297
 - statement identifier use 2-469
 - syntax 2-468
 - valid statement text 2-470
 - with external routines 2-471
 - with SPL routines 2-471
- Prepared statement
 - comment symbols in 2-470
 - DESCRIBE statement with 2-311
 - executing 2-348
 - parameterizing 2-353
 - prepared object limit 2-287, 2-469
 - setting PDQ priority 2-644
 - valid statement text 2-470
 - with DESCRIBE INPUT statement 2-316
- Preserving newline characters in quoted strings 4-191
- PREVIOUS keyword, in FETCH statement 2-368
- Primary key column, no NULL default 2-195
- Primary key constraint
 - moving 2-457
- PRIMARY KEY keywords
 - in ALTER TABLE statement 2-50, 2-64
 - in CREATE TABLE statement 2-197, 2-205, 2-227
 - in CREATE Temporary TABLE statement 2-234
- PRIMARY keyword
 - in CREATE ACCESS_METHOD statement 2-91
- Primary-key constraint
 - data type conversion 2-60
 - defining column as 2-199
 - dropping 2-66
 - requirements for 2-51, 2-199
 - rules of use 2-200
 - using 2-199
- PRINT ERROR command 6-6
- PRINT PARTITION command 6-6
- PRIOR keyword, in FETCH statement 2-368
- PRIVATE keyword
 - in CREATE SYNONYM statement 2-187
- Privilege 2-397, 2-496
 - Alter 2-401
 - chaining grantors 2-507
 - column-specific 2-499
 - Connect 2-398
 - DBA 2-399, 2-496
 - effect of NODEFDAC 2-402
 - Execute 2-405, 2-501
 - for triggered action 2-266
 - fragment-level 2-423
 - revoking 2-517
 - granting 2-395
 - in system calls 3-45
 - needed, to create a cast 2-95
 - on a synonym 2-188
 - on a view 2-275
 - on languages 2-406, 2-503
 - on named row type 2-404
 - on remote objects 2-648
 - on sequences 2-407, 2-503
 - on table fragments 2-423
 - on UDRs called by a trigger 2-266
 - Resource 2-398
 - table-level 2-497
 - ANSI-compliant 2-402
 - column-specific 2-400
 - effect on view 2-403
 - Usage 2-501
- PRIVILEGES keyword
 - in GRANT statement 2-401
 - in INFO statement 2-428
 - in MOVE TABLE statement 2-455
 - in REVOKE statement 2-497
- Procedural language 4-75
- Procedure
 - altering with ALTER PROCEDURE 2-32
 - creating from file 2-163
 - dropping with DROP PROCEDURE 2-332
 - dropping with DROP ROUTINE 2-335
 - modifying path to executable file 2-33
 - modifying routine modifiers 2-32
 - privileges 2-156
 - specific name 5-73
 - system catalog tables for 2-160
 - trigger 2-157
 - user-defined, definition 2-155
- Procedure cursor
 - opening 2-461
- PROCEDURE keyword
 - DEFINE statement 3-8
 - in ALTER PROCEDURE statement 2-32
 - in CREATE PROCEDURE statement 2-153
 - in DECLARE statement 2-286
 - in DROP PROCEDURE statement 2-332
 - in EXECUTE PROCEDURE statement 2-364
 - in GRANT statement 2-405
 - in REVOKE statement 2-501
 - in TRACE statement 3-47
 - in UPDATE STATISTICS statement 2-714
- Projection
 - clause 2-531
 - column with dot notation 4-61
 - field projection 4-61
- Projection clause 2-531

- Projection list 2-531
- Promotable lock 2-292
- Pseudo-table 2-454
- Pseudo-users 2-507
- PUBLIC keyword
 - in CREATE SYNONYM statement 2-187
 - in GRANT FRAGMENT statement 2-426
 - in GRANT statement 2-408, 2-421
 - in REVOKE FRAGMENT statement 2-518
 - in REVOKE statement 2-504, 2-515
- Purpose flags
 - adding and deleting 2-6
 - list 5-52
- Purpose functions
 - adding, changing, and dropping 2-6
 - for access methods 2-680, 5-52
 - for XA data source types 5-54
 - parallel-execution indicator 5-52
- Purpose options
 - specifying 5-51
 - valid settings 5-52
- Purpose values
 - adding, changing, and dropping 2-6
- PUT keyword
 - in ALTER TABLE statement 2-62
 - in CREATE TABLE statement 2-222
- PUT statement
 - FLUSH with 2-479
 - source of row values 2-480
 - syntax 2-479
 - use in transactions 2-480

Q

- Qualifier, field 4-38, 4-176, 4-181, 4-183
- Qualifying rows 2-528
- Query
 - distributed 2-536, 2-537, 2-612, 4-22, 4-33, 5-18
 - external databases 5-17
 - external directives 2-525
 - operator overflow 2-619
 - optimizing
 - optimizer directives 5-35
 - optimizing prepared statements 2-655
 - optimizing with SAVE EXTERNAL DIRECTIVES 2-522
 - optimizing with SET OPTIMIZATION 2-641
 - pipng results to another program 2-467
 - priority level 2-643
 - qualifying rows 2-528
 - remote databases 5-18
 - result set 2-528
 - scheduling level for 2-649
 - sending results to an operating-system file 2-467
 - sending results to another program 2-467
 - statistics 2-628
- Query optimizer
 - recalculating distributions 2-705
- Question mark (?)
 - as placeholder in PREPARE 2-468
 - as wildcard 4-14
 - dynamic parameters 2-317
 - generating unique large-object filename 4-125
 - naming variables in PUT 2-481
- Question Mark (?)
 - placeholder in PREPARE 2-463
- Quotation marks 4-189, 5-23
 - delimited identifiers 5-24

- Quotation marks (*continued*)
 - double 5-26
 - effects of DELIMIDENT environment variable 5-24
 - literal in a quoted string 4-192
 - literal nested collection 4-180
 - owner name 5-46
 - quoted string delimiter 4-192
 - single 5-24
- Quoted Pathname segment 5-70
- Quoted string 4-189
 - as constant expression 4-72
 - DATETIME values as strings 4-192
 - effects of DELIMIDENT environment variable 5-24
 - in INSERT 2-434, 4-192
 - INTERVAL values as strings 4-192
 - maximum length 4-192
 - newline characters 4-160
 - newline characters in 4-191
 - wildcards 4-192
 - with LIKE keywords 2-562
- quoted string delimiter 4-189

R

- R-tree index
 - creating 2-133, 2-143
 - default operator class 2-153
 - detached 2-136
 - dropping 2-331
 - rtree_ops operator class 2-153
 - uses 2-133
- R-tree secondary-access method 2-133, 2-149
- Radicand 4-91
- Radix-64 encryption format 2-612
- RAISE EXCEPTION statement 3-41
- Range fragmentation 2-44
- RANGE function 4-165, 4-172
- RANGE keyword, in CREATE TABLE statement 2-218
- RAW keyword
 - in ALTER TABLE statement 2-69
 - in CREATE TABLE 2-193
 - in CREATE TABLE statement 2-191
 - in SELECT statement 2-574
 - in SET Default Table Type statement 2-602
- Read access rules for label-based access 2-182
- Read Committed isolation level 2-622, 2-658
- READ COMMITTED keywords, in SET TRANSACTION statement 2-656
- READ keyword
 - in CREATE SECURITY POLICY statement 2-182
 - in GRANT statement 2-418
 - in REVOKE statement 2-513
 - in SET ENVIRONMENT statement 2-622
 - in SET ISOLATION statement 2-633, 2-634
 - in SET TRANSACTION statement 2-656
- READ ONLY keywords
 - in DECLARE statement 2-286
 - in SELECT statement 2-573
 - in SET TRANSACTION statement 2-656
- Read permission 6-59
- Read Uncommitted isolation level 2-622, 2-658
- READ UNCOMMITTED keywords, in SET TRANSACTION statement 2-656
- READ WRITE keywords, in SET TRANSACTION statement 2-656
- REAL data type 4-29
- Real numbers 4-29

- Receive support function 2-148
- RECORDEND environment variable 2-110, 2-578
- RECORDEND keyword
 - in CREATE EXTERNAL TABLE statement 2-109
 - in SELECT statement 2-578
- REFERENCES keyword 5-55
 - in ALTER TABLE statement 2-52
 - in CREATE TABLE statement 2-199
 - in DEFINE statement 3-8
 - in EXECUTE FUNCTION statement 2-356
 - in EXECUTE PROCEDURE statement 2-364
 - in GRANT statement 2-400
 - in INFO statement 2-428
 - in REVOKE statement 2-497
- References privilege
 - defined 2-400
 - displaying 2-429
 - revoking 2-497
- REFERENCING keyword
 - in CREATE FUNCTION statement 2-116
 - in CREATE PROCEDURE statement 2-157
 - in CREATE TRIGGER statement 2-271
 - Delete triggers 2-254
 - Insert triggers 2-255
 - Select triggers 2-256
 - Update triggers 2-256
 - view column values 2-270
- Referential constraint 2-104
 - B-tree index 2-126
 - Dataskip feature 2-598
 - defining 2-200
 - delete triggers 2-244
 - dropping 2-66
 - locking 2-201
- Referential integrity 2-306
- REJECTFILE keyword
 - in CREATE EXTERNAL TABLE statement 2-109, 2-111
- Relational operators 4-193
 - with WHERE keyword in SELECT 2-561
- RELATIVE keyword, in FETCH statement 2-368
- RELEASE keyword
 - in RELEASE SAVEPOINT statement 2-485
- RELEASE SAVEPOINT statement 2-485
- REMAINDER IN keywords
 - in ALTER FRAGMENT statement 2-20, 2-21, 2-23, 2-26
 - in CREATE INDEX statement 2-136
 - in CREATE TABLE statement 2-212, 2-218, 2-219
- Remote procedure
 - restrictions on optimizing 2-706
- Remote query 2-536, 2-537
- RENAME COLUMN statement 2-487
- RENAME DATABASE statement 2-488
- RENAME INDEX statement 2-489
- RENAME keyword in MOVE TABLE statement 2-455
- RENAME SECURITY LABEL COMPONENT statement 2-490
- RENAME SECURITY LABEL statement 2-490
- RENAME SECURITY POLICY statement 2-490
- RENAME SECURITY statement 2-490
- RENAME SEQUENCE statement 2-491
- RENAME SPACE command 6-6
- RENAME TABLE statement 2-492
- REOPTIMIZATION keyword in OPEN statement 2-460
- Reoptimizing query plans 2-706
- Repeatable Read isolation level 2-99, 2-621, 2-635, 2-659
- Repeatable Read isolation level, emulating during
 - update 2-374
- REPEATABLE READ keywords
 - in SET ISOLATION statement 2-631
 - in SET TRANSACTION statement 2-656
- REPLACE function 4-153
- REPLICATION keyword
 - in BEGIN WORK statement 2-74
- Reserved words
 - delimited identifiers 5-24
 - identifiers 5-23
 - SQL A-1, B-1
 - using in triggered action 2-260
- RESET SYSADMIN command 6-6, 6-49
- RESIDENT configuration parameter 6-23, 6-39, 6-41
- RESOLUTION keyword, in UPDATE STATISTICS
 - statement 2-711
- Resource Grant Manager 2-645, 2-650
- RESOURCE keyword 2-397, 2-496
- Resource privilege 2-398, 2-496
 - with CREATE ACCESS_METHOD statement 2-91
- Resource usage policy 2-581
- RESTART keyword, in ALTER SEQUENCE statement 2-42
- RESTRICT keyword 2-494
 - in CREATE SECURITY POLICY statement 2-182
 - in DROP ACCESS_METHOD statement 2-323
 - in DROP OPCLASS statement 2-331
 - in DROP ROW TYPE statement 2-336
 - in DROP SECURITY statement 2-338
 - in DROP TABLE statement 2-342
 - in DROP TYPE statement 2-345
 - in DROP VIEW statement 2-346
 - in DROP XADATASOURCE statement 2-347
 - in DROP XADATASOURCE TYPE statement 2-348
 - in MOVE TABLE statement 2-455
- RESTRICTED mode of UDRs 2-651
- Result sets 2-528, 2-551
- RESUME keyword
 - in ON EXCEPTION statement 3-38
 - in RETURN statement 3-42
- RETAIN UPDATE LOCKS keywords
 - in SET ISOLATION statement 2-631
- Return Clause segment 5-55
- RETURN statement 3-42
- Return value
 - declaring in CREATE FUNCTION 5-55
- RETURNED_SQLSTATE field 2-310, 2-375
- RETURNED_SQLSTATE keyword, in GET DIAGNOSTICS
 - statement 2-391
- RETURNING keyword 5-55
 - example 2-121
 - in CALL statement 3-4
- RETURNS keyword 5-55
 - in Java Shared-Object-File segment 5-72
- REUSE keyword
 - in TRUNCATE statement 2-678
- REVOKE FRAGMENT statement 2-517
- REVOKE statement 2-494
- revoking privileges 2-496
- RGM 2-645, 2-650
- RIGHT keyword
 - in ANSI Joined Tables segment 2-555
 - in SELECT statement 2-557
- Right outer joins 2-553
- ROBIN keyword
 - in ALTER FRAGMENT statement 2-20
 - in CREATE TABLE statement 2-212
- Role
 - activating with SET ROLE 2-647

- Role (*continued*)
 - built-in 2-165, 2-334
 - case-sensitivity 2-397
 - creating with CREATE ROLE 2-165
 - currently enabled 4-70
 - default 2-410, 2-506, 4-70
 - default roles 2-649
 - definition 2-165
 - dropping with DROP ROLE statement 2-334
 - enabling with SET ROLE 2-648
 - establishing with CREATE, GRANT, SET 2-409
 - EXTEND 2-411, 2-507
 - granting privileges with GRANT 2-410
 - granting role with GRANT 2-409
 - revoking privileges 2-505
 - scope of 2-648
- ROLE keyword
 - in CREATE ROLE statement 2-164
 - in DROP ROLE statement 2-334
 - in GRANT statement 2-410, 2-421
 - in REVOKE statement 2-506, 2-515
 - in SET ROLE statement 2-648
- ROLES keyword in MOVE TABLE statement 2-455
- ROLLBACK WORK statement 2-75, 2-520
 - with WHENEVER 2-81
- root dbspace 2-98, 6-49
- ROOT function 4-89, 4-91
- ROOT keyword
 - in CREATE SECURITY LABEL COMPONENT statement 2-177
- ROUND function 4-89, 4-92
- ROUND ROBIN keywords
 - in ALTER FRAGMENT statement 2-20
 - in CREATE TABLE statement 2-212
- Rounding error 4-196
- ROUTINE keyword
 - in ALTER ROUTINE statement 2-33
 - in DROP ROUTINE statement 2-335
 - in GRANT statement 2-405
 - in REVOKE statement 2-501
 - in UPDATE STATISTICS statement 2-714
- Routine manager 6-61
- Routine modifier
 - CLASS 5-60
 - COSTFUNC 5-62
 - HANDLESNULLS 5-62
 - INTERNAL 5-62
 - ITERATOR 5-63
 - NEGATOR 5-63
 - NOT VARIANT 5-66
 - PARALLELIZABLE 5-63
 - PERCALL_COST 5-64
 - SELCONST 5-65
 - SELFUNC 5-65
 - STACK 5-66
 - VARIANT 5-66
- Routine Parameter List segment 5-67
- Routine signature 5-19, 5-69
- Routines
 - altering with ALTER ROUTINE 2-33
 - built-in 6-1
 - checking references 2-266
 - creating with CREATE ROUTINE FROM 2-166
 - dropping with DROP ROUTINE 2-335
 - modifying
 - path to executable file 2-33
 - routine modifiers 2-34
- Routines (*continued*)
 - modifying path to executable file 2-35
 - overloading 5-19
 - privileges 2-156
 - restrictions in triggered action 2-265
 - specific name 5-73
 - trigger 2-116, 2-157
- ROW constructor, in Expression segment 4-82
- ROW data types
 - collection-derived tables 5-7
 - constructor syntax 4-82
 - dot notation 4-61
 - loading field values 2-445, 2-447
 - nested 4-189
 - privileges 2-404
 - selecting fields 2-540, 2-550
 - selecting from 2-550
 - unloading 2-685, 2-687
 - updating 2-696, 2-703
- ROW keyword 4-186
 - in ALLOCATE ROW statement 2-4
 - in ALTER TABLE statement 2-68
 - in CREATE ROW TYPE statement 2-168
 - in CREATE TABLE statement 2-225
 - in CREATE Temporary TABLE statement 2-175
 - in CREATE TRIGGER statement 2-251, 2-257
 - in DROP ROW TYPE statement 2-337
 - in Expression segment 4-82
- Row variable
 - accessing 5-14
 - allocating memory 2-4
 - deallocating memory for 2-285
 - inserting 2-437
 - inserting into 2-442
 - selecting from 2-550
 - updating 2-703
- ROW_COUNT keyword, in GET DIAGNOSTICS statement 2-390
- Row-column level encryption 4-108
- Row-level locking, in CREATE TABLE 2-226
- Row-type columns, generating values for 4-82
- ROWID
 - adding column with INIT clause 2-19
 - adding with ALTER TABLE 2-46
 - dropping from fragmented tables 2-46
 - specifying support 5-52
 - use in a column expression 4-64
 - use in fragmented tables 2-19
 - used as column name 5-28, 5-29
- rowid column 2-19, 2-214, 2-571
- ROWID keyword, in Expression segment 4-59
- ROWIDS keyword
 - in ALTER FRAGMENT statement 2-19
 - in ALTER TABLE statement 2-46
 - in CREATE TABLE statement 2-212
- Rows
 - deleting 2-303
 - finding location 4-64
 - inserting
 - through a view 2-432
 - with a cursor 2-433
 - order of qualifying rows 2-532
 - phantom row 2-633
 - retrieving with FETCH 2-370
 - rowid defined 2-370
 - uncommitted row 2-633
 - updating through a view 2-691

- Rows (*continued*)
 - waiting for a locked row 2-638
 - writing buffered rows with FLUSH 2-377
- ROWS keyword, in START VIOLATIONS TABLE statement 2-664
- RPAD function 4-155
- RSAM access method 2-575
- RTNPARAMS data type 2-536
- RTNPARAMTYPES data type 4-22
- RTO_SERVER_RESTART configuration parameter 6-23, 6-41
- RTRIM function 4-149
- RULE keyword
 - in GRANT statement 2-416
 - in REVOKE statement 2-511

S

- SAMEAS keyword
 - in CREATE EXTERNAL TABLE statement 2-105
- Sampled queries 2-548
- SAMPLES OF keywords, in SELECT statement 2-548
- Sampling data 2-548
- SAMPLING keyword, in UPDATE STATISTICS statement 2-711
- SAVE EXTERNAL DIRECTIVES statement 2-522
- SAVEPOINT keyword
 - in RELEASE SAVEPOINT statement 2-485
 - in ROLLBACK WORK statement 2-521
 - in SAVEPOINT statement 2-526
- SAVEPOINT statement 2-526
- Savepoints
 - destroying 2-485
 - setting 2-526
- SBS spacename configuration parameter 2-222, 4-114, 4-115
- sbspaces
 - specifying in ALTER TABLE 2-62
 - specifying in CREATE TABLE 2-222
- SCALE field, with DESCRIBE INPUT statement 2-318
- SCALE field, with DESCRIBE statement 2-313
- SCALE keyword 2-381
 - in SET DESCRIPTOR statement 2-606
- Scan cost 2-6
- Scan threads 2-618
- SCHEDULE keyword
 - in SET SCHEDULE LEVEL statement 2-649
- SCHEDULER START command 6-6
- Scheduling level 2-650
- Schema name 5-46
- Scope of reference
 - global 2-288, 2-349, 3-10
 - in subqueries with UNION 2-580
 - local 3-12
 - module 2-288, 2-349
 - static 2-261
- SCRATCH keyword
 - in CREATE Temporary TABLE statement 2-175, 2-231
 - in SELECT statement 2-574
 - in SET Default Table Type statement 2-602
- Scratch table
 - creating 2-232, 2-578
 - duration 2-238
- Screen reader
 - reading syntax diagrams C-1
- Scroll cursors
 - defined 2-293
 - with FETCH 2-369
 - WITH HOLD 2-637
- SCROLL keyword, in DECLARE statement 2-286
- SDS_TIMEOUT configuration parameter 6-46
- SECLABEL_BY_COMP function 4-120
- SECLABEL_BY_NAME function 4-120
- SECLABEL_TO_CHAR function 4-121
- SECOND keyword 4-38, 4-176
- SECONDARY keyword
 - in CREATE ACCESS_METHOD statement 2-91
- Secondary-access methods
 - B-tree 2-133, 2-151
 - default operator class 2-153
 - defined 2-125, 2-149
 - R-Tree 2-149
 - R-tree 2-133
 - registering 2-91
 - USING clause 2-132
- Secure auditing 4-109
- SECURED WITH keywords
 - in ALTER TABLE statement 2-55, 2-60
 - in CREATE TABLE statement 2-193
- SECURITY keyword
 - in ALTER SECURITY LABEL COMPONENT statement 2-36
 - in ALTER TABLE statement 2-60, 2-70
 - in CREATE SECURITY LABEL COMPONENT statement 2-177
 - in CREATE SECURITY LABEL statement 2-175
 - in CREATE SECURITY POLICY statement 2-182
 - in CREATE TABLE statement 2-229
 - in DROP SECURITY statement 2-338
 - in GRANT statement 2-418
 - in RENAME SECURITY statement 2-490
 - in REVOKE statement 2-513
- Security label
 - assigning to a column 2-55, 2-60
 - creating 2-175
 - dropping from a column 2-60
 - identifier 4-120
 - renaming 2-338, 2-490
 - string format 4-121
- Security label component
 - for a security policy 2-182
 - renaming 2-338, 2-490
- SECURITY LABEL COMPONENT keywords
 - in DROP SECURITY statement 2-338
 - in RENAME SECURITY statement 2-490
- Security label components
 - creating 2-36, 2-177
- SECURITY LABEL keywords
 - in ALTER SECURITY LABEL COMPONENT statement 2-36
 - in CREATE SECURITY LABEL COMPONENT statement 2-177
 - in CREATE SECURITY LABEL statement 2-175
 - in CREATE SECURITY POLICY statement 2-182
 - in DROP SECURITY statement 2-338
 - in GRANT statement 2-418
 - in RENAME SECURITY statement 2-490
 - in REVOKE statement 2-513
- Security label support functions
 - SECLABEL_BY_COMP 4-120
 - SECLABEL_BY_NAME 4-120
 - SECLABEL_TO_CHAR 4-121
 - syntax 4-119
- Security policy
 - associating with a table 2-70
 - creating 2-182

- Security policy (*continued*)
 - renaming 2-338, 2-490
- SECURITY POLICY keywords
 - in ALTER TABLE statement 2-70
 - in CREATE SECURITY POLICY statement 2-182
 - in CREATE TABLE statement 2-229
 - in DROP SECURITY statement 2-338
 - in RENAME SECURITY statement 2-490
- segment 4-189
- Segment
 - defined 4-1, 5-1
- SELCONST keyword routine modifier 5-60
- Select cursor
 - declaring 2-288
 - opening 2-460, 2-461
 - reopening 2-461
- SELECT INTO clause
 - no table expressions 2-546
- SELECT ITEM keywords, in Collection-Subquery segment 4-3
- SELECT keyword 2-430, 4-3
 - ambiguous use as routine variable 5-31
 - in Condition segment 4-17, 4-18
 - in CREATE INDEX statement 2-141
 - in CREATE TRIGGER statement 2-247
 - in CREATE VIEW statement 2-276
 - in DECLARE statement 2-286
 - in GRANT statement 2-408
 - in LET statement 3-32
 - in OUTPUT statement 2-467
 - in REVOKE statement 2-497, 2-504
 - in UNLOAD statement 2-684
- Select list 2-531
- Select privilege 2-400, 2-408, 2-497, 2-504
- SELECT statements 3-23
 - aggregate functions in 4-165
 - BETWEEN condition 2-562
 - collection with 2-549
 - column numbers 2-571
 - cursor for 2-572, 2-573
 - FIRST clause 2-533, 2-534
 - FOR READ ONLY clause 2-573
 - FOR UPDATE clause 2-572
 - FROM clause 2-544
 - GROUP BY clause 2-566
 - HAVING clause 2-567
 - IN condition 2-561
 - in FOR EACH ROW trigger 2-253
 - in INSERT 2-441
 - indicator variables 2-359
 - INTO clause with ESQL 2-541
 - INTO EXTERNAL clause 2-577
 - INTO SCRATCH clause 2-578
 - INTO TEMP clause 2-575
 - IS NULL condition 2-562
 - joining tables in WHERE clause 2-564
 - LIKE or MATCHES condition 2-563
 - null values in the ORDER BY clause 2-571
 - ORDER BY clause 2-568
 - outer join 2-560
 - Projection clause 2-531
 - relational-operator condition 2-561
 - restrictions in routine 5-76
 - restrictions with INTO clause 2-471
 - row type 2-540, 2-550
 - ROWID keyword 4-64
 - select numbers 2-571
- SELECT statements (*continued*)
 - singleton 2-542
 - SKIP option 2-532
 - smart large objects with 4-64
 - SPL routine in 2-540
 - subquery with WHERE keyword 2-561
 - syntax 2-528
 - UNION operator 2-578
 - use of expressions 2-538
 - user-defined routine in 2-540
 - with DECLARE 2-286
 - with LET 3-33
 - WITH NO LOG keywords 2-576
 - writing rows retrieved to an ASCII file 2-683
- Select triggers 2-248
- Selecting from a specific table in a table hierarchy 2-549
- SELECTING operator 2-118, 2-159, 4-12
- Selectivity
 - functions 5-65
 - of the WHERE clause 2-712
- Selectivity functions 2-708
- Selectivity of an index key 5-38
- Self-join
 - defined 2-565
 - path 2-628
 - with aliases 2-545
- SELFUNC keyword routine modifier 5-60
- SELFUNCARGS data type 2-536, 4-22, 5-65
- Semantic integrity 2-243, 2-440
- Semicolon (;)
 - SPL statement block delimiter 2-160
 - statement terminator 2-173
- Send support function 2-148
- SENDRECV data type 2-537, 4-23
- SEQ_CACHE_SIZE configuration parameter 2-42, 2-187
- Sequence
 - cache 2-187
 - creating a synonym for 2-187
 - generator 2-185
 - privileges on 2-407, 2-503
- SEQUENCE keyword
 - in ALTER SEQUENCE statement 2-41
 - in CREATE SEQUENCE statement 2-184
 - in DROP SEQUENCE statement 2-340
 - in RENAME SEQUENCE statement 2-491
- Sequential cursor
 - with DECLARE 2-294
 - with FETCH 2-369
- SERIAL columns
 - nonsequential numbers in 2-217
 - resetting counter 2-58
 - use with hash fragmentation 2-217
- SERIAL data type
 - inserting values 2-436
 - invalid default 2-195
 - last inserted value 4-105
 - length 4-27
 - resetting counter 2-58, 2-436
 - value range 4-27
- Serial key 2-599
- SERIAL8 data type
 - inserting values 2-436
 - invalid default 2-195
 - last inserted value 4-105
 - value range 4-27
- SERIALIZABLE keyword, in SET TRANSACTION statement 2-656

- SERVER_LOCALE environment variables 4-25
- SERVER_NAME keyword, in GET DIAGNOSTICS statement 2-391
- Session
 - get the client time zone 4-105
 - set initial environment for 6-51
- Session control block 4-103
- Session ID 4-103
- SESSION keyword, in SET SESSION AUTHORIZATION statement 2-650
- Session password 2-611
- SET ALL_MUTABLES statement 2-581
- SET AUTOFREE statement 2-582
- SET CHUNK OFFLINE command 6-6
- SET COLLATION statement 2-584
- SET columns, generating values for 4-83
- SET CONNECTION statement 2-587
- SET CONSTRAINTS statement 2-590
- SET data type
 - defined 4-84
 - deleting elements from 2-309
 - unloading 2-684
 - updating elements 2-703
- SET Database Object Mode statement
 - syntax 2-591
 - with
 - CREATE TRIGGER statement 2-268
- SET DATASKIP OFF command 6-6
- SET DATASKIP ON command 6-6
- SET DATASKIP statement
 - syntax 2-598
- SET DEBUG FILE statement
 - syntax 2-599
 - with TRACE statement 3-47
- SET Default Table Space statement 2-601
- SET Default Table Type statement 2-602
- SET DEFERRED_PREPARE statement 2-603
- SET DESCRIPTOR statement 2-605
- SET ENCRYPTION PASSWORD statement
 - audit-event mnemonic 4-109
 - syntax 2-611
- SET ENVIRONMENT statement 2-615
- SET EXPLAIN statement 2-627
- SET INDEX COMPRESSION command 6-6
- SET INDEX statement 2-630
- SET INDEXES statement 2-631
- SET ISOLATION statement
 - isolation levels defined 2-658
 - similarities to SET TRANSACTION statement 2-657
- SET keyword 4-178
 - in ALTER SECURITY LABEL COMPONENT statement 2-36
 - in CREATE SECURITY LABEL COMPONENT statement 2-177
 - in DEFINE statement 3-13
 - in Expression segment 4-83
 - in MERGE statement 2-452
 - in ON EXCEPTION statement 3-38
 - in UPDATE statement 2-693
- SET LOCK MODE statement 2-637
- SET LOG statement 2-639
- SET ONCONFIG command 6-6
- SET OPTIMIZATION statement
 - ALL_ROWS option 2-642
 - FIRST_ROWS option 2-642
 - HIGH option 2-641
 - LOW option 2-641
- SET OPTIMIZATION statement (*continued*)
 - syntax 2-641
- SET PDQPRIORITY statement 2-643
- SET PLOAD FILE statement 2-645
- SET Residency statement 2-646
- SET ROLE statement 2-647
- SET SBSPACE ACESSTIME OFF command 6-6
- SET SBSPACE ACESSTIME ON command 6-6
- SET SBSPACE AVG_LO_SIZE command 6-6
- SET SBSPACE LOGGING OFF command 6-6
- SET SBSPACE LOGGING ON command 6-6
- SET SCHEDULE LEVEL statement 2-649
- SET SESSION AUTHORIZATION statement 2-650
- SET SHUTDOWN command 6-6
- SET SQL TRACING command 6-6
- SET SQL TRACING DATABASE ADD command 6-6
- SET SQL TRACING DATABASE CLEAR command 6-6
- SET SQL TRACING DATABASE LIST command 6-6
- SET SQL TRACING DATABASE REMOVE command 6-6
- SET SQL TRACING OFF command 6-6
- SET SQL TRACING ON command 6-6
- SET SQL TRACING RESUME command 6-6
- SET SQL TRACING SESSION command 6-6
- SET SQL TRACING SUSPEND command 6-6
- SET SQL TRACING USER ADD command 6-6
- SET SQL TRACING USER CLEAR command 6-6
- SET SQL TRACING USER LIST command 6-6
- SET SQL TRACING USER REMOVE command 6-6
- SET STATEMENT CACHE statement 2-652
- SET TABLE statement 2-655
- SET Transaction Mode statement 2-661
- SET TRANSACTION statement 2-656
 - default database levels 2-659
 - effects of isolation 2-660
 - similarities to SET ISOLATION statement 2-657
- SET TRIGGERS statement 2-663
- Set-column level encryption 4-108
- setenv utility 2-86
- setnet32 utility 2-85
- SETSESSIONAUTH keyword
 - in GRANT statement 2-421
 - in REVOKE statement 2-515
- SETSESSIONAUTH privilege 2-650
- SETUDTEXTNAME procedure 6-61
- setUDTextName() procedure 2-120
- setUDTEExtName() procedure 2-161
- Shadow columns 2-46, 2-47
- SHARE keyword, in LOCK TABLE statement 2-449
- Shared library functions 2-411, 5-20
- Shared lock mode 2-449
- Shared memory
 - decision-support parameters 6-33
 - index fragments 2-630
 - table fragments 2-655
- Shared-object files 2-120, 5-20
- Shell script 3-46
- SHM VP class) 6-37
- SHMADD configuration parameter 6-30
- shortcut keys
 - keyboard C-1
- Side-effect index 2-152
- Signatures 5-19
- Simple assignment 3-32
- Simple join 2-553
- Simple large objects 4-29
 - loading 2-445, 2-446
 - unloading 2-685

- Simple table expression 2-546
- Simple view 2-272
- SIN function 4-139, 4-140
- Single quotation marks 4-189
 - literal in a quoted string 4-192
- Single-byte characters 4-24
- Single-threaded application 2-588
- Singleton SELECT statement 2-534, 2-542, 2-695
- SITENAME function
 - in ALTER TABLE statement 2-49
 - in Condition segment 4-10
 - in CREATE TABLE statement 2-195
 - in DEFINE statement 3-10
- SIZE keyword
 - in ALTER TABLE statement 2-66, 2-67
 - in CREATE EXTERNAL TABLE statement 2-109
 - in CREATE TABLE statement 2-222, 2-224
 - in SELECT statement 2-574
 - in UPDATE STATISTICS statement 2-711
- SKIP keyword in SELECT statement 2-532
- Slash and asterisk (/ * */) comment indicator 1-4, 2-470, 5-36
- Slot number 4-64
- SLV. 4-164
- SMALLFLOAT data type 4-29
 - literal values 4-186
 - systems not supporting 2-282
- SMALLINT data type, literal values 4-185
- Smart large objects 4-30
 - accessing column data 4-64
 - copying to a file 4-125
 - copying to a smart large object 4-126
 - creating from a file 4-122, 4-123, 6-54
 - data integrity 2-223
 - expressions with 4-64
 - extent size 2-223
 - functions for copying 4-122
 - generating filename for 4-126
 - handle values 4-64
 - loading values 2-445, 2-447
 - logging 2-223
 - storing 2-62, 2-222
 - unloading 2-685, 2-686
- SMI. 4-103
- SOC VP class 6-37
- Socket connections 6-37
- SOME keyword
 - beginning a subquery 2-564
 - in Condition segment 4-18
- Sort-merge join 2-621
- Sorting
 - in a combined query 2-579
 - in SELECT 2-568
- SOURCEID field 2-381
 - in SET DESCRIPTOR statement 2-606
- SOURCETYPE field 2-381
 - in SET DESCRIPTOR statement 2-606
- Spatial data 2-343, 2-698
- SPECIFIC FUNCTION keywords
 - in ALTER FUNCTION statement 2-28
 - in GRANT statement 2-405
 - in REVOKE statement 2-501
 - in UPDATE STATISTICS statement 2-714
- SPECIFIC keyword
 - in ALTER FUNCTION statement 2-28
 - in ALTER PROCEDURE statement 2-32
 - in ALTER ROUTINE statement 2-33
 - in CREATE FUNCTION statement 2-113
- SPECIFIC keyword (*continued*)
 - in CREATE PROCEDURE statement 2-153
 - in DROP FUNCTION statement 2-328
 - in DROP PROCEDURE statement 2-332
 - in DROP ROUTINE statement 2-335
 - in GRANT statement 2-405
 - in REVOKE statement 2-501
 - in UPDATE STATISTICS statement 2-714
- Specific Name segment 5-74
- SPECIFIC PROCEDURE keywords
 - in ALTER PROCEDURE statement 2-32
 - in GRANT statement 2-405
 - in REVOKE statement 2-501
 - in UPDATE STATISTICS statement 2-714
- SPECIFIC ROUTINE keywords
 - in ALTER ROUTINE statement 2-33
 - in GRANT statement 2-405
 - in REVOKE statement 2-501
 - in UPDATE STATISTICS statement 2-714
- SPL function
 - CREATE FUNCTION 2-119
 - cursors 3-23
 - dropping 2-327
 - dynamic routine-name specification 2-360
 - executing 2-356, 2-471
 - optimization 2-119
 - registering 2-119
 - registering from inside an external routine 2-123
- SPL keyword
 - in GRANT statement 2-406
 - in REVOKE statement 2-503
- SPL procedure
 - creating with CREATE PROCEDURE 2-160
 - dynamic routine-name specification 2-367
 - executing 2-471
 - optimization 2-160, 2-706
 - registering with CREATE PROCEDURE 2-160
 - sysdbclose() 6-51
 - sysdbopen() 6-51
- SPL routines
 - as triggered action 2-259
 - BYTE and TEXT data types 3-16
 - comment indicators 1-4
 - debugging 3-47
 - defined 3-1
 - definition 2-155
 - dropping with DROP PROCEDURE 2-333
 - executing operating-system commands 3-44
 - handling multiple rows 3-43
 - header 3-8
 - in SELECT statement 2-540
 - limits on parameters 5-68
 - output file for TRACE statement 2-599
 - ownership of created objects 2-162
 - preparing 2-471
 - receiving data from SELECT 2-541
 - reoptimizing 2-706
 - restrictions when used with DML statements 5-78
 - sending mail 3-45
 - setting environment variables 3-46
 - simulating errors 3-41
 - SQL-statement restrictions 5-76
- SPL statements, defined 3-1
- sqexplain.out file 2-525, 2-557, 2-626, 2-714, 5-43
- SQL
 - comments 1-3
 - compliance of statements with ANSI standard 1-10

- SQL (*continued*)
 - reserved words A-1, B-1
 - statement types 1-5
- SQL Admin API functions
 - ADMIN 6-2
 - TASK 6-2
- SQL Communications Area 2-78
 - result after CLOSE 2-78
 - result after DATABASE 2-282
 - result after DATASKIP event 2-598
 - result after DELETE 2-310, 4-102
 - result after DESCRIBE 2-312
 - result after DESCRIBE INPUT 2-316
 - result after EXECUTE 2-353, 2-476
 - result after FETCH 2-375
 - result after FLUSH 2-377
 - result after INSERT 4-102
 - result after OPEN 2-462
 - result after PUT 2-484
 - result after SELECT 2-543, 4-102
 - result after UPDATE 4-102
 - sqlca.sqlerrd1 4-102
 - sqlca.sqlerrd2 4-102
 - sysibm.SQLCAMessage function 6-64
- SQL DESCRIPTOR keywords
 - in DESCRIBE INPUT statement 2-315
 - in DESCRIBE statement 2-311
 - in EXECUTE statement 2-350, 2-353
 - in FETCH statement 2-368
 - in OPEN statement 2-460
 - in PUT statement 2-479
- SQL Function. 4-90
- SQL keyword
 - in DESCRIBE INPUT statement 2-317
 - in DESCRIBE statement 2-313
 - in EXECUTE statement 2-355
 - in OPEN statement 2-460
- SQL statement cache
 - disabling 2-654
 - enabling 2-653
 - prepared statements 2-655
 - qualifying criteria 2-654
- SQL statements
 - restrictions within SPL routines 5-76
- SQL_INFX_ATTR_DELIMIDENT connection attribute 5-25
- SQL-99 standard 1-4
- SQLCA. 2-282, 4-105, 4-109
- SQLCAMessage function 6-64
- SQLCODE function 4-98
- SQLCODE variable 2-78, 2-310, 2-375, 2-378, 2-462, 2-720
- sqld value 2-373
- sqlda structure
 - in DESCRIBE 2-311, 2-313
 - in DESCRIBE INPUT 2-316
 - in EXECUTE 2-351
 - in EXECUTE ... INTO 2-352
 - in FETCH 2-373
 - in OPEN 2-350, 2-354, 2-460, 2-479, 2-480
 - in OPEN...USING DESCRIPTOR 2-465
- SQLERRD array
 - last inserted BIGSERIAL value 4-105
 - last inserted SERIAL8 value 4-105
 - number of inserted rows 2-378
 - value of inserted SERIAL8 value 4-109
- SQLERROR keyword, in WHENEVER statement 2-718
- sqlhosts file 2-85
- SQLI client-server communication protocol 2-387
- SQLJ driver built-in procedures
 - ALTER_JAVA_PATH 6-60
 - JVPCONTROL 6-57
 - REMOVE_JAR 6-59
 - REPLACE_JAR 6-59
 - SETUDTEXTNAME 6-61
 - SQLJ.INSTALL_JAR 6-58
 - UNSETUDTEXTNAME 6-62
- sqlj schema 6-58
- SQLJ.ALTER_JAVA_PATH procedure 6-60
- SQLJ.INSTALL_JAR procedure 5-20, 5-73, 6-58
- SQLJ.REMOVE_JAR procedure 6-59
- SQLJ.REPLACE_JAR procedure 6-59
- SQLNOTFOUND
 - error conditions with EXECUTE statement 2-353
 - with INSERT statement 2-441
- SQLNOTFOUND value 2-575
- SQLSTATE
 - after FETCH 2-375
 - after FLUSH 2-378
 - after REVOKE 2-500
 - list of codes 2-387
 - not found condition 2-310, 2-375, 2-721
 - runtime errors 2-720
 - warnings 2-720
- sqlstypes.h header file 2-312, 2-317
- sqltypes.h file 2-385
- sqltypes.h header file 2-608
- SQLUNKNOWN data type 2-317
- sqlvar structures 2-373
- SQLWARNING keyword, in WHENEVER statement 2-718
- sqlxtype.h header file 2-609
- SQRT function 4-89, 4-91
- srvsendrecv data type 2-147
- STABILITY keyword, in SET ISOLATION statement 2-635
- STACK keyword
 - Routine Modifier segment 5-60
- STACKSIZE configuration parameter 5-66
- STANDARD keyword
 - in ALTER TABLE statement 2-69
 - in CREATE TABLE statement 2-191, 2-193
 - in SELECT statement 2-575
 - in SET Default Table Type statement 2-602
- START keyword
 - in CREATE SEQUENCE statement 2-186
- START MIRRORING command 6-6
- START VIOLATIONS TABLE statement 2-664
- Starting the database server online 6-35
- STAT data type 2-536, 4-22
- statcollect() function 2-708
- Statement block segment 5-74
- Statement identifier
 - cursor for 2-297
 - defined 2-469
 - in DECLARE 2-286
 - in FREE 2-379, 2-380
 - in PREPARE 2-469
 - releasing 2-470
- STATEMENT keyword, in SET STATEMENT CACHE statement 2-652
- Statement Local Variables
 - data type of 4-164
 - declaration 4-162
 - expression 4-164
 - name space of 4-164
 - OUT parameter 5-69
 - precedence of 4-164

- Statement Local Variables *(continued)*
 - scope of 4-164
 - using 4-164
- Statements
 - SQL
 - ANSI-compliant 1-11
 - entering 1-1
 - extensions to ANSI standard 1-12
- Statements, SQL
 - valid only in ESQL/C 4-56
- STATIC keyword
 - in ALTER TABLE statement 2-69
 - in CREATE TABLE statement 2-191, 2-193
 - in SELECT statement 2-575
 - in SET Default Table Type statement 2-602
- STATISTICS keyword
 - UPDATE STATISTICS statement 2-705
- statprint() function 2-708
- STATUS keyword, in INFO statement 2-428
- Status, displaying with INFO statement 2-429
- STDEV function 4-165, 4-173
- STEP keyword 3-19
 - audit-event mnemonic 4-109
- STMT_CACHE configuration parameter 2-652, 6-34
- STMT_CACHE environment variable 2-653
- STMT_CACHE_HITS configuration parameter 2-654, 6-40
- STMT_CACHE_NOLIMIT configuration parameter 2-654, 6-40
- STMT_CACHE_SIZE configuration parameter 2-654
- STOP keyword, in WHENEVER statement 2-718
- STOP MIRRORING command 6-6
- STOP VIOLATIONS TABLE statement 2-676
- STORAGE keyword
 - in TRUNCATE statement 2-678
- Storage options, CREATE Temporary TABLE 2-236
- Stored Procedure Languages 3-1
- Storing smart large objects 2-222
- STR VP class 6-37
- STRATEGIES keyword, in CREATE OPCLASS statement 2-151
- Strategy functions 2-151
- Stream pipe connections 6-37
- String-manipulation functions 4-141
- STYLE keyword
 - External Routine Reference segment 5-19
- SUBCLASS_ORIGIN keyword, in GET DIAGNOSTICS statement 2-391
- Subdiagram reference 4-1
- Subordinate table 2-553
- Subquery 4-3
 - beginning with ALL, ANY, SOME keywords 2-564
 - beginning with EXISTS keyword 2-563, 4-18
 - beginning with IN keyword 2-563, 4-17
 - correlated 4-16
 - defined 2-561
 - estimated cost 2-627
 - in a table hierarchy 2-696
 - in Condition segment 4-15
 - in WHERE clause of UPDATE statement 2-700
 - no FIRST keyword 2-533, 2-534
 - updating a column 2-694
 - updating multiple columns 2-695
 - with DISTINCT keyword 2-535
 - with UNION or UNION ALL 2-580
- Subscripting character columns 2-570, 4-63
- Subsecond precision 4-76
- SUBSTR function 4-152
- Substring
 - in ORDER BY clause of SELECT 2-570
 - operator in column expression 4-63
- SUBSTRING function 4-150
- Subtable
 - inherited properties 2-228
 - restrictions 2-171
- Subtype
 - creating 2-169
 - dropping 2-336
- SUM function 4-165, 4-172
- Supertable
 - updating 2-696
- Supertype
 - creating 2-169
 - dropping 2-337
- Support functions
 - assigning 2-149, 2-437, 2-448, 2-698
 - comparing 2-149
 - defined 2-152
 - defining 2-147
 - destroy 2-149, 2-309, 2-343
 - export 2-148
 - exportbinary 2-149
 - import 2-148
 - importbinary 2-148
 - input 2-148
 - lohandles 2-149
 - output 2-148
 - receive 2-148
 - send 2-148
 - specifying in CREATE OPCLASS 2-152
- SUPPORT keyword, in CREATE OPCLASS statement 2-149
- Synonym
 - chaining 2-190
 - creating 2-188
 - difference from alias 2-188
 - dropping 2-341
 - external 2-343
- SYNONYM keyword
 - in DROP SYNONYM statement 2-341
- syntax 4-183, 4-193, 5-23
- Syntax diagrams
 - reading in a screen reader C-1
- syntax for views 2-271
- sysadmin database 6-2, 6-49
- sysaggregates system catalog table 2-92, 2-324
- sysams system catalog table 2-5, 2-91, 2-133, 2-322
 - columns 5-51
- sysblobs system catalog table 2-230
- syscasts system catalog table 2-95, 2-324
- syschecks system catalog table 2-487
- syscolattribs system catalog table 2-222
- syscolauth system catalog table 2-169
- syscolumns system catalog table 2-230, 2-337, 2-432, 2-608, 2-706
- sysconstraints system catalog table 2-64, 2-126, 2-330, 2-489
- SYSDATE function
 - as constant expression 4-71
 - in ALTER TABLE statement 2-49
 - in CREATE TABLE statement 2-195
 - in DEFINE statement 3-10
 - in INSERT statement 2-439
- SYSDATE operator
 - defined 4-76
- sysdbclose() procedure 6-52
- sysdbopen() procedure 2-582, 2-601, 2-602, 2-644, 6-51, 6-52

sysdirectives system catalog table 2-525, 5-45
 sysdistrib system catalog table 2-144, 2-457, 2-706, 2-708
 syssecpolicies system catalog table 2-182, 2-338, 2-490
 sysfragauth system catalog table 2-165, 2-424, 2-427
 sysfragments system catalog table 2-9, 2-15, 2-230, 2-489, 4-102
 sysindexes system catalog table 2-203, 2-330, 2-489, 2-706, 2-709
 sysinherits system catalog table 2-230, 2-337
 sysmaster database 2-230, 2-456, 4-103, 4-169
 sysobjstate system catalog table 2-489, 2-591
 sysprocauth system catalog table 2-119, 2-160, 2-163, 2-165, 2-230
 sysprocbody system catalog table 2-119, 2-159, 4-109
 sysprocedures system catalog table 2-119, 2-160, 2-163, 2-651, 6-58
 sysprocplan system catalog table 2-119, 2-160, 2-163, 2-716
 sysroleauth system catalog table 2-165
 SYSSBSPACENAME onconfig parameter 2-708
 sysseclabelauth system catalog table 2-421, 2-515
 sysseclabelcomponentelements system catalog table 2-36, 2-178
 sysseclabelcomponents system catalog table 2-36, 2-178, 2-338, 2-490
 sysseclabels system catalog table 2-175, 2-338, 2-490
 syssecpolicycomponentrules system catalog table 2-182
 syssequences system catalog table 2-42, 2-185, 2-340
 sysssynonyms system catalog table 2-341
 sysssyntable system catalog table 2-341
 systabauth system catalog table 2-165, 2-169, 2-424, 2-458
 systables system catalog table 2-230, 2-330, 2-500, 2-706, 2-709, 4-101
 System catalog tables
 owner informix 5-17
 sysaggregates 2-92
 sysams 2-5, 2-90, 2-91
 syscasts 2-95, 2-324
 syschecks 2-487
 syscolauth 2-165
 syscolumns 2-337, 2-706
 sysconstraints 2-126, 2-203, 2-330, 2-489
 sysdepend 2-346
 sysdirectives 2-523, 5-45
 sysdistrib 2-144, 2-706
 syssecpolicies 2-182, 2-338, 2-490
 sysfragauth 2-165, 2-424, 2-427
 sysfragments 2-9, 2-230, 2-489
 sysindexes 2-203, 2-489, 2-706
 sysinherits 2-337
 sysobjstate 2-489, 2-591
 sysprocauth 2-119, 2-160, 2-163, 2-165
 sysprocbody 2-119, 2-160, 2-163, 4-109
 sysprocedures 2-119, 2-160, 2-163, 2-651
 sysprocplan 2-119, 2-160, 2-163, 2-716
 sysroleauth 2-165
 sysseclabelauth 2-421, 2-515
 sysseclabelcomponentelements 2-36, 2-178
 sysseclabelcomponents 2-36, 2-178, 2-338, 2-490
 sysseclabels 2-175, 2-338, 2-490
 syssecpolicycomponentrules 2-182
 syssequences 2-42, 2-43, 2-185, 2-340
 sysssynonyms 2-341
 sysssyntable 2-341
 systabauth 2-165, 2-169, 2-413, 2-424
 systables 2-330, 2-337, 2-500, 2-706
 systriggers 2-243, 2-344
 sysusers 2-165, 2-170

System catalog tables (*continued*)
 sysviews 2-275, 2-487, 2-493
 sysviolations 2-664, 2-677
 sysxdatasources 2-279, 2-347
 sysxasourcetypes 2-280, 2-348
 sysxdttypeauth 2-146, 2-165, 2-169, 2-230, 2-404
 sysxdtypes 2-146, 2-169, 2-337, 2-385, 2-404
 System catalogs
 creating 2-98
 dropping tables 2-343
 System clock 4-70, 4-78
 System constants 4-70
 System index 2-138, 5-37
 System name
 database qualifier 5-16
 SYSTEM statement 3-44
 System-descriptor area
 assigning values 2-605
 creating 2-3
 deallocating 2-284
 item descriptors 2-3
 OPEN using 2-354, 2-464, 2-482
 resizing 2-606
 use with EXECUTE statement 2-355
 with ALLOCATE DESCRIPTOR 2-3
 with DESCRIBE 2-313
 with DESCRIBE INPUT 2-318
 with EXECUTE ... INTO 2-351
 System-diagnostics area 2-375
 System-monitoring interface tables 4-103
 System-Monitoring Interface tables 2-681, 4-169
 systriggers system catalog table 2-243, 2-344
 sysuolauth system catalog table 2-458
 sysusers system catalog table 2-165, 2-170, 2-230, 2-429
 sysutils database 2-456
 sysviews system catalog table 2-275, 2-487, 2-493
 sysviolations system catalog table 2-664, 2-677
 sysxdatasources system catalog table 2-279, 2-347
 sysxasourcetypes system catalog table 2-280, 2-348
 sysxdttypeauth system catalog table 2-101, 2-146, 2-165, 2-230, 2-404
 sysxdtypes system catalog table 2-101, 2-146, 2-170, 2-230, 2-337, 2-345, 2-385, 2-404, 2-610
 DESCRIBE and GET DESCRIPTOR with 2-384

T

T keyword in SQL Admin API functions 6-2
 tabid column 2-19
 Table
 adding a constraint 2-63, 2-64
 adding a constraint to a column with data 2-61
 alias in SELECT 2-544
 build 2-628
 child 2-200
 creating 2-191
 creating a synonym for 2-187
 default privileges 2-424
 defining fragmentation strategy 2-212
 deleting all rows 2-677, 2-682
 derived 2-546, 4-5, 5-6
 diagnostics 2-242, 2-672
 diagnostics table 2-676
 dominant 2-553
 dropping 2-327, 2-342
 dropping a synonym 2-341
 dropping a system table 2-343

- Table (continued)
 - duplicate 2-103, 2-326
 - external 2-104, 2-242, 2-577, 5-37
 - fragmented 2-8
 - hash 2-628
 - inheritance 2-243, 2-250
 - inheritance hierarchy 2-227, 2-691
 - isolating 2-211
 - joins in Condition segment 2-564
 - loading data with the LOAD statement 2-443
 - locking
 - with ALTER INDEX 2-31
 - with LOCK TABLE 2-449
 - logging 2-232
 - memory-resident 2-646, 2-655
 - moving 2-455
 - nonfragmented 2-20
 - operational 2-242
 - parent 2-200
 - permanent 2-238
 - privilege 2-8
 - privileges
 - granting 2-399
 - moving 2-458
 - privileges on 2-229, 2-497
 - protected 2-17, 2-440
 - qualifiers 5-18
 - raw 2-242, 2-692
 - read-only 2-102
 - renaming 2-456, 2-457, 2-492
 - residency status 2-646, 2-656
 - scratch 2-232, 2-242, 2-578
 - static 2-242, 2-658
 - subordinate 2-553
 - system catalog 2-242
 - target 2-664, 2-668, 2-676
 - temporary 2-232, 2-242, 2-575, 2-706
 - typed 2-168, 2-228, 2-707
 - unlocking 2-688
 - untyped 2-228
 - updating statistics 2-706
 - violations 2-242, 2-667
 - violations table 2-676
 - virtual 5-5
 - waiting for a locked table 2-638
- Table expression 2-546, 4-5, 5-6
- Table format, in SET Database Object Mode statement 2-593
- TABLE keyword 2-705, 5-4
 - in ALTER FRAGMENT statement 2-7
 - in ALTER TABLE statement 2-68
 - in CREATE DUPLICATE statement 2-102
 - in CREATE EXTERNAL TABLE statement 2-105
 - in CREATE TABLE statement 2-225
 - in CREATE TEMP TABLE statement 2-230
 - in CREATE Temporary TABLE statement 2-175
 - in Data Type segment 4-29
 - in DROP DUPLICATE statement 2-326
 - in DROP TABLE statement 2-342
 - in LOCK TABLE statement 2-449
 - in MOVE TABLE statement 2-455
 - in RENAME TABLE statement 2-492
 - in SELECT statement 2-551
 - in SET Residency statement 2-646
 - in SET TABLE statement 2-656
 - in START VIOLATIONS TABLE statement 2-664
 - in STOP VIOLATIONS TABLE statement 2-676
 - in TRUNCATE (XPS) statement 2-682
- TABLE keyword (continued)
 - in TRUNCATE statement 2-678
 - in UNLOCK TABLE statement 2-688
- Table name
 - INFO statement 2-429
- TABLE_SPACE keyword
 - in SET Default Table Space statement 2-601
- TABLE_TYPE keyword
 - in SET Default Table Type statement 2-602
- Table-level locking, in CREATE TABLE 2-226
- Table-level privileges 2-411
- TABLES keyword, in INFO statement 2-428
- TAN function 4-139, 4-140
- Target table
 - relationship to diagnostics table 2-668, 2-676
 - relationship to violations table 2-668, 2-676
- TASK function 6-2
- TCP/IP connections 2-310, 2-439, 2-538, 2-699
- TDES (Triple Data Encryption Standard) 4-115
- TEMP keyword
 - in CREATE TEMP TABLE statement 2-230
 - in CREATE Temporary TABLE statement 2-231
 - in SELECT statement 2-574
 - in SET Default Table Space statement 2-601
 - in SET Default Table Type statement 2-602
- TEMP_TAB_EXT_SIZE keyword, in SET ENVIRONMENT statement 2-619
- TEMP_TAB_NEXT_SIZE keyword, in SET ENVIRONMENT statement 2-619
- Temporary dbspace 2-237
- Temporary tables
 - and fragmentation 2-237
 - constraints allowed 2-234
 - creating 2-231
 - creating constraints for 2-234
 - default location 2-601
 - default type 2-602
 - defining columns 2-233
 - differences from permanent tables 2-238
 - duration 2-238
 - extent size 2-619
 - implicit creation 2-602
 - INFO statement restrictions 2-238
 - system-generated 2-619
 - updating statistics 2-706
 - when deleted 2-238
 - where stored 2-237
- TEMPTAB_NOLOG configuration parameter 2-232
- TEST keyword, in SAVE EXTERNAL DIRECTIVES statement 2-523
- TEXT column, modifying 2-58
- TEXT data type
 - declaration syntax 4-30
 - default value 2-196
 - loading 2-445
 - SPL routines 3-9, 3-16
 - storage location 4-30
 - unloading 2-685
 - with SET DESCRIPTOR 2-611
- TEXT keyword 5-55
 - in CREATE EXTERNAL TABLE statement 2-105
 - in Data Type segment 4-29
 - in DEFINE statement 3-8
- THEN keyword
 - in CASE statement 3-5
 - in Expression segment 4-66, 4-67
 - in IF statement 3-28

- THEN keyword (*continued*)
 - in MERGE statement 2-452
- Thread-safe application
 - defined 2-321, 2-588, 2-590
- THREADS keyword, in EXECUTE FUNCTION statement 6-56
- Time 4-128
- Time and date, getting current 4-75
- Time data types 4-31
- Time function 4-128
 - restrictions with GROUP BY 2-567
 - use in SELECT 2-539
- TIME keyword
 - in ALTER TABLE statement 2-62
 - in CREATE TABLE statement 2-222
- Time unit 4-181
 - INTERVAL data types 4-177
- Time zone 2-617
 - of CURRENT operator 4-77
 - of DBINFO function 4-105
 - of TODAY operator 4-75
- times() operator function 4-51
- TLI VP class 6-37
- TMPSPACE_LIMIT keyword
 - in SET ENVIRONMENT statement 2-619
- TO CLUSTER keywords, in ALTER INDEX statement 2-30
- TO keyword 3-19
 - in ALTER FRAGMENT statement 2-26
 - in ALTER INDEX statement 2-30
 - in CONNECT statement 2-83
 - in CREATE SECURITY POLICY statement 2-182
 - in DATETIME field qualifier 4-39
 - in EXTEND function 4-128
 - in GRANT FRAGMENT statement 2-423
 - in GRANT statement 2-414, 2-416, 2-418, 2-421
 - in INTERVAL Field Qualifier segment 4-176
 - in MOVE TABLE statement 2-455
 - in OUTPUT statement 2-467
 - in RENAME COLUMN statement 2-487
 - in RENAME DATABASE statement 2-488
 - in RENAME INDEX statement 2-489
 - in RENAME SECURITY statement 2-490
 - in RENAME SEQUENCE statement 2-492
 - in RENAME TABLE statement 2-492
 - in ROLLBACK WORK statement 2-521
 - in SET ALL_MUTABLES statement 2-581
 - in SET DEBUG FILE statement 2-600
 - in SET Default Table Space statement 2-601
 - in SET Default Table Type statement 2-602
 - in SET EXPLAIN statement 2-624
 - in SET ISOLATION statement 2-631
 - in SET LOCK MODE statement 2-638
 - in SET PLOAD FILE statement 2-646
 - in SET SESSION AUTHORIZATION statement 2-650
 - in UNLOAD statement 2-684
 - in WHENEVER statement 2-721
- TO_CHAR function 4-128, 4-135
- TO_DATE function 4-128, 4-138
- TO_NUMBER function 4-139
- TODAY function
 - as expression 4-71
 - in ALTER TABLE statement 2-49
 - in Condition segment 4-10
 - in CREATE TABLE statement 2-195
 - in DEFINE statement 3-10
 - in INSERT 2-434, 2-439
- TRACE statement 3-46
 - TRACE statement (*continued*)
 - output file 2-600
- TRAILING keyword, in TRIM expressions 4-147
- TRANSACTION keyword
 - in CONNECT statement 2-83
 - in SET TRANSACTION statement 2-657
- Transaction mode
 - constraints 2-661
- Transactions
 - access mode 2-660
 - example 2-295
 - partial rollback 2-521
 - read-only 2-660
 - rollback 2-520
 - statements that initiate 2-652
 - using cursors in 2-300
 - without error handling 2-75
- TREE keyword
 - in ALTER SECURITY LABEL COMPONENT statement 2-36
 - in CREATE SECURITY LABEL COMPONENT statement 2-177
- Trigger
 - inherited 2-243, 2-250
 - overriding 2-250
- Trigger action 2-239
- Trigger event 2-239
 - DELETE 2-246, 2-270
 - INSERT 2-255, 2-273
 - privileges on 2-245
 - SELECT 2-247, 2-272
 - UPDATE 2-246, 2-270
- Trigger functions 2-116
- TRIGGER keyword
 - in DROP TRIGGER statement 2-345
 - in EXECUTE FUNCTION statement 2-356
 - in EXECUTE PROCEDURE statement 2-364
- Trigger procedures 2-157
- Trigger routines 2-116, 2-157
- Trigger UDR 2-116, 2-157
- Trigger-type Boolean operator 4-12
- Triggered action 2-271
 - action statements 2-259
 - cascading 2-253
 - correlation names 2-265
 - effect of cursors 2-245
 - for multiple triggers 2-253, 2-272
 - list of actions 2-258
 - WHEN condition 2-258
- Triggering statement
 - consistent results 2-259
 - performance 2-245
 - UPDATE 2-247
- Triggering view 2-271
- Triggers
 - affected by dropping a column from table 2-56
 - affected by modifying a column 2-61
 - enabling or disabling 2-595
 - overriding 2-269
- TRIGGERS keyword, in SET Database Object Mode statement 2-593, 2-663
- Trigonometric function
 - ACOS function 4-140
 - ASIN function 4-141
 - ATAN function 4-141
 - ATAN2 function 4-141
 - COS function 4-140

Trigonometric function (*continued*)
 SIN function 4-140
 TAN function 4-140
 TRIM function 4-147
 Triple Data Encryption Standard (TDES or DES3) 4-115
 TRUNC function 4-89
 TRUNCATE statement 2-677, 2-682
 Trusted Facility feature 4-109
 Two-phase commit operations 2-279
 Two-phase commit protocol, 6-44
 TXTIMEOUT configuration parameter 6-44
 TYPE field 2-381
 changing from BYTE or TEXT 2-611
 in SET DESCRIPTOR statement 2-606
 setting in X/Open programs 2-609
 with DESCRIBE INPUT statement 2-318
 with DESCRIBE statement 2-313
 with X/Open programs 2-383
 Type hierarchy 2-170
 TYPE keyword
 in ALTER TABLE statement 2-69, 2-72
 in CREATE DISTINCT TYPE statement 2-100
 in CREATE ROW TYPE statement 2-168
 in CREATE TABLE statement 2-227
 in CREATE VIEW statement 2-274
 in CREATE XADATASOURCE TYPE statement 2-280
 in DROP ROW TYPE statement 2-336
 in DROP TYPE statement 2-345
 in DROP XADATASOURCE TYPE statement 2-348
 in GRANT statement 2-403
 in REVOKE statement 2-500
 Typed collection variable 2-2, 3-14
 Typed table
 ADD TYPE clause 2-72
 altering 2-73
 altering serial columns 2-59, 2-171
 inheritance 2-228
 NOT NULL constraint 2-171
 Typed view 2-275

U

UDR definition routines
 IFX_REPLACE_MODULE 6-54
 JVPCONTROL 6-56
 UDR Definition Routines
 IFX_UNLOAD_MODULE 6-55
 Unary minus operator (-) 4-184, 4-185
 Unary plus operator (+) 4-184, 4-185
 Unbuffered logging 2-640
 UNCOMMITTED keyword
 in SET TRANSACTION statement 2-658
 Uncommitted row 2-633
 Uncorrelated subquery 2-546
 UNDEFINED parameter value 5-2
 undelimited identifiers 5-23
 UNDER keyword
 in ALTER SECURITY LABEL COMPONENT
 statement 2-36
 in CREATE ROW TYPE statement 2-168
 in CREATE SECURITY LABEL COMPONENT
 statement 2-177
 in CREATE TABLE statement 2-227
 in GRANT statement 2-403
 in REVOKE statement 2-497, 2-500
 UNDER ON TYPE keywords
 in GRANT statement 2-403

UNDER ON TYPE keywords (*continued*)
 in REVOKE statement 2-500
 Under privilege 2-401, 2-404, 2-497, 2-500
 Underscore (_) 5-23
 as wildcard 4-13
 in SQL identifiers 5-21
 Unicode 2-196, 2-586, 4-24
 UNION operator 2-528
 in collection subquery 4-5
 in SELECT statement 2-578
 OUT parameter and 4-164
 restrictions on use 2-441, 2-578
 UNION subquery 2-579
 Union view 2-276
 Unique constraint
 dropping 2-66
 moving 2-457
 rules of use 2-198
 UNIQUE keyword
 in ALTER TABLE statement 2-50, 2-64
 in CREATE INDEX statement 2-125, 2-141
 in CREATE TABLE statement 2-197, 2-205
 in CREATE Temporary TABLE statement 2-234
 in Expression segment 4-165, 4-175
 in SAVEPOINT statement 2-526
 in SELECT 2-535
 in subquery 4-17
 Units of time, INTERVAL values 4-184
 UNITS operator 4-71
 Universal Time (UT) 4-106
 UNIX operating system
 chmod utility 6-59
 home directory 2-119
 shell script 3-46
 UNKNOWN truth values 4-20
 UNLOAD statement 2-683, 2-684
 UNLOAD TO file 2-684
 UNLOCK TABLE statement 2-688
 Unnamed row data types
 field definition 4-35
 unloading 2-685, 2-687
 updating fields 2-703
 UNSETUDTEXTNAME procedure 6-62
 Untyped collection variable 2-2, 2-283, 2-550, 4-36, 5-7, 5-11
 Untyped row variable 2-4
 Untyped view 2-275
 Updatable view 2-278, 2-308
 Update cursor
 locking considerations 2-292
 opening 2-460
 restricted statements 2-296
 use in DELETE 2-307
 use in UPDATE 2-702
 UPDATE keyword 2-424, 2-518, 2-528, 2-705
 in CREATE TRIGGER statement 2-246
 in DECLARE statement 2-286
 in GRANT statement 2-400
 in MERGE statement 2-452
 in REVOKE FRAGMENT statement 2-518
 in REVOKE statement 2-497
 in SELECT statement 2-572
 in SET ISOLATION statement 2-636
 Update locks 2-637
 Update privilege 2-400, 2-497
 Update privilege, with a view 2-691
 UPDATE statement 2-689

- UPDATE statements
 - and triggers 2-259
 - collection variables 5-10
 - cursor with 2-290
 - distributed 2-438, 2-698
 - OUT parameter and SLVs 4-165
 - SET clause 2-693
 - single-column SET clause 2-693
 - smart large objects with 4-64
 - update triggers 2-244
 - updating through a view 2-691
 - with FETCH 2-373
 - with SELECT...FOR UPDATE 2-572
- UPDATE STATISTICS statement 2-705
 - dropping data distributions 2-709
 - examining index pages 2-708
 - specifying distributions only 2-713
 - upgrading the database server 2-717
- Update trigger 2-246, 2-271
- Updating a specific table in a table hierarchy 2-691
- UPDATING operator 2-118, 2-159, 4-12
- Upgrading the database server 2-706, 2-717
- UPPER function 4-156
- Upper index filter 2-628
- Uppercase characters
 - converting from lowercase 4-156
 - in database server names 2-85, 5-18
- USAGE keyword
 - in GRANT statement 2-403, 2-406
 - in REVOKE statement 2-500, 2-503
- USAGE ON LANGUAGE keywords
 - in GRANT statement 2-406
 - in REVOKE statement 2-503
- USAGE ON TYPE keywords
 - in GRANT statement 2-403
 - in REVOKE statement 2-500
- use in Function Expressions 4-128
- USE_HASH keyword, in optimizer directives 5-40
- USE_NL keyword, in optimizer directives 5-40
- USELASTCOMMITTED configuration parameter 2-450, 2-622, 2-634, 6-23, 6-41
- USELASTCOMMITTED configuration parameters 2-658
- USELASTCOMMITTED environment option 2-634, 2-639, 2-658
- USELASTCOMMITTED keyword, in SET ENVIRONMENT statement 2-450, 2-622
- USEOSTIME configuration parameter 4-75
- USER function
 - as constant expression 4-71
 - defined 4-73
 - in ALTER TABLE statement 2-49
 - in Condition segment 4-10
 - in CREATE TABLE statement 2-195
 - in DEFINE statement 3-10
 - in INSERT statement 2-434, 2-439
 - in Literal Row segment 4-186
- User informix 2-100, 2-324, 2-411, 2-506, 2-647, 5-47
 - privileges associated with 2-399
- USER keyword
 - in CONNECT statement 2-87
 - in GRANT statement 2-414, 2-416, 2-418, 2-421
 - in REVOKE statement 2-509, 2-511, 2-513, 2-515
- User name
 - case-sensitivity 2-397, 2-652
 - using another name 2-650
- User-defined access method
 - creating 2-91
- User-defined access method (*continued*)
 - modifying 2-5
- User-defined aggregates
 - creating 2-92
 - defined 4-167
 - dropping 2-323
 - invoking 4-175
- User-defined data types 4-32
 - maximum in one row 2-194, 4-34
 - privileges 2-403, 2-404, 2-501
- User-defined function 4-161
 - arguments 5-1
 - cursor 2-359
 - inserting data with 2-441
 - iterator 5-63
 - negator 2-357, 5-63
 - noncursor 2-358
 - OUT parameter 4-164
 - selectivity 5-65
 - variant 5-20, 5-66
- User-defined routines
 - arguments 2-147, 5-1
 - defined 2-155
 - dropping with DROP ROUTINE 2-335
 - EXTEND role 2-411, 2-507
 - ill-behaved 5-62
 - in SELECT statements 2-540
 - inserting data with 2-441
 - ownership of created objects 2-121, 5-19
 - privileges 2-405, 2-501
 - REFERENCES keyword with BYTE or TEXT data type 3-16
 - reoptimization 2-714
 - RESTRICTED mode 2-651
 - return values 5-55
 - VP class 5-61
- User-defined VP class 5-61, 6-38
- USETABLENAME environment variable 2-44, 2-343, 2-679
- USING BITMAP keywords, in CREATE INDEX statement 2-124
- USING DESCRIPTOR keywords
 - in EXECUTE 2-353
 - in FETCH 2-373
 - in OPEN 2-460
 - in PUT 2-355
- USING keyword
 - in CONNECT statement 2-87
 - in CREATE EXTERNAL TABLE statement 2-104
 - in CREATE INDEX statement 2-124, 2-132
 - in CREATE TABLE statement 2-224
 - in CREATE XADATASOURCE statement 2-279
 - in DELETE statement 2-303
 - in DESCRIBE INPUT statement 2-317
 - in DESCRIBE statement 2-313
 - in EXECUTE statement 2-350, 2-353, 2-355
 - in FETCH statement 2-368
 - in INTO EXTERNAL clause 2-574
 - in MERGE statement 2-452
 - in OPEN statement 2-460, 2-463
 - in PUT statement 2-479
 - in START VIOLATIONS TABLE statement 2-664
- USING SQL DESCRIPTOR keywords
 - in DESCRIBE INPUT statement 2-315, 2-317
 - in DESCRIBE statement 2-311, 2-313
 - in EXECUTE statement 2-355
- UTC_CURRENT option of DBINFO 4-106
- UTC_TO_DATETIME option of DBINFO 4-106

UTF-8 locale 2-196, 2-586, 4-24

Utilities

chmod 6-59
dbschema 2-42
oncheck 2-149, 4-118, 6-3
ondblog 2-640
oninit 4-104
onmode 2-582, 2-655, 6-27, 6-29, 6-43
onspaces 2-8, 2-135, 2-426, 2-518, 6-3
onstat 2-582, 2-598, 2-618, 2-655, 6-33
onutil 2-109, 2-113, 4-118
setenv 2-86
setnet32 2-85

V

V option of oninit 4-104

VALUE clause

after null value is fetched 2-384
relation to FETCH 2-383
use in GET DESCRIPTOR 2-382
use in SET DESCRIPTOR 2-606

VALUE keyword

in SET DESCRIPTOR statement 2-605

VALUES clause

effect with PUT 2-480
in INSERT statement 2-434
in MERGE statement 2-452

VALUES keyword

in MERGE statement 2-452

VALUES keyword, in INSERT statement

2-434
VARCHAR data type 4-22, 4-24
in LOAD statement 2-446
in UNLOAD statement 2-685
syntax 4-23

VARIABLE keyword, in CREATE OPAQUE TYPE statement

2-145
Variable-length UDT 2-194, 4-34

Variables

declaring in SPL 3-8
default values in SPL 3-10, 3-13
global 3-10
local 3-8, 3-12
PROCEDURE type 3-16
uninitialized 3-49, 4-9
unknown values in IF 3-30

VARIANCE function

4-165, 4-173

Variant function

2-54, 5-20, 5-66

VARIANT keyword

External Routine Reference segment 5-19
Routine Modifier segment 5-60

Varying-length opaque data type

2-146, 2-147

VERCOLS keyword

in ALTER TABLE statement 2-47
in CREATE TABLE statement 2-210

version option of DBINFO

4-104

View

affected by dropping a column 2-56
affected by modifying a column 2-62
creating a synonym for 2-187
creating a view 2-274
dependent 2-346
dropped by ALTER FRAGMENT statement 2-14
dropping 2-346
materialized 2-275
privileges 2-403
typed 2-168, 2-275

View (continued)

union 2-276
untyped 2-275
updatable 2-278
updating 2-691

VIEW keyword

in CREATE VIEW statement 2-274
in DROP VIEW statement 2-346

VIOLATIONS keyword

in START VIOLATIONS TABLE statement 2-664
in STOP VIOLATIONS TABLE statement 2-676

Violations table

creating 2-664
effect on transactions 2-665
examples 2-670, 2-677
how to stop 2-676
privileges 2-669
relationship to diagnostics table 2-205, 2-668
relationship to target table 2-668
restriction on dropping 2-344
restriction on moving 2-456
schema 2-667
security label protection 2-668

Virtual column

2-276

Virtual index

2-91

Virtual processors

2-618
adding or removing with SQL Admin API 6-36

Virtual table

5-5

Virtual-processor class

5-61

Visual disabilities

reading syntax diagrams C-1

VP_MEMORY_CACHE_KB configuration parameter

6-23, 6-41

VP.

2-618

VPCLASS configuration parameter

6-37, 6-38

W

WAIT keyword, in SET LOCK MODE statement

2-638

WARNING keyword

in SET ISOLATION statement 2-633

Weekday argument to ADD_MONTHS function

4-129, 4-131

Weekday argument to NEXT_DAY function

4-133

WEEKDAY function

4-128

Well-behaved C UDRs

5-62

WHEN keyword

in CASE statement 3-5
in CREATE TRIGGER statement 2-258
in EXIT statement 3-17

in Expression segment 4-66, 4-67
in MERGE statement 2-452

WHENEVER statement

syntax and use 2-718

WHERE clause

2-528
estimated selectivity 2-712
in system descriptor area 2-3
joining tables 2-564
with a subquery 2-561
with ALL keyword 2-564
with ANY keyword 2-564
with BETWEEN keyword 2-562
with IN keyword 2-561
with IS keyword 2-562
with relational operator 2-561
with SOME keyword 2-564

WHERE CURRENT OF keywords

2-689
in DELETE statement 2-303

WHERE CURRENT OF keywords (*continued*)
 optimizer directives 5-37

WHERE keyword 2-689
 in CREATE INDEX statement 2-143
 in DELETE statement 2-303
 in SELECT statement 2-560

WHILE keyword
 in CONTINUE statement 3-7
 in EXIT statement 3-17

WHILE statement 3-49

White space characters
 delimited identifiers 5-24
 SQL statements 1-1

Wildcard character
 asterisk (*) 4-14
 backslash (\) 4-13, 4-14
 brackets ([]) 4-14
 caret (^) 4-14
 percent sign (%) 4-13
 question mark (?) 4-14
 underscore (_) 4-13
 with LIKE 2-562, 4-13
 with LIKE or MATCHES 4-192
 with MATCHES 2-562, 4-14

Windows
 adding or removing virtual processors 6-38
 batch file 3-46
 sqlhosts subkey 2-85

WITH APPEND keywords
 in SET DEBUG FILE statement 2-600
 in SET EXPLAIN statement 2-624
 in SET PLOAD FILE statement 2-646

WITH BUFFERED LOG keywords
 in CREATE DATABASE statement 2-98

WITH CHECK OPTION keywords, in CREATE VIEW statement 2-274, 2-277

WITH CONCURRENT TRANSACTION keywords
 in CONNECT statement 2-89

WITH CRCOLS keywords
 in CREATE TABLE statement 2-209
 in CREATE Temporary TABLE statement 2-235

WITH ERROR keywords
 in ALTER TABLE statement 2-51
 in CREATE INDEX statement 2-138
 in CREATE TABLE statement 2-203, 2-205
 in SET Database Object Mode statement 2-594

with FOREACH 3-23

WITH GRANT OPTION keywords 2-396
 in GRANT FRAGMENT statement 2-423

WITH HOLD keywords 3-23
 in DECLARE statement 2-286

WITH IDSLBACRULES keywords
 in CREATE SECURITY POLICY statement 2-182

WITH keyword 2-396
 in ALLOCATE DESCRIPTOR statement 2-3
 in ALTER FRAGMENT statement 2-19
 in ALTER FUNCTION statement 2-28
 in ALTER PROCEDURE statement 2-32
 in ALTER ROUTINE statement 2-33
 in ALTER SEQUENCE statement 2-42
 in ALTER TABLE statement 2-55, 2-60
 in CONNECT statement 2-83
 in CREATE AGGREGATE statement 2-92
 in CREATE CAST statement 2-95
 in CREATE DATABASE statement 2-97
 in CREATE FUNCTION statement 2-113
 in CREATE INDEX statement 2-138

WITH keyword (*continued*)
 in CREATE PROCEDURE statement 2-153
 in CREATE SECURITY POLICY statement 2-182
 in CREATE SEQUENCE statement 2-186
 in CREATE TABLE statement 2-193, 2-209, 2-210
 in CREATE TEMP TABLE statement 2-230
 in CREATE Temporary TABLE statement 2-235
 in CREATE VIEW statement 2-274, 2-277
 in DECLARE statement 2-286
 in EXECUTE FUNCTION statement 2-356
 in EXECUTE PROCEDURE statement 2-364
 in GRANT FRAGMENT statement 2-423
 in MOVE TABLE statement 2-455
 in OPEN statement 2-460
 in SET CONSTRAINTS statement 2-590
 in SET Database Object Mode statement 2-594
 in SET DEBUG FILE statement 2-600
 in SET ENCRYPTION PASSWORD statement 2-611
 in SET EXPLAIN statement 2-624
 in SET INDEXES statement 2-631
 in SET ISOLATION statement 2-633
 in SET PLOAD FILE statement 2-646

WITH LISTING IN keywords
 in CREATE FUNCTION statement 2-113
 in CREATE PROCEDURE statement 2-153

WITH LOG keyword
 in CREATE DATABASE statement 2-97

WITH LOG MODE ANSI keywords
 in CREATE DATABASE statement 2-98

WITH MAX keywords
 in ALLOCATE DESCRIPTOR statement 2-3

WITH NO LOG keywords
 in CREATE TEMP TABLE statement 2-230
 in CREATE Temporary TABLE statement 2-231
 in SELECT statement 2-574, 2-576

WITH REOPTIMIZATION keywords in OPEN statement 2-460

WITH RESUME keywords
 in ON EXCEPTION statement 3-38
 in RETURN statement 3-42

WITH ROWIDS keywords
 in ALTER FRAGMENT statement 2-19
 in CREATE TABLE statement 2-212

WITH TRIGGER REFERENCES keywords
 in EXECUTE FUNCTION statement 2-356
 in EXECUTE PROCEDURE statement 2-364

with UNION operator 2-528

WITH VERCOLS keywords
 in CREATE TABLE statement 2-210

WITH WARNING keywords
 in SET ISOLATION statement 2-633

WITHOUT ERROR keywords
 in ALTER TABLE statement 2-51
 in CREATE INDEX statement 2-138
 in CREATE TABLE statement 2-203, 2-205
 in SET Database Object Mode statement 2-594

WITHOUT HEADINGS keywords
 in OUTPUT statement 2-467

WITHOUT keyword
 in ALTER TABLE statement 2-51
 in BEGIN WORK statement 2-74
 in CREATE INDEX statement 2-138
 in CREATE TABLE statement 2-203
 in OUTPUT statement 2-467
 in SET CONSTRAINTS statement 2-590
 in SET Database Object Mode statement 2-594
 in SET INDEXES statement 2-631

Word length (32-bit or 64-bit) 4-104
 WORK keyword
 in BEGIN WORK statement 2-74
 in COMMIT WORK statement 2-81
 in ROLLBACK WORK statement 2-521
 WORK WITHOUT REPLICATION keywords, in BEGIN WORK statement 2-74
 WRITE keyword
 in CREATE SECURITY POLICY statement 2-182
 in GRANT statement 2-418
 in REVOKE statement 2-513
 in SET TRANSACTION statement 2-656
 Write lock 2-292
 WRITEDOWN keyword
 in GRANT statement 2-416
 in REVOKE statement 2-511
 WRITEUP keyword
 in GRANT statement 2-416
 in REVOKE statement 2-511
 Writing direction 4-151

X

X for storage in an extent space 2-91, 5-52
 x option to onstat utility 6-44
 X/Open DTP XA standard 2-279, 2-280
 X/Open mode
 CONNECT statement 2-87
 FETCH statement 2-369
 GET DESCRIPTOR statement 2-383
 OPEN statement 2-464
 SET DESCRIPTOR statement 2-609
 XA data source
 privileges to create 2-279
 privileges to drop 2-347
 XA data source type
 creating 2-280
 dropping 2-348
 XA Switch Structure 5-54
 xa.h file 5-54
 XADATASOURCE keyword
 in CREATE XADATASOURCE statement 2-279
 in CREATE XADATASOURCE TYPE statement 2-280
 in DROP XADATASOURCE statement 2-347
 in DROP XADATASOURCE TYPE statement 2-348
 XID data type 2-536, 4-22
 xopen compiler option 2-609
 XOR bitwise logical operation 4-54

Y

Y
 option to onmode utility 6-43
 value of MORE field 2-390
 YEAR function 4-131
 YEAR keyword 4-38, 4-176, 4-181, 4-183
 YES
 NODEFDAC setting 2-402

Z

Z option to onmode utility 6-44
 Zero
 AUTO_REPREPARE setting 2-620
 CLIENT_TZ setting 2-617
 DIRECTIVES setting 5-36

Zero (*continued*)
 EXT_DIRECTIVES setting 2-523
 IFX_AUTO_REPREPARE setting 2-620
 IFX_EXTDIRECTIVES setting 2-524
 in UNLOAD file 2-685
 invalid divisor 2-389
 longitude 2-617
 OPTCOMPIND setting 2-621
 prohibited MOD divisor 4-91
 returned ADMIN value 6-3
 returned IFX_REPLACE_MODULE value 6-55
 returned IFX_UNLOAD_MODULE value 6-56
 scale and ROUND function 4-92
 scale and TRUNC function 4-95
 setting of STMT_CACHE_NOLIMIT 2-654
 sqlca value after INSERT 2-441
 sqlcode value after ALLOCATE COLLECTION 2-2
 subseconds and CURRENT function 4-76
 Sunday and WEEKDAY function 4-131
 sysdirectives.active value 2-525
 time unit value returned by EXTEND 4-134
 to specify next serial value 2-436
 variance and STDEV function 4-173
 variance and VARIANCE function 4-173
 Zoned decimal 2-106
 ZONED keyword
 in CREATE EXTERNAL TABLE statement 2-105



Printed in USA

SC23-7751-02



Spine information:

IBM Informix **Version 11.50**

IBM Informix Dynamic Server Guide to SQL: Syntax

