

Note

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

This edition replaces SC23-7758-01.

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	ix
In This Introduction	ix
About This Publication	ix
Assumptions About Your Locale	ix
What's New in Migration for Dynamic Server, Version 11.50	x
Documentation Conventions	xi
Typographical Conventions	xi
Feature, Product, and Platform Markup	xi
Example Code Conventions	xii
Additional Documentation	xii
Compliance with Industry Standards	xii
Syntax Diagrams	xiii
How to Read a Command-Line Syntax Diagram	xiv
Keywords and Punctuation	xv
Identifiers and Names	xv
How to Provide Documentation Feedback	xv

Part 1. Overview of Migration to IBM Informix Dynamic Server Version 11.50

Chapter 1. Overview of Dynamic Server Migration 1-1

Dynamic Server Enterprise or Express Edition Migration	1-1
The Migration Process	1-1
Migration Effort	1-1
Migration Skills	1-1
Migration Plans	1-2
Types of Migration	1-2
Migration Tools	1-3
Upgrading Dynamic Server (In-place Migration)	1-3
Migrating Dynamic Server (Non-in-place Migration)	1-4
Hardware Prerequisites	1-5
Software Prerequisites	1-5
Operating Systems Supported by Dynamic Server Version 11.50	1-5
Operating Systems that Version 11.50 Does Not Support	1-5
Migration Paths to Version 11.50	1-5
Naming Conventions for Dynamic Server Releases and Fix Packs	1-7
Migration Paths on UNIX Operating Systems	1-7
Migration Paths on Linux Operating Systems	1-8
Migration Paths on Windows Operating Systems	1-8
Changes in Dynamic Server	1-9

Chapter 2. Overview of Moving Data 2-1

Automatic Data Migration	2-1
Prerequisites Before Moving Data	2-1
Data-Migration Tools	2-1
High-Performance Loader Performance Advantages for Large Databases	2-4
Use Nonlogging Raw Tables with Utilities to Speed Up Data Loading Time	2-5
TEXT and BYTE Data Scanned, Not Compressed	2-6
Moving Non-Informix Data Between Computers and Dbspaces	2-6
Importing Data from a Non-Informix Source	2-6
Importing Data with IBM Informix Enterprise Gateway Products	2-6
Overview of Moving Data by Using Distributed SQL	2-7

Part 2. Migration to and Reversion from Version 11.50

Chapter 3. Preparing for Migration to Dynamic Server Version 11.50	3-1
Preparing for Migration	3-1
Check and Configure Available Space	3-2
Save Copies of the Current Configuration Files	3-4
Save a Copy of the Storage Manager sm_versions File	3-5
Close All Transactions and Shut Down the Source Database Server	3-5
Initiate Fast Recovery to Verify that No Open Transactions Exist	3-5
Verify the Integrity of the Data	3-6
Verify that the Database Server Is in Quiescent Mode	3-7
Disable High-Availability Data Replication.	3-7
Make a Final Backup of the Source Database Server	3-7
Verify That the Source Database Server Is Offline	3-7
Modify Kernel Parameters (UNIX, Linux)	3-8
Checklist of Diagnostic Information for Performance and Troubleshooting	3-8
Migrating from 32-bit to 64-Bit Database Servers.	3-9
 Chapter 4. Migrating to and Reverting from Dynamic Server Version 11.50 with Enterprise Replication.	 4-1
Preparing to Migrate to Version 11.50 with Enterprise Replication	4-1
Migrating to Version 11.50 with Enterprise Replication.	4-1
Converting Replication of 9.21 User-Defined Data Types	4-2
Reverting from Dynamic Server Version 11.50 with Enterprise Replication	4-2
Express Edition Migration and Enterprise Replication	4-4
 Chapter 5. Migrating to Dynamic Server Version 11.50	 5-1
Migrating to Version 11.50	5-1
Install Dynamic Server Version 11.50.	5-2
Set Environment Variables	5-4
Customize Configuration Files.	5-4
Optionally Add Communications Support Modules.	5-5
Install or Upgrade Any DataBlade Modules	5-5
Initialize Dynamic Server Version 11.50	5-5
Upgrade the High-Performance Loader onpload Database	5-6
Completing Required Post-Migration Tasks	5-6
For ON-Bar, Rename the sm_versions.std File.	5-7
Update Statistics After Migrating	5-7
Review Client Applications and Registry Keys	5-8
Verify the Integrity of the Data	5-8
Back Up Dynamic Server Version 11.50	5-9
Tune Dynamic Server Version 11.50 for Performance	5-9
Prepare HDR Secondary Servers	5-9
Register DataBlade Modules	5-9
 Chapter 6. Reverting from Dynamic Server Version 11.50.	 6-1
Overview of Reversion	6-1
Preparing to Revert from Dynamic Server Version 11.50	6-1
Ascertain that Reversion Is Possible and Identify Reversion Requirements	6-2
Check and Configure Available Space for Reversion	6-8
Save Copies of the Current Configuration Files	6-9
Save System Catalog Information	6-9
Verify the Integrity of the Version 11.50 Data	6-9
Back Up Dynamic Server Version 11.50	6-9
Remove Version 11.50 Features	6-10
Remove New BladeManager Extensions	6-10
Disable HDR Before Reversion	6-10
Reverting from Dynamic Server Version 11.50	6-10
Run the Reversion Utility	6-11
Restore Original Configuration Parameters	6-12
Restore Original Environment Variables	6-12
Remove Any Communications Support Module Settings.	6-12

Recompile Any Java UDRs that Were Compiled Using JDK 5.0	6-12
Reinstall and Start the Earlier Database Server	6-12
Update Statistics After Reverting	6-13
Verify the Integrity of the Reverted Data	6-13
Back Up the Database Server	6-13
Return the Database Server to Online Mode	6-13
Prepare to Restart HDR.	6-14

Part 3. Migration of Data Between Database Servers

Chapter 7. Migrating Database Servers to a New Operating System	7-1
Choosing a Tool for Moving Data Before Migrating Between Operating Systems	7-1
Adjusting Database Tables for File-System Variations	7-2
Moving Data to a Database Server on a Different Operating System	7-2
Adapting Your Programs for a Different Operating System	7-2
Ensuring the Successful Creation of System Databases	7-3
Moving Data Between Dynamic Server and Workgroup Edition on Different Operating Systems.	7-3

Part 4. Data Migration Utilities

Chapter 8. The dbexport and dbimport Utilities	8-1
Overview of the dbexport and dbimport Utilities	8-1
The dbexport Utility	8-1
Syntax of the dbexport Command	8-2
Contents of the Schema File that dbexport Creates	8-5
The dbimport Utility	8-6
Syntax of the dbimport Command	8-7
Changing the Database Locale with dbimport	8-12
Simple Large Objects (Version 9.21 or Later Versions).	8-12

Chapter 9. The dbload Utility	9-1
Overview of the dbload Utility	9-1
Syntax of the dbload Command	9-2
Table Locking During a Load Operation	9-3
Rows to Ignore During a Load Operation	9-4
Bad-Row Limit During a Load Operation	9-4
Termination of the dbload Utility	9-4
Name and Object Guidelines for the dbload Utility	9-4
Command File for the dbload Utility	9-5
Delimiter Form of the FILE and INSERT Statements	9-5
Character-Position Form of the FILE and INSERT Statements	9-9
Command File to Load Complex Data Types (Version 9.21 or Later Versions).	9-13
Using the dbload Utility with Named Row Types	9-13
Using the dbload Utility with Unnamed Row Types	9-14
Using the dbload Utility with Collection Data Types	9-15

Chapter 10. The dbschema Utility	10-1
Overview of the dbschema Utility	10-1
Object Modes and Violation Detection	10-1
Guidelines for Using dbschema	10-1
Syntax of the dbschema Command	10-2
Database Schema Creation.	10-3
dbschema Server-Specific Information	10-4
User-Defined and Complex Data Types (Version 9.21 or Later Versions).	10-5
Sequence Creation	10-6
Synonym Creation	10-6
Privileges	10-7
Granting Privileges	10-7
Displaying Privilege Information for a Role	10-7

Table, View, or Procedure Creation	10-8
Table Information.	10-9
Role Creation	10-9
Distribution Information for Tables	10-10
Example of Output Showing Distribution Information	10-10
Distribution Information	10-11
Overflow Information	10-12
Use dbschema Output as DB-Access Input	10-12
Inserting a Table into a Database Example	10-12
Re-Creating the Schema of a Database	10-13

Chapter 11. The LOAD and UNLOAD Statements 11-1

Overview of the LOAD and UNLOAD Statements.	11-1
Syntax of the UNLOAD Statement	11-1
Syntax of the LOAD Statement	11-2
Load and Unload Statements for Locales That Support Multibyte Code Sets	11-2
Load and Unload Statements for Non-Default Locales and GL_DATETIME Environment Variable	11-2

Chapter 12. The onmode Utility Reversion Option 12-1

Use of the onmode -b Command for Reversion.	12-1
Preparation for Reversion	12-1
Syntax of the onmode -b Command	12-1

Chapter 13. The onunload and onload Utilities 13-1

Overview of the onunload and onload Utilities.	13-1
Requirements for Using onload and onunload	13-2
How onunload and onload Work	13-2
Syntax of the onunload Command	13-3
onunload Destination Parameters	13-4
Constraints That Affect onunload	13-4
Privileges for Database or Table Unloading	13-5
Tables that are Unloaded with a Database	13-5
Data that is Unloaded with a Table	13-5
Locking During Unload Operation	13-5
Logging Mode.	13-5
Syntax of the onload Command	13-6
onload Source Parameters	13-7
onload Create Options	13-7
Constraints That Affect onload	13-8
Logging During Loading	13-9
Movement of Simple Large Objects to a BlobSpace.	13-9
Ownership and Privileges	13-9
Exclusive Locking During Load Operation	13-10
Using onunload and onload.	13-10

Part 5. Appendixes

Appendix A. New Environment Variables A-1

Appendix B. New Configuration Parameters B-1

Appendix C. Configuration Parameters That Have Been Changed or Removed C-1

Configuration Parameter Changes in the Version 11.50 onconfig.std File.	C-1
Configuration Parameters that Have Been Changed or Removed In Versions 9.30 through 11.10.	C-6

Appendix D. New Reserved Words D-1

Appendix E. System Catalog and System Database Changes E-1

Changes for Dynamic Server 11.50	E-1
Changes for Dynamic Server 11.10	E-1
Changes for Dynamic Server 10.0.	E-2
Changes for Dynamic Server 9.40.	E-2
Changes for Dynamic Server 9.30.	E-2
Column-Width Changes in sysmaster Tables in Version 9.20 and Later Versions	E-3
Data Type Changes in sysmaster tables in Version 9.20 and Later Versions	E-4
Changes in Treatment of Null Values in sysmaster Tables in Version 9.30	E-5
Other sysmaster Database Table and Column Changes in Version 9.30	E-5
Remote Queries on System Catalog Tables Between Version 7.31 and Later Versions.	E-5
Difference in sysindexes between Version 7.31 and Later Versions	E-6
Appendix F. New and Changed Features	F-1
Server Library Name Changes in Version 11.50	F-1
Appendix G. Accessibility	G-1
Accessibility features for IBM Informix Dynamic Server	G-1
Accessibility Features	G-1
Keyboard Navigation	G-1
Related Accessibility Information.	G-1
IBM and Accessibility	G-1
Dotted Decimal Syntax Diagrams	G-1
Notices	H-1
Trademarks	H-3
Index	X-1

Introduction

In This Introduction

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

About This Publication

This publication describes how to migrate to either Enterprise or Express editions of Version 11.50 of IBM® Informix® Dynamic Server. Migration sometimes involves reverting to the database server that you migrated from, or moving data manually between databases, servers, and computers. This publication also contains information on how to use the **dbexport**, **dbimport**, **dbload**, **dbschema**, **onload**, and **onunload** data-migration utilities and the LOAD and UNLOAD SQL statements.

This publication does not contain information on using the DSN (data source name) Migration Utility (**dsnigrate.exe**) to migrate from one version of CSDK to another. For information on migrating DSN, see the "DSN Migration Utility" section in the *IBM Informix ODBC Driver Programmer's Manual*.

For information about migrating to previous versions of IBM Informix database servers, see the *Migration Guide* in the documentation set for that version of the server.

For information on the Informix client-server environment, read *IBM Informix Dynamic Server Getting Started Guide*. That book contains information on the differences between Informix database servers plus information on network and server configurations.

Migration includes conversion (upgrading) to a later version of a database server, reversion to an earlier version of a database server, and movement of data between databases, database servers on the same operating system, database servers on different operating systems, and different kinds of database servers. Conversion or reversion often involves changing connectivity information in the **sqlhosts** file or registry key, host environment variables, configuration parameters, and other database server features.

Assumptions About Your Locale

IBM Informix products can support many languages, cultures, and code sets. All the information related to character set, collation, and representation of numeric data, currency, date, and time is brought together in a single environment, called a GLS (Global Language Support) locale.

The examples in this publication are for the default locale, **en_us.8859-1**. This locale supports U.S. English format conventions for date, time, and currency. In addition, this locale supports the ISO 8859-1 code set, which includes the ASCII code set plus many 8-bit characters such as è, é, and ñ.

If you plan to use nondefault characters in your data or your SQL identifiers, or if you want to conform to the nondefault collation rules of character data, you must specify the 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*.

What's New in Migration 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 the *IBM Informix Migration Guide*.

Table 1. What's New for Migration in Version 11.50xC1

Overview	Reference
<p>Configuration file improved</p> <p>Starting in Version 11.50, the configuration parameters in the onconfig.std file are organized by functional area. Some configuration parameters have new default values. Also, deprecated configuration parameters were removed from the file. Before you migrate to the current version, ensure that you save a copy of your old file and compare your existing configuration settings with those in the new configuration file.</p>	<p><i>IBM Informix Dynamic Server Getting Started Guide</i></p> <p>"Configuration Parameter Changes in the Version 11.50 onconfig.std File" on page C-1</p> <p>"Preparing for Migration" on page 3-1</p>
<p>Support added for Mac OS X operating systems</p> <p>Informix Dynamic Server Version 11.50, Enterprise Edition runs on the Mac OS X operating system. If you want to move your environment from your current operating system to the Mac OS X, you must migrate your data. You cannot simply upgrade from another operating system to a Mac OS X operating system.</p>	<p>Chapter 7, "Migrating Database Servers to a New Operating System," on page 7-1</p>
<p>Informix Dynamic Server Developer Edition enhancements</p> <p>If you are using Developer Edition Version 11.10, you can upgrade to take advantage of Version 11.50 enhancements. Also, you can now run Developer Edition on the Mac OS X and Ubuntu operating systems.</p> <p>Developer Edition is not for production environments. You can use it only for developing database applications. Developer Edition is free, and supported only at the IDS Developer Edition wiki.</p>	<p>http://www.informix-zone.com/idswiki/doku.php</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.

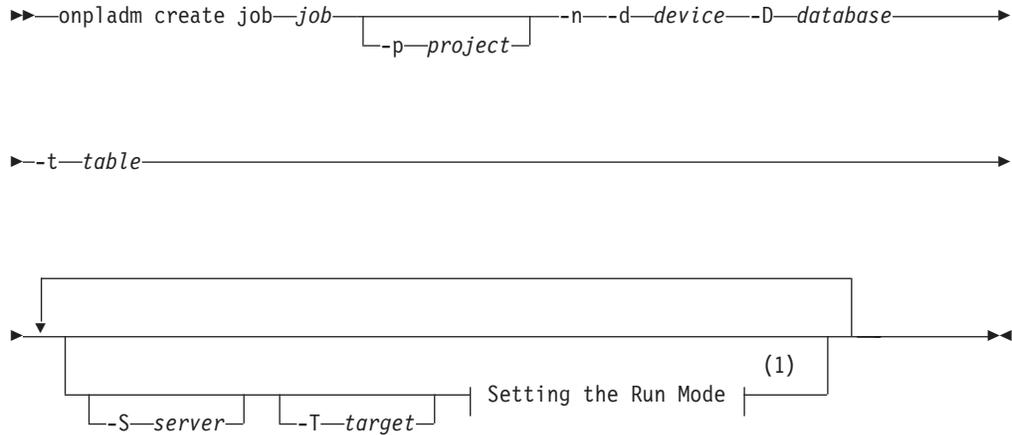
Table 2. Syntax Diagram Components

Component represented in PDF	Component represented in HTML	Meaning
	>>-----	Statement begins.
	----->	Statement continues on next line.
	>-----	Statement continues from previous line.
	----->>	Statement ends.
	-----SELECT-----	Required item.
	--+-----LOCAL-----+	Optional item.
	---+-----ALL-----+--- +--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 	---+-----view-----+--- +-----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

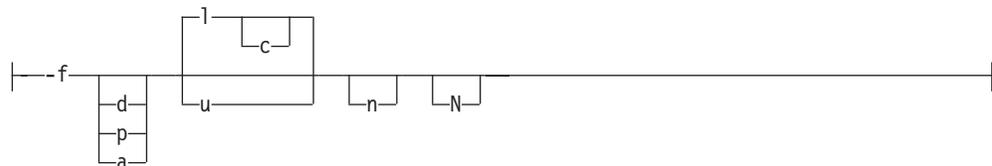


Notes:

- 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.

Part 1. Overview of Migration to IBM Informix Dynamic Server Version 11.50

Chapter 1. Overview of Dynamic Server Migration

This chapter provides an overview of database server migration, a list of the operating systems on which you can run Version 11.50, and the paths for migrating from previous IBM Informix database servers to Version 11.50 of IBM Informix Dynamic Server.

Dynamic Server Enterprise or Express Edition Migration

If you have IBM Informix Dynamic Server Version 11.10, 10.00, 9.40, 9.30, 9.21, or 7.31, you can migrate directly to either the Enterprise Edition or Express Edition of Dynamic Server Version 11.50.

If you have another version of Dynamic Server, you must first migrate to an interim version of Dynamic Server. For more information, see “Migration Paths to Version 11.50” on page 1-5.

The process for migrating to the Enterprise Edition and the Express Edition is the same, except that some particular tasks are not applicable for the Express Edition, because the Express Edition does not support Enterprise Replication and High-Availability Data Replication (HDR).

The Migration Process

This overview of the migration process describes what you need to know to plan your migration and the resources that you can use to assist you.

Careful planning will ensure minimal impact on your business.

- “Migration Effort”
- “Migration Skills”
- “Migration Plans” on page 1-2
- “Types of Migration” on page 1-2
- “Migration Tools” on page 1-3

Migration Effort

Depending on your environment, the migration process can take a few hours or several weeks.

The migration effort is determined by many factors:

- Your current version of Dynamic Server. The older the version, the greater the effort.
- The site architecture and configuration, before and after migration.
- The level of site customization, before and after migration.
- Integration of additional software products.
- To some extent, the size of the database.

Migration Skills

Your Dynamic Server migration team needs the following skills:

- Database administration skills, to help migrate custom database extensions.

- System administration skills, to perform various system tasks: operating system installation, configuration and maintenance, installation and configuration of Dynamic Server and any additional software products.
- Application programming skills, to create and maintain scripts to evaluate and modify application programs.

If you prefer, highly-skilled IBM Services personnel and business partners are available to assist you in migrating your environment. Contact your IBM representative for further information.

Migration Plans

Before you begin to migrate to a new version of the database server, you can plan for migration.

To plan your migration requirements, complete these tasks:

1. Inventory the existing Dynamic Server environment assets, such as machines, instances, databases, database customization, custom code, IBM software, and third-party software.
2. Itemize the requirements for the post-migrated environment. New requirements can include upgrading or adding hardware, using new features, or replacing custom-code with new built-in function.
3. Plan the migration activities. Typical activities include:
 - Performing a level-0 backup of the database.
 - Quiescing the database server and preventing connections to the database until migration completes.

Important: Any connection attempts (for example, from cron jobs or monitoring scripts) to the database after you quiesce the database server and during migration will cause migration to fail.

 - Installing Dynamic Server Version 11.50.
 - Migrating database data.
 - Reverting to the previous version.
 - Migrating applications before using them with the new database server.

Depending on your environment, you might need to perform some of these activities more than once. You might not need to restore the level-0 backup; however, if you encounter problems you can always restore the backup of your current server.

Types of Migration

There are basically three ways to migrate to Dynamic Server Version 11.50:

Upgrading (In-place migration)

Upgrading is a special case of migration that uses your existing hardware and operating system. You install a new or improved version of the product in a different location from your current version on the same machine. You can copy your configuration file and add new parameters. When you start the new Dynamic Server instance, the database data is automatically converted. For example, you can upgrade from Version 11.10 to Version 11.50.

Migrating (Non-in-place migration)

The process of “switching over” your environment from one computer to another. This type of migration requires more planning and setup time

compared to upgrading on your existing computer. Non-in-place migration requires that you modify and copy the database schema, user data, and user objects from one server to another server. Use this type of migration if you are moving to Dynamic Server Version 11.50 from an early version of Dynamic Server that has a different architecture, page size, optimization of dbspaces, and extent allocations.

Migrating from a non-IBM database

The process of moving your data from another database management system (DBMS) such as Oracle or Sybase SQL Anywhere to Dynamic Server Version 11.50. This type of migration is especially useful if you are currently using various products. You can consolidate to take advantage of the Dynamic Server features and total cost of ownership.

Migration Tools

You can choose from various migration tools, depending on the task that you must perform.

- For in-place upgrades: You do not use any migration tools. Simply start the server by using the **oninit** utility. The data from the source database server is converted to the target database server automatically.
- For non-in-place migration: You can use distributed queries to move your data, or you can pick from a number of data transfer tools and utilities, such as:
 - **dbexport** and **dbimport**
 - **onload** and **onunload**
 - **dbload** and **dbunload**
 - High Performance Loader (HPL)

Each of these tools and utilities has specific advantages and limitations. Consider all of the variables and pick a tool or utility that is appropriate for your situation.

- For migrating your data from non-IBM products: To migrate from a variety of source databases to Informix Dynamic Server, regardless of platform, download the free, easy-to-use <http://www.ibm.com/software/data/db2/migration/mtk/>. You can also use the HPL to move your data from non-IBM products to Dynamic Server.

Upgrading Dynamic Server (In-place Migration)

Depending on your existing Database Server setup, you might be able to upgrade directly to the current version by installing the product in a new directory, copying a few configuration files, and starting the new server to convert your database data.

You can upgrade directly from any of the following products: Dynamic Server Version 11.10, 10.00, 9.40, 9.30, 9.21, or 7.31.

Upgrading is an in-place migration method that uses your existing test and production hardware. The operating system on those machines must be supported by Dynamic Server Version 11.50. Also, you must have enough space for the system database data conversion.

Upgrading consists of these steps:

1. Prepare your system. That includes removing outstanding in-place alters, closing all transactions, verifying the integrity of the data with **oncheck**, and

performing a level-0 backup. If you are using Enterprise Replication or High Availability Data Replication, stop replication.

2. Install the new product on the machine.
Important: Do not install it over the existing product.
3. Copy the ONCONFIG file to the target and set parameters that are new for the current release.
4. Start the Dynamic Server Version 11.50 instance. The database data is automatically converted.
5. Run UPDATE STATISTICS MEDIUM for non-leading index columns, then run UPDATE STATISTICS FOR PROCEDURE.

This type of migration minimizes the risk of introducing errors. You can always revert from the new server to the old one. In the event of a problem during reversion, you can restore the level-0 backup.

Related Information

- “Operating Systems Supported by Dynamic Server Version 11.50” on page 1-5
- Technical document <http://www.ibm.com/support/docview.wss?uid=swg21144602>
- Chapter 3, “Preparing for Migration to Dynamic Server Version 11.50,” on page 3-1

Migrating Dynamic Server (Non-in-place Migration)

Depending on your existing Database Server setup, you might have to move to a new architecture or a different server. This type of migration is more complicated than upgrading. It requires more planning and it is more time consuming.

The non-in-place type of migration consists of these steps:

1. Prepare your system. That includes removing outstanding in-place alters, closing all transactions, verifying the integrity of the data with **oncheck**, and performing a level-0 backup. If you are using Enterprise Replication or High-Availability Data Replication, stop replication.
2. Install the new product on a new machine.
3. Create a database with the current schema (dbschema -d dbname -ss).
4. Modify the schema for new extent allocations and lock mode changes. If applicable, modify schema for new dbspaces.
5. Move data by using the appropriate utility or tool, or by using distributed queries.
6. Start the Dynamic Server Version 11.50 instance.
7. After Dynamic Server migration, upgrade applications before running them. Also, if you use Enterprise Replication or HDR, start replication.

Related Information

- “Operating Systems that Version 11.50 Does Not Support” on page 1-5
- Chapter 3, “Preparing for Migration to Dynamic Server Version 11.50,” on page 3-1
- “Data-Migration Tools” on page 2-1

Hardware Prerequisites

Your hardware must support the operating systems that Dynamic Server Version 11.50 supports and must provide enough disk space for Version 11.50 and all your other software applications.

Software Prerequisites

You can download the Dynamic Server installation package from the Web or install from the product CD.

To download the installation package, log on to Passport Advantage® at <http://www.ibm.com/software/howtobuy/passportadvantage> and follow the directions provided.

Operating Systems Supported by Dynamic Server Version 11.50

You can run IBM Informix Dynamic Server, Version 11.50 on various Linux®, UNIX®, and Windows® operating systems.

Table 1-1. Operating Systems on Which to Run Version 11.50

IBM Informix Dynamic Server	Operating Systems
Enterprise Edition	Linux UNIX Windows
Express Edition	Linux Windows

For a detailed list of the operating systems supported by the current version of Dynamic Server and by other IBM Informix products, download the platform availability spreadsheet from <http://www.ibm.com/software/data/informix/pubs/roadmaps.html>. Search for the product name, or sort the spreadsheet by name.

Prepare the operating system by completing any required machine or operating system changes. See the machine notes, which are in file **ids_machine_notes_11.50.txt** on the installation media, for information about platform-specific actions that you must take to configure and use IBM Informix products.

Operating Systems that Version 11.50 Does Not Support

Version 11.50 does not run on the Windows 2000, Windows NT®, or Windows 95 operating systems. If you are currently using a version of Dynamic Server on one of those operating systems, you must migrate to an operating system on which Version 11.50 runs.

For specific information on migrating between operating systems, see Chapter 7, “Migrating Database Servers to a New Operating System,” on page 7-1.

Migration Paths to Version 11.50

Dynamic Server Version 11.50 is the superset of function of all previous versions. If you have Dynamic Server Version 11.10, 10.0, 9.40, 9.30, 9.21 or 7.31, you can migrate directly to Dynamic Server Version 11.50.

If you are migrating from another database server version, you must migrate to an earlier, interim version of Dynamic Server before you migrate to Version 11.50.

Table 1-2. Migration Paths for Moving to a Newer Database Server

Source Database Server	Target Database Server	Reference
Dynamic Server 11.10	You can migrate directly to Dynamic Server Version 11.50.	The Version 11.50 Migration Guide
Dynamic Server 10.00		
Dynamic Server 9.40		
Dynamic Server 9.30		
Dynamic Server 9.21		
Dynamic Server 7.31		
Dynamic Server 7.30	You must first migrate to Dynamic Server 10.00 or Dynamic Server 9.40.	The Version 10.00 or 9.40 Migration Guide
Dynamic Server 7.24		
Workgroup Edition 7.24		
Universal Server 9.14	You must first migrate to Dynamic Server 9.30 or Dynamic Server 9.21.	The Version 9.30 Migration Guide
Dynamic Server 9.20		
OnLine Dynamic Server 7.23	You must first migrate to Dynamic Server 7.31.	The Version 7.31 Migration Guide
OnLine Dynamic Server 7.22		
OnLine 5.1 or earlier versions		

Figure 1-1 illustrates the paths for converting from specific older database servers to Dynamic Server Version 11.50.

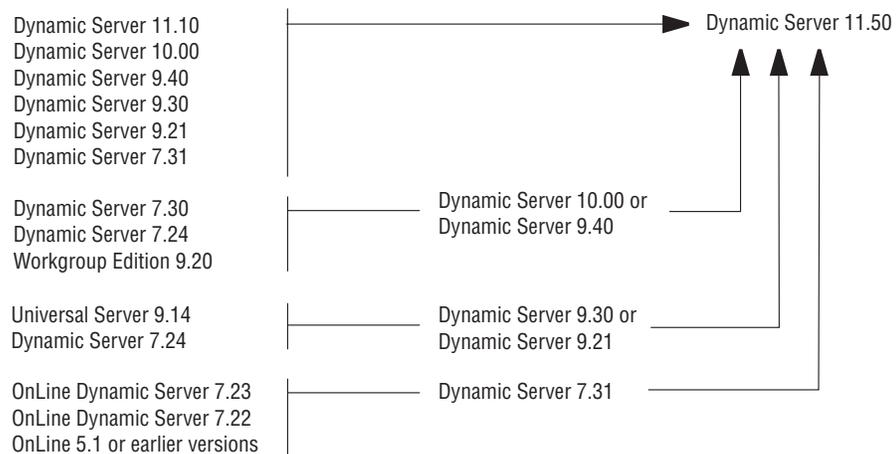


Figure 1-1. Migration Paths to Dynamic Server Version 11.50

If you need to migrate to another version of Dynamic Server before you migrate to Version 11.50, see the *Migration Guide* that is included in the documentation set for that database server. For example, for information about migrating to Version 7.31 before migrating to Version 11.50, see the *Version 7.31 IBM Informix Migration Guide*. You must follow the steps in that guide before you follow the instructions in the current guide for migrating to Version 11.50.

If necessary, you can revert to the server from which you upgraded.

Naming Conventions for Dynamic Server Releases and Fix Packs

Dynamic Server releases and fix packs contain version names that appear in the format `aa.bb.xCn`.

In this format:

- aa = major release number
- bb = minor release number
- x = all operating system platforms, unless one of the following characters appears in the position of x:
 - F = 64-bit on any UNIX, Linux, or Windows platform
 - H = 32-bit build on any HP 11.x platform; also runs on HP 11.x 64-bit
 - J = Java™
 - T = 32-bit on Windows platforms
 - U = 32-bit on any UNIX or Linux platform
- C = GA release
- n = fix pack level

A version with the letter E at the end means the version is for the Express Edition; a version with the letters DE means the version is for the Developer Edition.

For example, in Dynamic Server Version 11.10.xC2, 11 is the major release number, 10 is the minor release number, x means any platform, C means GA release, and 2 means fix pack 2, for either the Enterprise or Workgroup.

Migration Paths on UNIX Operating Systems

If you have Dynamic Server Version 11.10, 10.0, 9.40, 9.30, 9.21 or 7.31 on a UNIX operating system, you can migrate directly to Dynamic Server Version 11.50 on a UNIX operating system.

The following table lists the source and target versions for the migration and reversion of database servers on UNIX. You can migrate from the source versions shown in the first column to the target version shown in the second column.

Important: If the target version is not Version 11.50, you must first migrate to one of the interim servers shown in the Target Version column before you can migrate to the current version.

Table 1-3. Source and Target Versions for Migration and Reversion on UNIX

Source Version	Target Version
11.10	11.50
10.00	11.50
9.40, 9.30, 9.21	11.50
9.20, 9.14	9.30, 9.21
8.50, 8.40, 8.32, 8.31, 8.30, 8.21	7.31, 7.30, 7.24 (without Enterprise replication)
8.21	7.24 (without Enterprise replication)
7.31	11.50

Table 1-3. Source and Target Versions for Migration and Reversion on UNIX (continued)

Source Version	Target Version
7.30, 7.24 (without Enterprise replication)	10.0, 9.40
7.24 (with Enterprise replication), 7.23, 7.22	7.31, 7.30, 7.24 (without Enterprise replication)
5.1 or earlier versions	7.31, 7.30, 7.24 (without Enterprise replication)

Migration Paths on Linux Operating Systems

If you have Dynamic Server Version 11.10, 10.0, 9.40, 9.30, 9.21 or 7.31 on a Linux operating system, you can migrate directly to Dynamic Server Version 11.50 on a Linux operating system.

The following table lists the source and target versions for the migration of database servers on Linux. You can migrate from the source versions shown in the first column to the target version shown in the second column.

Table 1-4. Source and Target Versions for Migration and Reversion on Linux

Source Version	Target Version
11.10	11.50
10.00	11.50
9.40, 9.30, 9.21	11.50
7.31	11.50

Migration Paths on Windows Operating Systems

You can migrate from several Windows operating systems to Dynamic Server Version 11.50 on Windows XP, 2003, x64, or Vista.

The following table lists the source and target versions for the migration of database servers on Windows. You can migrate from the source versions shown in the first column to the target version shown in the second column.

Dynamic Server Version 11.50 does not run on Windows 2000, Windows NT, or Windows 95 operating systems.

Table 1-5. Source and Target Versions for Migration and Reversion on Windows

Source Version	Target Version
Dynamic Server version 11.10, 10.00, 9.4, or 7.31.TD9 and higher on Windows XP	Version 11.50 on Windows XP
Dynamic Server version 11.10, 10.00, 9.4.TC4, or 7.31.TC9 and higher on Windows 2003	Version 11.50 on Windows 2003
Dynamic Server version 10.00.xC6 or later, or version 11.10 on Windows x64, Windows 2003 or Windows XP	Version 11.50 on Windows x64
Dynamic Server version 10.00 (32-bit) on Windows XP	Version 11.50 on Windows Vista
Dynamic Server version 11.10 (32-bit or 64-bit server) on Windows XP	

Changes in Dynamic Server

Each version of Dynamic Server contains new features, and new and changed environment variables, configuration parameters, SQL reserved words, system catalogs, and system databases. Some of these changes might affect your applications.

Related concepts

Appendix F, “New and Changed Features,” on page F-1

Each version of Dynamic Server contains many new and changed features.

Related reference

Appendix A, “New Environment Variables,” on page A-1

Each version of Dynamic Server contains new environment variables that could affect your installation. You might also need to adjust the values of existing environment variables.

Appendix B, “New Configuration Parameters,” on page B-1

Each version of Dynamic Server contains new configuration parameters that might affect your installation. If you need to revert to a prior version of the server, you must either replace the Dynamic Server Version 11.50 ONCONFIG configuration file with the ONCONFIG file that you used before you converted, or you must remove configuration parameters that the earlier database server does not support.

Appendix C, “Configuration Parameters That Have Been Changed or Removed,” on page C-1

Dynamic Server Version 11.50 includes an improved **onconfig.std** file, with new default values for configuration parameters. In addition, some configuration parameters used with earlier versions of Dynamic Server have been changed or removed.

Appendix D, “New Reserved Words,” on page D-1

Each version of Dynamic Server supports new SQL keywords that are reserved words and might affect migration of your applications.

Appendix E, “System Catalog and System Database Changes,” on page E-1

New versions of Dynamic Server contain system catalog table changes and **sysmaster** database changes.

Chapter 2. Overview of Moving Data

Automatic data migration occurs when you migrate from one version of the database server to another. If you need to move data during non-inplace migration or before and after reversion, you can use one of several IBM Informix tools and utilities.

For example, suppose you migrated to the current version of Dynamic Server and created a few new databases, but decide to revert to the previous version. Before you revert, you can use one of the data-migration tools to save the data you added. After reverting, you can reload the data.

Before you move data, consider these issues:

- Changes in the configuration parameters and environment variables
- Amount of memory and dbspace space that is required
- Organization of the data
- Whether you want to change the database schema to accommodate more information, to provide for growth, or to enhance performance

For information about how to move data between database servers on different operating systems, also see Chapter 7, “Migrating Database Servers to a New Operating System,” on page 7-1

For information about how to move to a different GLS locale, see the *IBM Informix GLS User's Guide*.

Automatic Data Migration

You do not need to use data movement tools or utilities for in-place migration. The data is converted automatically from the source database server to the target database server after you start the target database server.

Prerequisites Before Moving Data

Before you use any data migration utility, you must set your PATH, INFORMIXDIR, and INFORMIXSERVER environment variables.

For information about environment variables, see the *IBM Informix Guide to SQL: Reference*.

Data-Migration Tools

Dynamic Server provides tools, utilities, and SQL statements that you can use to move data from one IBM Informix database to another. Normally, if you are migrating on the same operating system, you do not need to load and unload data.

You might also want to use a data-migration tool when you have different operation systems, page sizes, or code pages. For example, UNIX or Linux and Windows store data in different page sizes. When your migration involves migrating between different operating systems, you must export data and its schema information from one database server and import the exported data into the other database server.

You can use the following tools to move data:

- The **dbexport** and **dbimport** utilities
- The **dbload** utility
- The **onunload** and **onload** utilities
- UNLOAD and LOAD statements
- The High-Performance Loader (HPL)
- Nonlogging raw tables

When you import data from non-Informix sources, you can use the following tools:

- The **dbimport** and **dbload** utilities
- The High-Performance Loader (HPL)
- IBM Informix Enterprise Gateway products
- External tables that you create. For information on creating external tables to import data from non-Informix sources, see the description of the CREATE EXTERNAL TABLE statement in the *IBM Informix Guide to SQL: Syntax*.

The best method for moving data depends on your operating system and whether you want to move an entire database, selected tables, or selected columns from a table. The following table summarizes the characteristics of the methods for loading data and the advantages and disadvantages of each method. The table also shows the database servers on which you can use the tools.

Table 2-1. Comparison of Tools for Moving Data

Tool	Description	Advantages	Disadvantages	Availability
dbexport and dbimport utility	Imports or exports a database to a text file that is stored on disk or tape	<ul style="list-style-type: none"> Can modify the database schema and change the data format Can move data between operating systems Optional logging Can import data from non-Informix sources 	<ul style="list-style-type: none"> Faster performance than the dbload utility, but slower performance than the onload utility Moves the entire database 	
dbload utility	Transfers data from one or more text files into one or more existing tables	<ul style="list-style-type: none"> Can modify database schema Can move data between operating systems Optional logging Moderately easy to use Can import data from non-Informix sources 	<ul style="list-style-type: none"> Slower performance than the dbexport, dbimport, and onload utilities 	

Table 2-1. Comparison of Tools for Moving Data (continued)

Tool	Description	Advantages	Disadvantages	Availability
onunload and onload utilities	Unloads data from a database into a file on tape or disk; loads data, which was created with the onunload command, into the database server	Fast performance Optional logging	Only moves data between database servers of the same version on the same operating system Cannot modify the database schema Logging must be turned off Difficult to use	Not available on: <ul style="list-style-type: none"> SE 7.22-7.25 SE 5.1 or earlier versions OnLine 5.1 or earlier versions
UNLOAD and LOAD statements	Unloads and loads specified rows	Can modify database schema Can move data between operating systems Easy to use Optional logging	Only accepts specified data formats	
HPL	Loads data from any ASCII or COBOL file that meets certain format requirements	For extremely large databases, has a performance advantage over other IBM Informix data-migration utilities, because it performs I/O and code-set conversions in parallel Can modify database schema Can move data between operating systems Can import data from non-Informix sources	Requires significant preparation time	Not available on: <ul style="list-style-type: none"> SE 7.22-7.25 SE 5.1x OnLine 5.1x
Nonlogging raw tables	Loads certain kinds of large tables	Can load very large data warehousing tables quickly	Does not support primary constraints, unique constraints, and rollback Requires SQL Not recommended for use within a transaction	

If you are choosing a tool for loading data, the questions shown in Figure 2-1 on page 2-4 will help you make a decision.

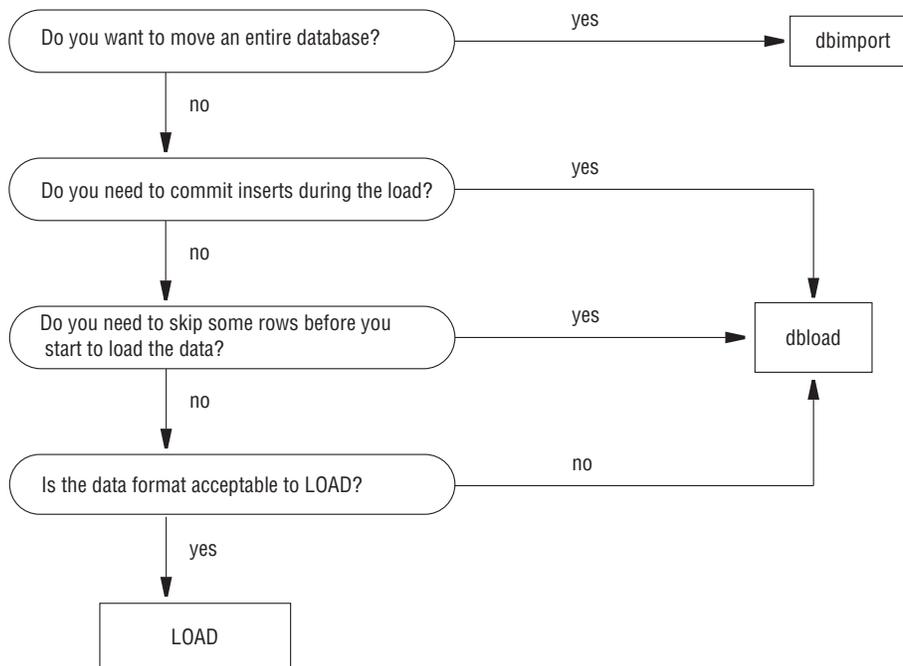


Figure 2-1. Choosing Among dbimport, dbload, and LOAD

In addition to the tools that move data, you can use the **dbschema** utility, which gets the schema of a database and redirects the output to a file, so you can provide the file to DB–Access to re-create the database.

Related concepts

Chapter 13, “The onunload and onload Utilities,” on page 13-1

This chapter contains information about using the **onunload** and **onload** utilities to unload and load databases and tables from and into IBM Informix database servers.

Chapter 9, “The dbload Utility,” on page 9-1

This chapter contains information on using the **dbload** utility, which loads data into databases or tables that IBM Informix products created.

Chapter 10, “The dbschema Utility,” on page 10-1

This chapter contains information about the **dbschema** utility and how to use it.

Chapter 11, “The LOAD and UNLOAD Statements,” on page 11-1

This chapter shows the syntax of the SQL UNLOAD and LOAD statements.

“High-Performance Loader Performance Advantages for Large Databases”

The High-Performance Loader (HPL) utility, which can load data from any ASCII or COBOL file that meets certain format prerequisites, uses parallel processing to perform fast data loading and unloading. However, the HPL requires significant preparation time.

High-Performance Loader Performance Advantages for Large Databases

The High-Performance Loader (HPL) utility, which can load data from any ASCII or COBOL file that meets certain format prerequisites, uses parallel processing to perform fast data loading and unloading. However, the HPL requires significant preparation time.

For extremely large databases, the HPL has a performance advantage over other IBM Informix data-migration utilities because it performs I/O and code-set

conversions in parallel. Use the HPL only for large databases, for which the time savings in the actual loading or unloading of data makes the preparation time worthwhile.

The following HPL features provide powerful tools for handling data from non-Informix sources:

- Drivers to handle different database types
- Filters and functions to manipulate data
- Code-set conversion
- The **ipload** GUI for UNIX
- The **onpladm** command-line utility for UNIX and Windows

For more information about the HPL, refer to the *IBM Informix High-Performance Loader User's Guide*.

Use Nonlogging Raw Tables with Utilities to Speed Up Data Loading Time

If you use a Dynamic Server utility to load data, you can use nonlogging raw tables in a logging database to speed up the initial loading and validation of data if you are moving data to or from Version 9.21 or any later version.

Data warehousing and other applications can have very large tables that take a long time to load. Nonlogging tables are faster to load than logging tables.

In a logged database, Dynamic Server creates standard tables that use logging by default.

To create a nonlogging table, use the CREATE RAW TABLE statement, or use the ALTER TABLE statement to change the table type from STANDARD to RAW. After the loading of a raw table is complete, you can change the table to a logging table (in a logging database) by changing the table type to STANDARD. Then you can use ALTER TABLE statements to add referential constraints to the table and CREATE INDEX statements to add indexes. For more information on these SQL statements, see the *IBM Informix Guide to SQL: Syntax*.

To load raw tables, you can use any data loading utility, such as **dbimport** or HPL in express mode. After you load data, perform a level-0 backup. Before you modify any data in a raw table or use it in a transaction, change the table type to STANDARD.

If an error or failure occurs during the loading of a raw table, the resulting data is whatever was on the disk at the time of the failure.

The **dbexport** and **dbschema** utilities support the CREATE RAW TABLE and ALTER TABLE...TYPE (RAW) statements.

For more information on nonlogging tables, see your *IBM Informix Administrator's Guide*. For more information on how to improve the performance of loading very large tables, see your *IBM Informix Performance Guide*. For more information on the ALTER TABLE statement, see the *IBM Informix Guide to SQL: Syntax*.

TEXT and BYTE Data Scanned, Not Compressed

An Informix database server scans TEXT and BYTE data into an existing table when you load data by using the SQL LOAD statement, the **dbload** utility, the Informix ESQ/C program, the HPL, or external tables.

Informix database servers do not have any mechanisms for compressing TEXT and BYTE data after the data has been scanned into a database.

Moving Non-Informix Data Between Computers and Dbspaces

Except when you use the HPL or external tables, you must unload your data to ASCII files before you move the data to another computer.

If you are moving to an Informix database server on another computer, you can use the **dbimport** and **dbload** utilities to load the data that you exported.

If you are moving data to a non-Informix application, you might need to use the UNLOAD statement because you can specify the delimiter that is used in the data files.

Importing Data from a Non-Informix Source

The **dbimport** and **dbload** utilities can import data from any ASCII file that is properly formatted.

Most applications that produce data can export the data into files that have a suitable format for **dbimport**. If the format of the data is not suitable, use UNIX, Linux, or Windows utilities to reformat the data before you import it.

In addition to **dbimport** and **dbload**, the IBM Informix Enterprise Gateway products and the HPL provide ways to access information from non-Informix sources.

Importing Data with IBM Informix Enterprise Gateway Products

You can use IBM Informix Enterprise Gateway with DRDA[®] to query a DRDA database and then insert the results into an Informix database.

For example, to import data, run a SELECT statement to select data from the non-Informix database and then run an INSERT statement to insert data into the Informix database. For more information, refer to the *IBM Informix Enterprise Gateway with DRDA User Manual*.

IBM Informix Enterprise Gateway provides a single, standards-based gateway to multiple data sources. Gateway Manager connects the Informix environment with that of any shared-library ODBC Level 2-compliant driver manager and driver on UNIX or Linux. For instance, you can use Gateway Manager with the IBM Informix Enterprise Gateway driver products to access UNIX or Linux database server products. For more information, refer to the *IBM Informix Enterprise Gateway Manager User Manual*.

Overview of Moving Data by Using Distributed SQL

If you want to move data with different binary pages and page sizes across platforms and you have expertise in using distributed SQL, you can use INSERT and SELECT SQL statements to transfer the data.

Important: Do not use INSERT and SELECT statements to move data if the database contains BLOB data types.

Prerequisites: A network connection must exist between database server instances.

To move data using INSERT and SELECT statements with fully qualified table names:

1. Capture the complete database schema from the source database server.
2. Alter the extent sizing and, if necessary, the lock modes on tables from page to row.
3. Create and verify the schema on the target database server.
4. Disable logging on both source and target servers where necessary.
5. Create and run the following scripts:
 - a. Create and run separate scripts for:
 - Disabling select triggers on the source server
 - Disabling indexes, triggers and constraints for each table on the target database server.
 - b. Create and run one script per table for the fully-qualified INSERT and SELECT statements.

For example:

```
INSERT INTO dbname@target:owner.table SELECT *  
FROM dbname@source:owner.table
```

You can run the scripts in parallel. In addition, for larger tables, you can create multiple scripts that can partition the table to run in parallel.

- c. Create and run separate scripts for enabling indexes, triggers and constraints for each table
6. Run UPDATE STATISTICS on system catalog tables and stored procedures and functions on the target database server.
7. Adjust starting values for all tables that have serial columns on the target database server.
8. Turn on transaction logging on the source and target database servers.
9. Return the source and target database servers to multi-user mode.
10. Validate the data that was transferred to the target database server.

For information on INSERT and SELECT statements, refer to the *IBM Informix Guide to SQL: Syntax*. For information on distributed transactions, refer to the *IBM Informix Administrator's Guide* and the *IBM Informix Dynamic Server Administrator's Reference*.

Part 2. Migration to and Reversion from Version 11.50

Chapter 3. Preparing for Migration to Dynamic Server Version 11.50

This chapter contains information about the tasks that you must perform before you install the new version of Dynamic Server.

This chapter includes:

- A list of all required pre-migration tasks with links to subtopics containing more information on the tasks.
- A checklist of performance and diagnostic information that you can gather before you migrate to another version of Dynamic Server.
- Additional considerations if you are migrating from 32-bit to 64-bit database servers

Preparing for Migration

Preparing for migration includes gathering information about and backing up your data, so that you can reinstall the previous version of the server and restore your data if you have a migration problem.

Prerequisites:

- Check the release notes for information about the new version of Dynamic Server. Also refer to the following topics in this guide:
 - Appendix A, “New Environment Variables,” on page A-1
 - Appendix B, “New Configuration Parameters,” on page B-1
 - Appendix C, “Configuration Parameters That Have Been Changed or Removed,” on page C-1
 - Appendix D, “New Reserved Words,” on page D-1
 - Appendix E, “System Catalog and System Database Changes,” on page E-1
- Check the machine notes for information about the correct operating-system release and any patches that you need for successful installation and operation of the database server.
- On UNIX or Linux, plan to retain both versions of the IBM Informix product software on disk, if you have enough disk resources, so you can revert to the source version. You cannot retain both servers.
- If you use Enterprise Replication, perform Enterprise Replication additional preparation tasks before you complete the tasks in this topic. For more information, see Chapter 4, “Migrating to and Reverting from Dynamic Server Version 11.50 with Enterprise Replication,” on page 4-1
- If the source version of the database server contains the `IFX_EXTEND_ROLE` configuration parameter, which controls authorization to register DataBlade[®] modules or external user-defined routines (UDRs), disable the parameter by setting it to 0 (off).
- If you are upgrading from Dynamic Server Version 9.21 and you have columns for smart large objects (BLOB data types) that were defined with a default value of an empty string, you must redefine the columns to have NULL as the default value.
- If you are upgrading from Dynamic Server Version 7.3, you must drop all stored procedures that have embedded CREATE TRIGGER statements, CREATE TABLE

statements with fragmentation expressions, ALTER TABLE or CREATE TABLE statements with constraint expressions, and nested CREATE PROCEDURE statements. After migrating to the current database server, you must recreate all stored procedures that were dropped.

- If you are migrating from Dynamic Server 7.31, the number of columns of the VARCHAR or NVARCHAR data type per table for Dynamic Server 11.50 has been reduced from 231 to 195, within a row size of 32762 bytes and based on a page size of 2K on UNIX or LINUX. The same 195-column restriction also applies to BYTE and TEXT columns.

On platforms where the page size is 4K (Windows and AIX®), the limit for the number of columns limit is approximately 450 columns.

Optionally, you can use the checklist in “Checklist of Diagnostic Information for Performance and Troubleshooting” on page 3-8 for gathering additional performance information that can be useful if you have large or complex applications. This additional information is also useful if you need to troubleshoot problems or issues after migration and need help from Technical Support.

To prepare for migration to Version 11.50 from Version 11.10, 10.0, 9.40, 9.30, 9.21 or 7.31

1. Check and configure available space to be sure you have enough space to move data and for any other software and network tools that you use. See “Check and Configure Available Space.”
2. Save copies of the current configuration files. See “Save Copies of the Current Configuration Files” on page 3-4.
3. Save a copy of your current Storage Manager **sm_versions** file. See “Save a Copy of the Storage Manager sm_versions File” on page 3-5.
4. Close all transactions and shut down the source database server. See “Close All Transactions and Shut Down the Source Database Server” on page 3-5.
5. Initiate fast recovery to verify that no open transactions remain after you shut down the source database server. See “Initiate Fast Recovery to Verify that No Open Transactions Exist” on page 3-5.
6. Verify the integrity of the data. See “Verify the Integrity of the Data” on page 3-6.
7. Verify that the database server is in quiescent mode. See “Verify that the Database Server Is in Quiescent Mode” on page 3-7.
8. If you use High-Availability Data Replication (HDR), disable it. See “Disable High-Availability Data Replication” on page 3-7.
9. Make a final level-0 backup of the source database server, including all storage spaces, in case you need to revert to the source database server. See “Make a Final Backup of the Source Database Server” on page 3-7.
10. Run the **ontape -a** command after the backup is complete.
11. Verify that the source database server is offline. See “Verify That the Source Database Server Is Offline” on page 3-7.
12. On UNIX or Linux only, modify kernel parameters. See “Modify Kernel Parameters (UNIX, Linux)” on page 3-8.

Check and Configure Available Space

Before you migrate to Version 11.50, you must make sure that you have enough available space for the new server, your data, and any other network and data tools that you use.

During migration, Dynamic Server drops and then recreates the **sysmaster** database. Depending on which version of Dynamic Server you migrate from, the **sysmaster** database in the current version can be significantly larger.

UNIX/Linux Only

Dynamic Server Version 11.50 requires 3000 free pages of logical-log space (approximately 6000 kilobytes for a 2 KB page size) to build the **sysmaster** database on UNIX or Linux.

Windows Only

Dynamic Server Version 11.50 requires 1500 to 3000 free pages of logical-log space (approximately 6000 kilobytes for a 4 KB page size) to build the **sysmaster** database on Windows.

During migration, a second database, the **sysadmin** database, is created in the **root** dbspace. As you work after migrating, the **sysadmin** database, could grow dramatically. You can move the **sysadmin** database to a different dbspace.

You might need to increase the physical log size to accommodate new features, and you might consider adding a new chunk.

Partition header pages should not be full; key descriptors and other new features might require more space after migration to Dynamic Server Version 11.50. If partition header pages are full, you can merge extents in the partition to create some free space for the partition header. You can also unload the table, recreate the table to reduce the large number of extents, and then reload the table.

The root chunk should contain at least ten percent free space when converting to Dynamic Server Version 11.50.

In some cases, even if the database server migration is successful, internal conversion of some databases might fail because of insufficient space for system catalog tables. For more information, see the release notes for this version of Dynamic Server.

Add any additional free space to the system prior to the migration. If the dbspaces are nearly full, add space before you start the migration procedure. When you start Dynamic Server Version 11.50 on the same root dbspace of the earlier database server, Dynamic Server Version 11.50 automatically converts the **sysmaster** database and then each database individually. For a successful conversion of each database, ensure that 2000 KB of free space per database is available in each dbspace.

To ensure enough free space is available

1. Calculate the amount of free space that each dbspace requires.

In the following equation, n is the number of databases in the dbspace and X is the amount of free space they require:

$$X \text{ kilobytes free space} = 2000 \text{ kilobytes} * n$$

The minimum number of databases is 2 (for the **sysmaster** and **sysadmin** databases).

2. Check the amount of free space in each dbspace to determine whether you need to add more space.

You can run SQL statements to determine the free space that each dbspace requires and the free space available. These statements return the free-space calculation in

page-size units. The **free_space_req** column value is the free-space requirement, and the **free_space_avail** column value is the free space available.

The following SQL statement shows how to determine the free space that each dbspace requires:

```

DATABASE sysmaster;
SELECT partdbsnum(partnum) dbspace_num,
       trunc(count(*) * 2000) free_space_req
  FROM sysdatabases
 GROUP BY 1
 ORDER BY 1;

```

The following SQL statement queries the **syschunks** table and displays the free space available for each dbspace:

```

SELECT dbsnum dbspace_num, sum(nfree) free_space_avail
  FROM syschunks
 GROUP BY 1
 ORDER BY 1;

```

Important: If less free space is available than the dbspace requires, either move a table from the dbspace to another dbspace or add a chunk to the dbspace.

The dbspace estimates could be higher if you have an unusually large number of SPL routines or indexes in the database.

Save Copies of the Current Configuration Files

Save copies of the configuration files that exist for each instance of your source database server. Keep the copies available in case you decide to use the files after migrating or you need to revert to the source database server.

Although you can use an old ONCONFIG configuration file with Version 11.50, you should use the new Version 11.50 ONCONFIG file, or at least examine the file for new parameters. For information on Version 11.50 changes to the ONCONFIG file, see Appendix C, “Configuration Parameters That Have Been Changed or Removed,” on page C-1.

Configuration files that you might have are listed in Table 3-1.

Table 3-1. Configuration Files to Save from the Source Database Server

UNIX or Linux	Windows
\$INFORMIXDIR/etc/\$ONCONFIG	%INFORMIXDIR%\etc\%ONCONFIG%
\$INFORMIXDIR/etc/onconfig.std	%INFORMIXDIR%\etc\onconfig.std
\$INFORMIXDIR/etc/oncfg*	%INFORMIXDIR%\etc\oncfg*
\$INFORMIXDIR/etc/sm_versions	%INFORMIXDIR%\etc\sm_versions
\$INFORMIXDIR/aaodir/adtcfg	%INFORMIXDIR%\aaodir\adtcfg.*
\$INFORMIXDIR/dbssodir/adtmasks	%INFORMIXDIR%\dbssodir\adtmasks.*
\$INFORMIXDIR/etc/sqlhosts	
\$INFORMIXDIR/etc/tctermcap	
\$INFORMIXDIR/etc/termcap	

If you use ON-Bar to back up your source database server and the logical logs, you must also save a copy of any important storage manager files as well as the following file:

UNIX or Linux:
\$INFORMIXDIR/etc/ixbar.servernum

Windows:
%INFORMIXDIR%\etc\ixbar.servernum

Save a Copy of the Storage Manager `sm_versions` File

Before you migrate to a later version of the database server, save a copy of your current `sm_versions` file, which should be in the `$INFORMIXDIR/etc` directory.

If you are using a different directory as `INFORMIXDIR` for the new database server, copy `sm_versions` to the new `$INFORMIXDIR/etc`, or copy `sm_versions.std` to `sm_versions` in the new directory, and then edit the `sm_versions` file with appropriate values before starting the migration.

For information on how to install and use the Storage Manager, see the *IBM Informix Storage Manager Administrator's Guide*.

Close All Transactions and Shut Down the Source Database Server

Before migrating, terminate all database server processes and shut down your source database server. This lets users exit and shuts down the database server gracefully. If you have long running sessions, you must also shut those down.

Inform client users that migration time is typically five to ten minutes. However, if migration fails, you must restore from a level-0 backup, so ensure that you include this possibility when you estimate how long the server will be offline.

Before you migrate from the original source database server, make sure that no open transactions exist. Otherwise, fast recovery will fail when rolling back open transactions during the migration.

To let users exit and shut down the database server gracefully

1. Run the `onmode -sy` command to put the database server in quiescent mode.
2. Wait for all users to exit.
3. Run the `onmode -l` command to move to the next logical log.
4. Run the `onmode -c` to force a checkpoint.
5. Make a level-0 backup of the database server.
6. Run the `ontape -a` command after the level-0 backup is complete.
7. Run the `onmode -yuk` command to shut down the system.

If you need to perform an immediate shutdown of the database server, run these commands:

```
onmode -l  
onmode -c  
onmode -ky
```

Initiate Fast Recovery to Verify that No Open Transactions Exist

A shutdown procedure does not guarantee a rollback of all open transactions. To guarantee that the source database server has no open transactions, put the source database server in quiescent mode and initiate fast recovery.

Run the following command to enter quiescent mode and initiate a fast recovery:

```
oninit -s
```

UNIX/Linux Only

On UNIX or Linux, the **oninit -s** command rolls forward all committed transactions and rolls back all incomplete transactions since the last checkpoint and then leaves a new checkpoint record in the log with no open transactions pending.

You must run the **oninit -s** command before you initialize Dynamic Server Version 11.50. If any transactions remain when you try to initialize the new database server, you will receive the following error when you try to initialize the new database server, and it goes offline:

```
Open transaction detected when changing log versions.
```

For more information about fast recovery, see your *IBM Informix Dynamic Server Administrator's Guide*.

After you put the database server in quiescent mode and initiate fast recovery, issue the **onmode -yuk** command to shut down the database server. Then review the **online.log** file for any possible problems and fix them.

Only after proper shutdown can you bring the new database server (Dynamic Server Version 11.50) through the migration path. Any transaction that is open during the migration will cause an execution failure in fast recovery.

Verify the Integrity of the Data

After verifying that no open transactions exist, verify the integrity of your data by running the **oncheck** utility. You can also verify the integrity of the reserve pages, extents, system catalog tables, data, and indexes. If you find any problems with the data, fix the problems before you make a final backup of the source database server.

To obtain the database names, use the following statements with DB–Access:

```
DATABASE sysmaster;  
SELECT name FROM sysdatabases;
```

Alternatively, to obtain the database names, run the **oncheck -cc** command without any arguments and filter the result to remove unwanted lines, as shown in this example:

```
oncheck -cc | grep "ting database"
```

Table 3-2 lists the **oncheck** commands that verify the data integrity.

Table 3-2. Commands for Verifying the Data Integrity

Action	oncheck Command
Check reserve pages	oncheck -cr
Check extents	oncheck -ce
Check system catalog tables	oncheck -cc <i>database_name</i>
Check data	oncheck -cD <i>database_name</i>
Check indexes	oncheck -cI <i>database_name</i>

For information about the **oncheck** commands, see the *IBM Informix Dynamic Server Administrator's Guide*.

Verify that the Database Server Is in Quiescent Mode

Before you make a final backup, verify that your source database server is in quiescent mode.

Run the following command to verify that the database server is in quiescent mode:

```
onstat -
```

The first line of the onstat output shows the status of your source database server:

```
IBM Informix Dynamic Server Version X.XX.XXX -- Quiescent -- Up
```

Disable High-Availability Data Replication

If you use High-Availability Data Replication (HDR), disable it before you migrate to Dynamic Server Version 11.50.

To disable HDR, set the primary database server to standard with the following command on the primary database server:

```
onmode -d standard
```

You do not need to perform the whole conversion process on an HDR secondary database server. For more information, see “Prepare HDR Secondary Servers” on page 5-9.

For more information on HDR, see the *IBM Informix Dynamic Server Administrator's Guide*.

Make a Final Backup of the Source Database Server

Use ON-Bar or **ontape** to make a level-0 backup of the source database server, including all storage spaces and all used logs. After you make a level-0 backup, also perform a complete backup of the logical log, including the current logical-log file.

Be sure to retain and properly label the tape volume that contains the backup.

Important: You must also make a final backup of each source database server instance that you plan to convert.

For ON-Bar, remove the **ixbar** file, if any, from the **\$INFORMIXDIR%/etc** or **%INFORMIXDIR%\etc** directory after the final backup. Removing the **ixbar** file ensures that backups for the original source database server are not confused with backups about to be done for the new database server. Follow the instructions regarding expiration in your storage manager documentation.

For more information about making backups, see the *IBM Informix Backup and Restore Guide*.

Verify That the Source Database Server Is Offline

Before you install the new database server, verify that the source database server is offline. You must do this because the new database server uses the same files.

You cannot install the new database server if any of the files that it uses are active.

You can also use the **onstat** utility to determine that shared memory was not initialized.

Modify Kernel Parameters (UNIX, Linux)

You might need to change some of the kernel parameters for your UNIX or Linux operating system before you install Dynamic Server Version 11.50.

To reconfigure the operating system, follow the directions in both of these resources:

- Machine notes file included on your database server distribution media
- Kernel-configuration instructions for your operating system

Checklist of Diagnostic Information for Performance and Troubleshooting

Before you migrate to a newer version of Dynamic Server, you can gather diagnostic information. The information will be useful if you have large, complex applications or if you have problems or issues after migration and need help from Technical Support.

If you have problems, you or Technical Support can compare the information that you gather with information obtained after migration.

The following Table 3-3 contains a list of the diagnostic information that you can gather. You can print the checklist. Then, after you get the information specified in each row, check the second column of the row.

Table 3-3. Checklist of Information to Get Before Migrating

Information to Get Before Migrating	Done
Get the SQL query plans for all regularly used queries, especially complex queries, by using SET EXPLAIN ON.	
Run the dbschema -d -hd command for all critical tables.	
The output contains distribution information.	
Get oncheck -pr output that dumps all of the root reserved pages.	
Make a copy of the ONCONFIG configuration file.	
A copy of the ONCONFIG file is essential if you need to revert to an earlier version of Dynamic Server. In addition, a copy of this file is useful because oncheck -pr does not dump all of the configuration parameters.	
Prepare a list of all the environment variables that are set using the env command.	
During times of peak usage: <ul style="list-style-type: none">• Obtain an online.log snippet, with some checkpoint durations in it• Run onstat -aF, -g all, and -g stk all.	

Table 3-3. Checklist of Information to Get Before Migrating (continued)

Information to Get Before Migrating	Done
<p>During times of peak usage, run the following onstat commands repeatedly with the -r repeat option for a period of about three to five minutes:</p> <ul style="list-style-type: none"> • onstat -u, to see the total number of sqlexecs used • onstat -p, for read and write cache rates, to detect deadlocks and the number of sequential scans • onstat -g nta, a consolidated output of -g ntu, ntt, ntm and ntd • onstat -g nsc, -g nsd, and -g nss for the status of shared memory connections • onstat -P, -g tpf, and -g ppf • vmstat, iostat and sar, for cpu utilization • timex of all queries that you regularly run 	

Migrating from 32-bit to 64-Bit Database Servers

If you are migrating from a 32-bit version of Dynamic Server to a 64-bit version of Dynamic Server or reverting from a 64-bit version of Dynamic Server, you might need to follow additional steps to update certain internal tables.

These steps are documented in the platform-specific machine notes that are provided with your database server.

For 32- to 64-bit migrations, change **SHMBASE** and **STACKSIZE** according to the **onconfig.std** configuration file for the new version.

All UDRs and DataBlade modules that were built in 32-bit mode must be recompiled in 64-bit mode because they will not work with the 64-bit database server. If you have any UDRs that were developed in 32-bit mode, make sure that proper size and alignment of the data structures are used to work correctly on a 64-bit computer after recompiling in 64-bit mode. For more information, refer to the machine notes.

Chapter 4. Migrating to and Reverting from Dynamic Server Version 11.50 with Enterprise Replication

This chapter describes the additional tasks that you must perform when migrating to and reverting from Dynamic Server Version 11.50 if you are running Enterprise Replication.

The IBM Informix Dynamic Server Express Edition does not support Enterprise Replication. If you are migrating to the Express Edition and you currently use Enterprise Replication, see “Express Edition Migration and Enterprise Replication” on page 4-4.

Preparing to Migrate to Version 11.50 with Enterprise Replication

If you use Enterprise Replication, you must complete additional tasks before you prepare for migration.

Prerequisites: You must perform all migration operations as user **informix**.

To prepare for migration with Enterprise Replication

1. Remove all replicate groups that Dynamic Server Version 11.50 does not support.
2. Stop applications that are performing replicable transactions.
3. Make sure that control and TRG send queues are empty:
 - Run **onstat -g grp** to ensure that the Enterprise Replication grouper does not have any pending transactions. The grouper evaluates the log records, rebuilds the individual log records into the original transaction, packages the transaction, and queues the transaction for transmission.
 - Run **onstat -g rqm** to check for queued messages.
4. Replace old event class identifiers (IDs) used in the ALARMPROGRAM with the new event class IDs.
5. Shut down Enterprise Replication with the following command:

```
cdr stop
```

Now you can complete the steps in “Preparing for Migration” on page 3-1 and, if necessary, in “Migrating from 32-bit to 64-Bit Database Servers” on page 3-9.

Migrating to Version 11.50 with Enterprise Replication

If you use Enterprise Replication, you must complete additional tasks when you migrate to Dynamic Server Version 11.50.

Prerequisites:

- Complete the steps in “Preparing to Migrate to Version 11.50 with Enterprise Replication.”
- Complete the steps in “Preparing for Migration” on page 3-1
- Perform all migration operations as user **informix**.

To migrate to Dynamic Server Version 11.50 with Enterprise Replication

1. Perform the tasks described in “Migrating to Version 11.50” on page 5-1, including starting Dynamic Server Version 11.50.
2. If the CDR_QDATA_SBSpace configuration parameter is not set, set it in the ONCONFIG file to the sbspaces for Enterprise Replication to use for storing spooled row data.
3. For each node involved in Enterprise Replication, back up the **syscdr** databases by using the **dbexport -ss** command or the **dbschema -ss** command and the UNLOAD statement, or by a combination of these methods. The **-ss** option prevents backup tables from using default extent sizes and row-level locking, which is not an appropriate lock mode with Enterprise Replication.
4. Make sure that no replicable transactions occur before Enterprise Replication starts.
5. Run the conversion script, named **concdr.sh**, in the **\$INFORMIXDIR/etc/conv** directory on UNIX, or **concdr.bat**, in the **%INFORMIXDIR%\etc\conv** directory on Windows. For example, specify:


```
% sh concdr.sh from_version 11.50
```

 Valid *from_version* values are: 11.10, 10.00, 9.40, 9.30, 9.21, 9.20, and 7.31.
6. Wait for one of the following messages:


```
'syscdr' conversion completed successfully.
'syscdr' conversion failed.
```

 For details about the conversion, see either of the following files:
 - **\$INFORMIXDIR/etc/concdr.out**
 - **%INFORMIXDIR%\etc\concdr.out**
7. If conversion fails, resolve the problem reported in the **concdr.out** file, restore the **syscdr** database from a backup, and then attempt conversion again.
8. After successful conversion, start Enterprise Replication:


```
% cdr start
```

Important: After you convert to Dynamic Server Version 11.50 with Enterprise Replication from Dynamic Server 9.30, 9.21, or 7.31, do not drop the **syscdr** database. If **syscdr** is dropped, you cannot revert to the older database server with Enterprise Replication because the data required to carry out the reversion is stored in the **syscdr** database.

Converting Replication of 9.21 User-Defined Data Types

If you are migrating from Dynamic Server Version 9.21 with Enterprise Replication, you must complete additional tasks if you have user-defined data types (UDTs).

Dynamic Server Version 9.21 has limited support for the replication of UDTs. To take advantage of Version 11.50 UDT replication, the user-defined routines (UDRs) for a UDT must contain **streamwrite()** and **streamread()** functions.

After you migrate to Dynamic Server Version 11.50, implement the **streamwrite()** and **streamread()** functions for any currently replicated UDTs on all database servers within the enterprise.

Reverting from Dynamic Server Version 11.50 with Enterprise Replication

If you use Enterprise Replication, you must complete additional tasks when you revert from Version 11.50.

During reversion to an earlier version of Dynamic Server with Enterprise Replication:

- Master replicates become standard replicates and tables that were added to the **syscdr** database are removed.
- Tables created with templates are dropped.
- The table containing replicated table-version information, which was created during conversion, is dropped.

The procedure for reverting to Dynamic Server 11.10, 10.00, 9.40 or 9.30 is slightly different than the procedure for reverting to Dynamic Server 9.21 or 7.31. To revert to Dynamic Server 7.31, the database server must be an Enterprise Replication root server (the uppermost level in a hierarchically organized set of database servers).

Prerequisites: Perform all reversion operations as user **informix**.

To revert to Dynamic Server 11.10, 10.00, 9.40 or 9.30 from Dynamic Server Version 11.50 with Enterprise Replication

1. Stop applications doing replicable transactions.
 2. You cannot revert if Enterprise Replication is in Alter mode, so make sure Enterprise Replication is not in Alter Mode. Use **onstat -g cat repls** to see if Enterprise Replication is in Alter mode. If it is in Alter mode, change the mode.
 3. Delete shadow replicates.
 4. Make sure that control and TRG send queues are empty:
 - Run **onstat -g grp** to ensure that the Enterprise Replication grouper does not have any pending transactions.
 - Run **onstat -g rqm** to check for queued messages.
 5. Shut down Enterprise Replication with the following command:
`cdr stop`
 6. Back up the **syscdr** databases with **dbschema** or **UNLOAD**.
 7. Run the reversion script, named **revcdr.sh**, in the **\$INFORMIXDIR/etc/conv** directory on UNIX, or **revcdr.bat**, in the **%INFORMIXDIR%\etc\conv** directory on Windows, as shown in this example:

```
% sh revcdr.sh 11.50 9.40
```
- Valid *to_version* values are 11.10, 10.00, 9.40, and 9.30. This script does a reversion test followed by the actual Enterprise Replication reversion.
8. If the reversion test or actual reversion fails, check the file **\$INFORMIXDIR/etc/revtestcdr.out** or **revcdr.out**, respectively. Attempt reversion after resolving problems reported.
 9. Perform database server reversion tasks, as described in “Reverting from Dynamic Server Version 11.50” on page 6-10.
 10. Run **onmode -I** and **onmode -c**. If you do not do this after reversion and before starting Enterprise Replication, the database server might fail when you start Enterprise Replication.
 11. Start Enterprise Replication:

```
% cdr start
```

To revert to Dynamic Server 9.21 or 7.31 from Dynamic Server Version 11.50 with Enterprise Replication

1. Drop all replicate sets.

2. Drop any replicate that contains a smart large object or user-defined data type.
3. Stop applications doing replicable transactions.
4. Make sure that control and TRG send queues are empty:
 - Run **onstat -g grp** to ensure that the Enterprise Replication grouper does not have any pending transactions.
 - Run **onstat -g rqm** to check for queued messages.
5. Shut down Enterprise Replication with the following command:

```
cdr stop
```
6. Back up the **syscdr** databases with **dbschema** or UNLOAD.
7. Run the reversion script, named **revcdr.sh**, in the **\$INFORMIXDIR/etc/conv** directory on UNIX, or **revcdr.bat**, in the **%INFORMIXDIR%\etc\conv** directory on Windows:

```
% sh revcdr.sh 11.50 to_version
```

 Valid *to_version* values are 9.21, 9.20, and 7.31.
 This script does a reversion test followed by the actual Enterprise Replication reversion.
8. If the reversion test or actual reversion fails, check the file **\$INFORMIXDIR/etc/revtestcdr.out** or **revcdr.out**, respectively. Attempt reversion after resolving the problems reported.
9. Drop the Enterprise Replication sbspaces.
10. Perform database server reversion tasks, as described in “Reverting from Dynamic Server Version 11.50” on page 6-10.
11. Run **onmode -I** and **onmode -c**. If you do not do this after reversion and before starting Enterprise Replication, the database server might fail when you start Enterprise Replication.
12. Start Enterprise Replication:

```
% cdr start
```

Express Edition Migration and Enterprise Replication

IBM Informix Dynamic Server Express does not support Enterprise Replication. Before migrating your database server to Dynamic Server Express, drop existing Enterprise Replication configuration using the **cdr delete server** command option.

After upgrading to Dynamic Server Express, Enterprise Replication command-line interface options will not be available. Any remote Enterprise Replication clients trying to connect to an Express Edition server will receive an error stating that Enterprise Replication is not supported in the Express Edition server.

The Express Edition server does not support any of the Enterprise Replication configuration parameters. Examples of these configuration parameters are **CDR_ENV**, **CDR_SERIAL**, and **ENCRYPT_CDR**.

If you did not delete Enterprise Replication before converting to Dynamic Server Express and then revert the database server to the Enterprise Edition, the integrity of the **syscdr** database is not restored. In this case, you must drop the **syscdr** database. You can drop the **syscdr** database with the **cdr remove** command and by issuing **cdr delete serv** at other nodes in the topology.

Chapter 5. Migrating to Dynamic Server Version 11.50

This chapter contains the procedures for migrating to Dynamic Server Version 11.50 and for completing required post-migration tasks.

The process for migrating to the Enterprise Edition and the Express Edition is the same, except that some particular tasks are not applicable for the Express Edition, because the Express Edition does not support Enterprise Replication and High-Availability Data Replication (HDR).

Migrating to Version 11.50

After you prepare your databases for migration, you can migrate to the new version of Dynamic Server.

Recommendation: Install a new version of the database server in a new directory and then test a database server instance with the similar configuration settings and the same `ssqlhosts` information that you used for your old database server. If you do not have space on your machine to accommodate the new Dynamic Server installation, back up your old directory and then remove it.

You can also migrate on a database server dedicated to testing your migration to Version 11.50 before you migrate on your production database server.

Prerequisites:

- Read the release notes and the machine notes for any new information.
- Complete the steps in “Preparing for Migration” on page 3-1
- Refer to the *IBM Informix Installation Guide* for detail installation prerequisites, options, and procedures.

Important: Do not connect applications to a database server instance until migration has successfully completed.

To migrate to Dynamic Server Version 11.50

1. Install Dynamic Server Version 11.50 in a new directory. Do **not** install the new database server over the old database server. If you are migrating from a Dynamic Server version that does not support the custom installation of components to a version that does support it, you must choose the Typical installation option. After the typical components are installed, you can selectively remove the components that you do not need. For more information, see “Install Dynamic Server Version 11.50” on page 5-2.
2. Set environment variables. For more information, see “Set Environment Variables” on page 5-4.
3. Adjust configuration parameters as necessary. If the ALARMPROGRAM configuration parameter is set to the script `alarmprogram.sh`, set the value of BACKUPLOGS in `alarmprogram.sh` to N. For more information, see “Customize Configuration Files” on page 5-4.
4. Optionally add Communications Support Modules. See “Optionally Add Communications Support Modules” on page 5-5.
5. Optionally upgrade DataBlade modules to correspond to the newer server. See “Install or Upgrade Any DataBlade Modules” on page 5-5.

6. Switch to user **informix** and initialize Dynamic Server to trigger the migration. For more information, see “Initialize Dynamic Server Version 11.50” on page 5-5. When the migration starts, the **online.log** displays the message “Conversion from version <version number> Started.” The log continues to display start and end messages for all components. When the migration of all components are complete, the message “Conversion Completed Successfully” appears. For more information on this log, see “Migration Status Messages in online.log” on page 5-3.
7. If you successfully migrated to the new server, see “Completing Required Post-Migration Tasks” on page 5-6 for information on preparing the new server for use. If you successfully migrated, but the conversion of the High-Performance Loader **onpload** database failed, upgrade the **onpload** database. For more information, see “Upgrade the High-Performance Loader onpload Database” on page 5-6.
8. If the log indicates that migration failed, you must install the old database server and restore your database from a level-0 backup.

If you are migrating the database server from a version that does not support label-based access control, users who held the DBA privilege are automatically granted the SETSESSIONAUTH access privilege for PUBLIC during the migration process. For more information on SETSESSIONAUTH, see the *IBM Informix Guide to SQL: Syntax*. For information on label-based access control, see the *IBM Informix Security Guide*.

Install Dynamic Server Version 11.50

Install and configure the current version of Dynamic Server. On UNIX or Linux, you must be logged in as user **root**. On Windows, you must be a member of the Administrators group.

Important: Do not install the new version over the existing product. Preserve your old database server files by installing the new database server in a different directory. If you do not have space on your machine for the new Dynamic Server installation, back up your old directory and then remove it. If you install the new database server in the same directory where the old database server resided, the installation script overwrites the older files.

When you change from the old server to the new server, change the **INFORMIXDIR** environment variable to ensure that it points to the location of the installed database server. If you do not do this, the older version of the database server will start up when you reboot.

If you install a new instance into a directory that already contains an instance and there is insufficient free space, the install program will request that you confirm removal of the older product prior to extracting the new one. Otherwise, if sufficient free space exists, the install program will add or replace files without deleting the existing instance.

During subsequent installations, the behavior of the install program depends on the target directory, as follows.

- If you install from the media into the same install location, only the files chosen for replacement or installation will be installed.
- If you are installing into a different location on the hard drive, you must verify that you have enough free disk space tests prior to file extraction.

After you install a new instance in a new directory, complete data migration and verification.

The installation script installs Dynamic Server into the **INFORMIXDIR** directory specified for user **root** on UNIX or Linux.

On Windows, the install application suggests a default **INFORMIXDIR** directory, which you can change by typing a different path. You do not need to create the directory before installation.

Important: Monitor the database server message log, **online.log**, for any error messages. If you see an error message, solve the problem before you continue the migration procedure.

Refer to the *IBM Informix Installation Guide* for detailed installation prerequisites, options, and procedures. Refer to the *IBM Informix Dynamic Server Administrator's Guide* for additional information on configuring the new server.

Tips for Installing on the Server on Windows

If you are installing on Windows:

1. Be sure to choose to install to a different directory.
2. Do not initialize the server when installing.
3. Copy the **ONCONFIG** file to the target database server and set parameters that are new for the current release.
4. Start the new server without initializing it. See information on initializing and starting the database server in your *IBM Informix Dynamic Server Installation Guide for Windows*.
5. Monitor the **online.log** for the Conversion Successful message.
6. Once the upgrade has completed successfully, you can remove the old instance. When you run the uninstaller make sure that you select this option:
Retains all databases, but removes server binaries.

Migration Status Messages in online.log

When the migration starts, the **online.log** displays the message "Conversion from version <version number> Started." The log continues to display start and end messages for all components.

When conversions of all components are complete, the message "Conversion Completed Successfully" displays. This message indicates that the migration process completed successfully, but it does not guarantee that each individual database was converted successfully. The message log could contain additional information regarding the success or failure of the migration of each individual database. If migration of a particular database fails, then try to connect to the database to find out the exact cause of the failure.

At the end of the migration of each individual database, Dynamic Server runs a script to update some system catalog table entries. The message log includes messages related to this script. The success or failure of the script does not prevent the usage of a database.

For migration to Dynamic Server Version 11.50 from Dynamic Server 7.31, if the script fails for a database, run the following script as user **informix** while connected to the database:

`$INFORMIXDIR/etc/dummyupds7x.sql`

For information about any messages in the message log, see the *IBM Informix Dynamic Server Administrator's Guide*.

Set Environment Variables

After you install the current version of Dynamic Server, verify that the **INFORMIXDIR**, **INFORMIXSERVER**, **ONCONFIG**, **PATH**, and **INFORMIXSQLHOSTS** (if used) environment variables are set to the correct values.

On UNIX or Linux, the client application looks for the **sqlhosts** file in the **etc** directory in the **INFORMIXDIR** directory. However, you can use the **INFORMIXSQLHOSTS** environment variable to change the location or name of the **sqlhosts** file.

Important: Before you start the Version 11.50 database server, you must set the **DBONPLOAD** environment variable to the name of the **pload** database if the name is not **onpload**, the default name.

For information about environment variables, see the *IBM Informix Guide to SQL: Reference*.

Customize Configuration Files

When you initialize Dynamic Server Version 11.50, use the same configuration that the old database server used. After you observe the performance of Dynamic Server Version 11.50, use the new **ONCONFIG** file to obtain the benefits from the new or changed configuration parameters.

Alternatively, you can examine the new file for new configuration parameters that you might want to use.

Set the **ALARMPROGRAM** configuration parameter to either nothing or **no_log.sh** to prevent the generation of errors if the logical log fills during the migration. For more details, see "Initialize Dynamic Server Version 11.50" on page 5-5. After the migration, change the value of **ALARMPROGRAM** to **log_full.sh**.

Important: To facilitate migration (and reversion), use the same values for your new database server for **ROOTOFFSET**, **ROOTSIZE**, and **ROOTPATH** that you used for the old database server. Also, keep the same size for physical logs and logical logs, including the same number of logical logs, and the same **sqlhosts** file.

If you use an optical storage manager, set the **OPTICAL_LIB_PATH** configuration parameter in the **ONCONFIG** file to the location of the optical storage manager library. For more information, see the *IBM Informix Optical Subsystem Guide*.

If you use custom-code files with the High-Performance Loader, set the **HPL_DYNAMIC_LIB_PATH** configuration parameter in the **plconfig** file to the location of the shared library. For Dynamic Server Version 11.50, the location of this library is `$INFORMIXDIR/lib/ipldd11a.SOLIBSUFFIX`, where **SOLIBSUFFIX** is the shared-library suffix for your operating system.

For information on how to configure Dynamic Server, see your *IBM Informix Dynamic Server Administrator's Guide*. For information about how to tune the configuration parameters, see the *IBM Informix Dynamic Server Performance Guide*.

Optionally Add Communications Support Modules

You can optionally use a Communications Support Module (CSM) with the current version of Dynamic Server. After you install the CSM components, create entries in the `concsm.cfg` file and in the options field of the `sqlhosts` file to configure the CSM.

Existing client applications do not need to be recompiled or relinked if your database server does not use CSMs. If your database server uses a CSM, client applications must relink with new Informix libraries. The client applications must install and configure the CSM.

For information on how to set up the CSM, see the *IBM Informix Dynamic Server Administrator's Guide*.

Install or Upgrade Any DataBlade Modules

After you install the new version of Dynamic Server, you might need to install or upgrade any DataBlade modules that you want to add to the database server.

Register the DataBlade modules after you initialize the database server.

Initialize Dynamic Server Version 11.50

After installing the new database server, perform shared-memory initialization to start the server. Do not perform disk-space initialization, which overwrites whatever is on the disk space.

Prerequisite: If you installed Dynamic Server as user `root`, you must switch to user `informix` before initializing the server.

Important: Dynamic Server writes to the logical logs with the transactions that result from creating the `sysmaster` database. If you run out of log space before the creation of the `sysmaster` database is complete, Dynamic Server stops and indicates that you must back up the logical logs. After you back up the logical logs, the database server can finish building the `sysmaster` database. You cannot use ON-Bar to back up the logical logs because the database has not been converted yet. If you have `ALARMPROGRAM` set to `log_full.sh` in the `ONCONFIG` configuration file, errors are generated as each log file fills during the migration. Set the value of `ALARMPROGRAM` to either nothing or `no_log.sh` so that these errors are not generated. If your logical log does fill up during the migration, you must back it up with `ontape`, the only backup tool you can use at this point. Issue the `ontape -a` command.

Bring Dynamic Server Version 11.50 online for the first time by executing `oninit` on UNIX or by using the `Service` control application on Windows. For more information, see the *IBM Informix Dynamic Server Administrator's Guide*.

As Dynamic Server Version 11.50 comes online for the first time, it modifies certain disk structures. This operation should extend the initialization process by only a minute or two. If your disks cannot accommodate the growth in disk structures, you will find a message in the message-log file that instructs you to run `oncheck` on a table. The `oncheck` utility will tell you that you need to rebuild an index. You must rebuild the index as instructed.

Upgrade the High-Performance Loader onpload Database

If **onpload** conversion failed during database server migration, you can manually upgrade the **onpload** database.

Starting with Version 9.40.xC3, Dynamic Server has a new version of the **onpload** database with longer column lengths. The **onpload** database now requires slightly more disk space than it did before Version 9.40.xC3.

When you migrate to a new version of Dynamic Server, you must also upgrade the **onpload** database.

To upgrade the onpload database

1. If you are upgrading from a version of Dynamic Server that is prior to Version 9.40, run the **conploadlegacy.sh** script, as shown in this example:

```
conploadlegacy.sh 7.31 9.40
```

2. If are upgrading from a version of Dynamic Server that is prior to Version 9.40xC3 and, if necessary, have already run the **conploadlegacy.sh** script, you must also perform one of the following tasks:

- Run the **conpload.sh** script, as shown in this example:

```
conpload.sh 9.40 11.50
```

- Set the new environment variable `IFX_ONPLOAD_AUTO_UPGRADE` to 1 for the upgrade to happen automatically the first time you start an HPL utility using the **ipload** or **onpladm** command, after you migrate to a new database server version. You cannot use the `IFX_ONPLOAD_AUTO_UPGRADE` environment variable with the **onpload** utility.

If you start an HPL utility before upgrading the **onpload** database, then you receive an error stating that the **onpload** database must be converted.

Starting with Dynamic Server Version 9.40.xC3, the **ipload** utility does not support object names that contain more than 18 characters. The utility will continue to operate properly if legacy applications do not use long object names.

Completing Required Post-Migration Tasks

After you migrate, you must complete a series of post-migration tasks to prepare the new version of the server for use.

To complete post-migration tasks

1. For ON-Bar, rename or edit the **sm_versions.std** file. For more information, see “For ON-Bar, Rename the sm_versions.std File” on page 5-7.
2. Run UPDATE STATISTICS. For more information, see “Update Statistics After Migrating” on page 5-7.
3. Review client applications and registry keys. For more information, see “Review Client Applications and Registry Keys” on page 5-8.
4. Verify the integrity of the data. For more information, see “Verify the Integrity of the Data” on page 5-8.
5. Make an initial backup of Dynamic Server Version 11.50. For more information, see “Back Up Dynamic Server Version 11.50” on page 5-9.
6. Tune Dynamic Server Version 11.50 for performance. For more information, see “Tune Dynamic Server Version 11.50 for Performance” on page 5-9.

7. If you use Enterprise Replication, perform Enterprise Replication migration tasks. For more information, see “Migrating to Version 11.50 with Enterprise Replication” on page 4-1.
8. If you use High-Availability Data Replication (HDR) clusters, start HDR. For more information, see “Prepare HDR Secondary Servers” on page 5-9
9. Register any DataBlade modules that you installed. For more information, see “Register DataBlade Modules” on page 5-9.

After you migrate and start using Version 11.50, refer to Chapter 6, “Reverting from Dynamic Server Version 11.50,” on page 6-1 in case you need to revert.

Repeat the migration and post-migration procedures for each instance of Dynamic Server Version 11.50 that you plan to run on the computer.

If a serious error occurs during the migration, you might need to return to the previous version of the server, restore from a level-0 backup, and then correct the problem prior to restarting the migration tasks.

Important: Do not connect applications to a database server instance until the migration has successfully completed.

After successful migration to Dynamic Server Version 11.50, you might want to modify configuration files and environment variables to take advantage of Dynamic Server Version 11.50 features. For more information, see the *IBM Informix Dynamic Server Getting Started Guide* and your *IBM Informix Dynamic Server Administrator's Guide*.

For ON-Bar, Rename the `sm_versions.std` File

After migration, rename the `sm_versions.std` file to `sm_versions` for the ON-Bar backup and restore system to run.

Use one of the following methods:

- If you are using the same version of ISM, copy the same `sm_versions` file from your old database server to the new database server installation.
- If you are using other storage managers, copy your previous `sm_versions` file from the old `$INFORMIXDIR/etc` directory to the new `$INFORMIXDIR/etc` directory.
- If you are upgrading from Version 7.31, unload the contents of the `sysutils:bar_version` table.

Update Statistics After Migrating

Run `UPDATE STATISTICS` after migrating to Version 11.50.

In addition, after you convert, revert, or migrate data to a database server on a different operating system, run `UPDATE STATISTICS` to update the information that the target database server uses to plan efficient queries. This recommendation applies both to tables and to UDRs that perform queries. An unqualified `UPDATE STATISTICS` statement that specifies no table and column scope clause and no resolution clause updates all tables and all UDRs that are written in SPL.

If the database has UDRs that are written in the C or Java languages, you must run `UPDATE STATISTICS` statements that explicitly include a `FOR` clause to specify each external language UDR.

Update Statistics on Some System Catalog Tables After Migrating

After migrating successfully to Version 11.50, run UPDATE STATISTICS on some of the system catalog tables in your databases.

If you are migrating from a Version 7.31 or 7.24 database server, be sure to run UPDATE STATISTICS on the following system catalog tables in Dynamic Server Version 11.50:

sysblobs	syscolauth	sysfragments	sysindices	syssynonyms	sysstable
syscolumns	sysconstraints	sysobjstate	sysopclstr	sysstabauth	sysables
sysdefaults	sysdistrib	sysprocauth	sysprocedures	systriggers	sysusers
sysfragauth		sysroleauth			

Review Client Applications and Registry Keys

After you migrate a database server on the same operating system or move the database server to another compatible computer, review the client applications and **sqlhosts** file or registry-key connections.

If necessary, recompile or modify client applications.

Verify that the client-application version you use is compatible with your database server version. If necessary, update the **sqlhosts** file or registry key for the client applications with the new database server information.

For more information about interactions between client applications and different database servers, refer to a client manual.

Verify the Integrity of the Data

Open each database with DB–Access and use **oncheck** to verify that data was not corrupted during the migration process.

You can also verify the integrity of the reserve pages, extents, system catalog tables, data, indexes, and smart large objects, as Table 5-1 shows.

Table 5-1. Commands for Verifying the Data Integrity

Action	oncheck Command
Check reserve pages	oncheck -cr
Check extents	oncheck -ce
Check system catalog tables	oncheck -cc database_name
Check data	oncheck -cD database_name
Check indexes	oncheck -cI database_name
Check smart large objects	oncheck -cs sbspace_name
Check smart large objects plus extents	oncheck -cS sbspace_name

If the **oncheck** utility finds any problems, the utility prompts you to respond to corrective action that it can perform. If you respond Yes to the suggested corrective action, run the **oncheck** command again to make sure the problem has been fixed.

The **oncheck** utility cannot fix data that has become corrupt. If the oncheck utility is unable to fix a corruption problem, you might need to contact Technical Support before you proceed.

Back Up Dynamic Server Version 11.50

Use a backup and restore tool (ON-Bar or **ontape**) to make a level-0 backup of the new database server. Do not overwrite the tapes that contain the final backup of the old database server.

For more information, see the *IBM Informix Backup and Restore Guide*.

Important: Do not restore the backed up logical-log files from your old database server for your new database server.

Tune Dynamic Server Version 11.50 for Performance

After backing up the new server, you can tune the database server to maximize performance.

If you created sample queries for comparison, use them to analyze the performance differences between the old and new database servers and to determine if you need to adjust any configuration parameters or the layout of databases, tables, and chunks. For details on performance topics, see the *IBM Informix Dynamic Server Performance Guide*.

Prepare HDR Secondary Servers

If you use High-Availability Data Replication (HDR), install the new database server and any user-defined objects or DataBlade modules on HDR secondary or remote standalone (RS) secondary servers, make the secondary or RS server the primary server, migrate to that server the same way you migrated to the original primary server, and restart replication.

When you start HDR, follow the procedure for starting HDR for the first time as described in the *IBM Informix Dynamic Server Administrator's Guide*.

For information on configuring a system for HDR, see the *IBM Informix Dynamic Server Administrator's Guide*.

Register DataBlade Modules

Register any DataBlade modules that you installed.

Registration is the process that makes the DataBlade module code available to use in a particular database. For more information on how to use DataBlade modules, see the *IBM Informix DataBlade Module Installation and Registration Guide*.

Chapter 6. Reverting from Dynamic Server Version 11.50

This chapter contains procedures for preparing your databases for reverting from Dynamic Server Version 11.50 to the version of Dynamic Server from which you migrated.

See “Preparing to Revert from Dynamic Server Version 11.50” and “Reverting from Dynamic Server Version 11.50” on page 6-10.

Overview of Reversion

You can revert to the version of the database server from which you migrated. When you run the reversion utility, you specify the target server for reversion and then Dynamic Server checks your database. If necessary, Dynamic Server might tell you to drop new objects, before automatically converting your data into the target server.

If Dynamic Server cannot revert a database, Dynamic Server prevents reversion.

Normally, reversion takes only a few minutes.

If you used the new Version 11.50 features, reversion time is longer, because you must prepare your database and data for reversion, and you must remove the features that are not supported in the earlier version of the server. The more work you complete in the new version, the more time consuming the reversion. See “Preparing to Revert from Dynamic Server Version 11.50” before you revert.

If you did not use any of the new Version 11.50 features and have not completed much work using the new server, you can run the reversion utility and modify the values of the configuration parameters. See “Reverting from Dynamic Server Version 11.50” on page 6-10.

Preparing to Revert from Dynamic Server Version 11.50

Preparing for reversion includes ascertaining that reversion is possible, backing up Version 11.50, and removing new features and objects that are not supported in your original database server.

Prerequisites: Before you revert:

- Read the release notes and the machine notes for new information.

To prepare to revert to your source database server

1. Review the database schema to ascertain that reversion is possible and identify reversion requirements. See “Ascertain that Reversion Is Possible and Identify Reversion Requirements” on page 6-2.
2. Check and configure available space. See “Check and Configure Available Space for Reversion” on page 6-8.
3. Save copies of the current configuration files. See “Save Copies of the Current Configuration Files” on page 6-9.
4. Save system catalog information. See “Save System Catalog Information” on page 6-9.

5. Verify the integrity of the data, if you did not do this after you migrated. See “Verify the Integrity of the Version 11.50 Data” on page 6-9.
6. Back up Dynamic Server Version 11.50. See “Back Up Dynamic Server Version 11.50” on page 6-9.
7. Export or save your data.
8. Remove new features and any new in-place alters that were created using Version 11.50. See “Remove Version 11.50 Features” on page 6-10.
9. Remove any new objects (such as triggers or stored procedures) that you created in the 11.50 database and that are not supported in the version to which you are reverting. Do not remove objects that you did not create, such as the boot scripts (**boot90.sql** and **boot901.sql**) created in the system catalog, because the reversion utility uses them.
10. If you are reverting to the following versions, drop indexes if necessary:
 - Dynamic Server Version 7.3: You must drop any index with a key size that is greater than 254.
 - Dynamic Server Version 9.4: You must drop any index whose key size is greater than 390.

The maximum key size in Versions 11.50, 11.10, and 10.00 is 3200.
11. If you ran BladeManager against a Version 11.50 database, remove any BladeManager extensions. See “Remove New BladeManager Extensions” on page 6-10.
12. If you use High-Availability Data-replication (HDR), disable HDR. See “Disable HDR Before Reversion” on page 6-10.
13. If you use Enterprise Replication, perform additional Enterprise Replication prerequisite reversion tasks. See “Reverting from Dynamic Server Version 11.50 with Enterprise Replication” on page 4-2

After preparing to revert to your source database server, see “Reverting from Dynamic Server Version 11.50” on page 6-10.

Ascertain that Reversion Is Possible and Identify Reversion Requirements

You can revert from Dynamic Server Version 11.50 to the database server from which you migrated, if you have not added any extensions to the Version 11.50 database server and you are not reverting from a newly created instance. You must review your database schema to determine if reversion is possible.

See “Reversion Requirements and Limitations” on page 6-3 for limitations on reversion to previous databases and prerequisite steps you must take before you revert.

To review the database schema to determine if reversion is possible

1. Run the **dbschema** utility command.

For example, run the following command to display information about the database **db1**:

```
dbschema -d db1 -ss
```
2. Determine if the schema file contains SQL statements that the earlier database server does not support.
3. Determine if the database contains features, such as long identifiers, that the earlier database server does not support. See Appendix F, “New and Changed Features,” on page F-1.

4. Determine if any new SPL routines have been created in Dynamic Server Version 11.50 or if any routines were imported using **dbimport**.
5. Determine if tables or indexes using expression fragmentation had expressions changed or new fragments added.
6. Identify any new triggers, procedures, or check constraints.

Reversion Requirements and Limitations

The following table lists reversion requirements and limitations. If you used the new database server, review the list of reversion requirements and limitations, and then complete any prerequisite tasks before you revert. If the reversion restrictions indicate that you must drop objects from the database, you can unload your data and then reload it in the prior database server.

Table 6-1. Reversion Requirements and Limitations

Reversion Requirement or Limitation	If Reverting to this Server or Earlier Servers
UDRs that Use Methods or SQL Statements that Reference Savepoints: You must drop these UDRs, because they include new SQL syntax that earlier Dynamic Server versions do not support. (Before you can compile these UDRs, you must rewrite their error-handling code, so that no savepoint objects are referenced.)	11.50xC2
New indexes in sbspaces: If you built indexes in sbspaces so you could search the sbspaces with the Basic Search Text DataBlade module, you must drop the indexes before reverting.	11.50xC2
Revert Only to the Version From Which You Migrated: If you need to revert, you must revert to the Dynamic Server version that was your source version before you migrated to Version 11.50.	11.10
New Databases Created in Version 11.50: If you created a new database on Dynamic Server Version 11.50, you cannot revert the database back to an earlier version of the server. If the data is required, you can unload the data and reload it in the prior version of the server.	11.10
New Procedures, Fragmented Tables, Check Constraints, and Triggers: These cannot be reverted. You must remove any new procedures, fragmented tables, check constraints, and triggers.	11.10
New Configuration Parameters: These cannot be reverted.	11.10
<p>New or Outstanding In-place Alters: In-place ALTER TABLE statements performed in Dynamic Server Version 11.50 must not be outstanding against any table.</p> <p>If a table has an incomplete new in-place ALTER operation, you must ensure that the in-place ALTER operation is complete by running a dummy UPDATE statement against the table. If the reversion process does not complete successfully because of in-place ALTER operations, the reversion process lists all the tables that need dummy updates. You must perform a dummy update on each of the tables in the list before you can revert to the older database server.</p> <p>If an in-place ALTER operation is incomplete against a system table, run the following script while connected to the database for reversion to a 7.31 or 7.24 database server:</p> <pre>\$INFORMIXDIR/etc/dummyupds7x.sql</pre> <p>Important: Any in-place alter that was completed in a Dynamic Server Version prior to the current version will successfully revert and dummy updates are not necessary for them.</p>	11.10
Version Columns in Tables: If you have version columns in tables, you must remove them.	11.10
BIGINT and BIGSERIAL columns: If you have any BIGINT or BIGSERIAL columns, you must modify or remove them.	11.10
Extended Data Types or Attributes based on BIGINT and BIGSERIAL Data Types: If you have these, you must remove them.	11.10

Table 6-1. Reversion Requirements and Limitations (continued)

Reversion Requirement or Limitation	If Reverting to this Server or Earlier Servers
Casts based on BIGINT and BIGSERIAL Data Types: If you have these, you must remove them.	11.10
Components Installed With the Custom Installation Option: If you installed components with the custom installation option, you can uninstall a component only if you are not breaking any component dependencies.	11.10
JAVA UDRs That Have Been Compiled Using Newer Versions of JDK: These must be recompiled with older JDK versions. For details, see “Recompile Any Java UDRs that Were Compiled Using JDK 5.0” on page 6-12.	11.10
<p>Subqueries in DELETE and UPDATE Statements: If a condition with a subquery in the WHERE clause of DELETE or UPDATE references the same table that the DELETE or UPDATE statement modifies, before you revert, you must rewrite the INSERT or DELETE operation as two separate SQL statements:</p> <ul style="list-style-type: none"> • A SELECT statement that returns qualifying rows of the original table to a temporary table • A DELETE or INSERT statement that modifies the original table by inserting or deleting rows that match rows in the temporary table 	11.10
<p>Returned Data Type from CONCAT and Other SQL String-Manipulation Functions: Because these built-in functions now support promotion of their return value to longer data types, some operations on VARCHAR or NVARCHAR values might fail with overflow error -881 after reversion.</p>	11.10
<p>Automatic Update Statistics Feature: Dynamic Server versions that are earlier than version 11.10.xC1 do not support the Scheduler. Therefore, the functionality of the Automatic Update Statistics feature, which is implemented by the Scheduler, is not available after reversion. To enforce any Automatic Update Statistics policies that you intend to apply to your databases, you must manually issue the corresponding UPDATE STATISTICS statements after reversion to Version 10.00 or to an earlier version.</p>	10.00
<p>ANSI Joins in Distributed Queries: Distributed queries that use ANSI-compliant LEFT OUTER syntax for specifying joined tables and nested loop joins run more efficiently in Dynamic Server Version 10.00.UC4 than in earlier releases. This occurs through sending the query to each participating database server for operations on local tables of those servers. If you revert from Version 10.00.UC4 or later to an earlier release that does not support this implementation of the ANSI-compliant syntax, such queries might show reduced performance, because the Dynamic Server instance from which the query originates will perform the joins locally.</p>	10.00.xC4
<p>The INDEX_SJ and AVOID_INDEX_SJ Optimizer Directives: When queries explicitly use the new INDEX_SJ and AVOID_INDEX_SJ optimizer directives, these directives have no effect when the query runs. You must run UPDATE STATISTICS on stored procedures to force re-compilation of stored procedures.</p> <p>In addition, reversion removes the effect of these directives in SAVE EXTERNAL DIRECTIVES statements and on output from the SET EXPLAIN statement. If you revert to a version of the database server that supports the sysdirectives system catalog table, but does not support the AVOID_INDEX_SJ or INDEX_SJ directives, user informix must delete any active row of sysdirectives that includes AVOID_INDEX_SJ or INDEX_SJ in the directives column.</p>	10.00
<p>sysdbopen() and sysdbclose() procedures: Earlier versions of Dynamic Server do not support these procedures, and any UDRs with these names are not automatically started.</p>	10.00
<p>UDRs include a collection-derived table in the FROM clause of a query: These will not work correctly after reversion to an earlier release.</p>	10.00

Table 6-1. Reversion Requirements and Limitations (continued)

Reversion Requirement or Limitation	If Reverting to this Server or Earlier Servers
<p>Multiple BEFORE, FOR EACH ROW and AFTER triggers for the same INSERT, UPDATE, DELETE, or SELECT event on a table or view, and trigger routine UDRs: Before reverting, you must drop any of the following triggers and UDRs, if they exist in the database:</p> <ul style="list-style-type: none"> • Delete triggers defined on the same table or defined on a view as another Delete trigger • Insert triggers defined on the same table or defined on a view as another Insert trigger • Update triggers defined on the same table or view (or on the same subset of columns) as another Update trigger • Select triggers defined on the same table or same subset of columns as another Select trigger • Trigger routines defined with the FOR TRIGGER keywords • Triggers that use the DELETING, INSERTING, SELECTING, or UPDATING operators in their triggered action 	10.00
<p>Cross-server operations on BOOLEAN, LVARCHAR, or DISTINCT columns: If you revert to a database server version that does not support cross-server operations on BOOLEAN, LVARCHAR, or DISTINCT columns in databases, applications that use this feature will fail.</p>	10.00
<p>sysadmin Database: This database is automatically dropped during reversion.</p>	10.00
<p>Queries That Use SKIP and LIMIT: Dynamic Server Version 10.00xc3 supports queries that use the keywords SKIP and LIMIT. A query that uses either of these keywords will fail with an error after reversion to any earlier version of Dynamic Server.</p>	10.00
<p>FIRST Clause with an ORDER BY Clause: Dynamic Server Version 10.00xc3 supports the ORDER BY clause of the SELECT statement in a subquery whose result set is a collection-derived table (CDT), but only in conjunction with the SKIP keyword or the FIRST keyword (or its keyword synonym LIMIT) in the Projection clause of the same SELECT statement. Queries that use this syntax will fail with an error after reversion to an earlier version of Dynamic Server.</p>	10.00
<p>Label-based access control (LBAC): Before reverting, you must drop any security policy from tables. In addition, because IDSSECURITYLABEL is a new built-in type for Version 11.10, you must remove any columns of that type before you can revert to versions of Dynamic Server that are earlier than Version 11.10.</p>	10.00
<p>Support for Distributed Relational Database Architecture™ (DRDA): Dynamic Server drops stored procedures for metadata that Dynamic Server created automatically. You cannot manually drop these built-in stored procedures.</p>	10.00
<p>Reversion in Remote Standalone (RS) Secondary Environments: If your network contains active RS secondary servers that participate in a high-availability configuration, you must convert the RS secondary servers to standard servers or remove them from the network. Handle reversions the same way you handle reversions in High-Availability Data-replication (HDR) environments.</p> <p>Upon reversion, the syssha database is removed.</p>	10.00
<p>Reversion in Shared Disk (SD) Secondary Server Environments: If your network contains active SD secondary servers that participate in a high-availability configuration, you can revert only the primary server of the cluster. Before you revert the primary server, you must run the onmode -d clear command to disable the shared disk environment.</p> <p>Upon reversion, the syssha database is removed.</p> <p>You cannot revert SD secondary servers.</p>	10.00

Table 6-1. Reversion Requirements and Limitations (continued)

Reversion Requirement or Limitation	If Reverting to this Server or Earlier Servers
UDRs and Applications That Use the TRUNCATE Keyword or the am_truncate() Method: You cannot revert these to a pre-10.00 version of Dynamic Server. During reversion, you must drop or revise any routines that use the TRUNCATE statement, including any newly registered Virtual-Table Interface or Virtual-Index Interface purpose functions.	9.40
Column-Level Encryption: If your tables contain encrypted data, do not revert to a version of the server that does not contain encryption support (any version prior to Version 10.0) because you will not be able to interpret column data without writing a custom DataBlade module that is equivalent to the facilities provided by Dynamic Server. Dynamic Server does not record whether encrypted data is stored in a database.	9.40
XA Data Sources and XA Data Source Types: If you create any XA data sources and XA data source types, you must drop these.	9.40
DataBlade User-Defined Routines (UDRs) that Include the EXTERNAL Clause: The database server administrator (DBSA), user informix by default, must revoke the "extend" role from all users to whom the role has been granted.	9.40
New or Changed Built-in UDRs: Many system catalog tables use built-in UDRs. If you changed the definition of a built-in UDR, you must drop the UDR before reverting.	9.40
Multiple INOUT Parameter Support: You must drop any new UDRs that were created using INOUT parameters.	9.40
Default Roles: During reversion, the defrole column is dropped from the sysusers table. You must revoke default roles from users before reverting.	9.40
Tables and Indexes that Use Fragment Partition Syntax: If you created tables using the new fragment partition syntax, you must drop the tables or you must use the ALTER FRAGMENT INIT statement to change the syntax before reverting to a pre-10.00 version of Dynamic Server. Dynamic Server drops the partition column from the sysfragments table during reversion.	9.40
External Optimizer Directives: You cannot revert if external optimizer directives have been created.	9.40
Non-default Page Size: If you specified the page size for a standard or temporary dbspace, instead of using the default dbspace page size, you must drop all non-default-size dbspaces before you revert to a pre-10.00 version of Dynamic Server.	9.40
IPv6 Addresses: If you used an IPv6 address in the SQLHOSTS file during reversion, you must replace the IPv6 address with either the machine name or IPv4 address assigned to the machine.	9.40
Procedures, Functions, and Triggers Created with New Version: You must drop all triggers, procedures, and functions created with the new version before reverting.	9.40
Maximum chunk size: Dbspaces cannot have chunks larger than 2 GB or in the new chunk format, chunks that extend further than 2 GB into their device or file, or contain more than 2047 chunks.	9.40
Key Length of B-tree Indexes: All indexes must be B-tree indexes with a total key length less than or equal to 390.	9.40 (not applicable for 7.31)
Maximum file size: No files larger than 2 GB can be stored in or used by the database server.	9.30
TAPESIZE and LTAPESIZE Limitation: The TAPESIZE or LTAPESIZE configuration parameters cannot be set to 0.	9.30
ALARMPROGRAM Limitation: The ALARMPROGRAM configuration parameter cannot be set to the alarmprogram.sh file.	9.30
LRU_MAX_DIRTY and LRU_MIN_DIRTY Limitation: If reverting, set these configuration parameters to integers. (These configuration parameters were removed in Version 10.0.)	9.30

Table 6-1. Reversion Requirements and Limitations (continued)

Reversion Requirement or Limitation	If Reverting to this Server or Earlier Servers
UDRs with IN or OUT Parameters: UDRs must not use multiple IN or OUT parameters. Drop all such UDRs before reversion.	9.30
UDRs with Named Return Parameters: UDRs and stored procedures must not use named return parameters.	9.30
Sequence Objects: Sequence objects must not be in use. Drop all sequences before reversion.	9.30
Triggers Created with the INSTEAD OF Clause: Triggers created with the INSTEAD OF clause must not be in use. Drop all such triggers before reversion.	9.30
Multiple Collations: Multiple collations among indexes, stored procedures, triggers, and constraints must not be in use.	9.30
Functional Index Limitations: Functional indexes cannot contain more than 16 parameters.	9.30
High-Data Availability Replication and Enterprise Replication in the Same Server: High-Data Availability Replication and Enterprise Replication cannot co-exist on the same database server.	9.30
LVARCHAR Limitation: The LVARCHAR(<i>n</i>) data types must not be in use if <i>n</i> is not equal to 2042.	9.30
<p>Extensions Added to Sever: You cannot revert to an earlier database server from a database server that has had extensions added unless you remove the extensions.</p> <p>You must remove any new data types or routines that you created either explicitly or by registering a different version of a DataBlade module.</p> <p>To be able to revert, downgrade any DataBlade module to the version that was registered prior to reversion and explicitly drop any data types and routines that were created outside of any DataBlade registration. For information on how to use DataBlade modules, see the DataBlade documentation.</p>	9.30
UDR Limitations: New user-defined or SPL routines must not be created in new databases (either implicitly or explicitly). If you plan to use dbexport to export a database containing existing user-defined or SPL routines, you must drop these routines prior to reversion.	9.30
New Trigger Limitations: New triggers must not be defined in the upgraded databases.	9.30
<p>New Fragment Expressions and Check Constraints: New fragment expressions and check constraints must not exist in the databases. To revert, convert fragmented tables to nonfragmented tables by detaching fragment expressions.</p> <p>You cannot use ALTER TABLE or ALTER INDEX statements to change fragment strategies that existed before the migration to Dynamic Server 9.40.</p>	9.30
item_nvarchar Limitation: Reversion fails if, for an index, the value of item_nvarchar is 255 or higher.	7.31
Newly Added Log Files: If Dynamic Server uses a newly added log file, you cannot reset the status of the file to “newly added” after reversion to the earlier database server.	7.31
PER_STMT_EXEC or PER_STMT_PREP in a DataBlade module: A DataBlade module that uses the PER_STMT_EXEC or PER_STMT_PREP memory duration cannot be used with the earlier database server.	7.31
Select Triggers: Select triggers must not be in use.	7.31
User-Defined Statistics: User-defined statistics must not be in use.	7.31

Table 6-1. Reversion Requirements and Limitations (continued)

Reversion Requirement or Limitation	If Reverting to this Server or Earlier Servers
<p>Long Identifiers: Long identifiers or long user names must not be in use.</p> <p>Before reversion, make sure that the R-tree indexes do not use long identifiers as indexed column names, operator class (opclass) names, or operator class function names.</p> <p>Also, make sure that the following disk structures do not use long identifiers:</p> <ul style="list-style-type: none"> • Dbspace tablespaces (owner and database name length) • Tblspace tablespaces (owner and table space name length) • Dbspaces (owner and dbspace name length) and chunks (path length) 	7.31
<p>Storage Space Names: Each storage space must not have a name that is more than 18 characters long.</p>	7.31
<p>New Routine Languages: New routine languages must not be defined in the upgraded databases. In addition, new language authorizations must not have been completed in the upgraded databases.</p>	7.31
<p>New Operator Classes, Casts, and Extended Types: New operator classes, casts, or extended types must not be defined in the new database server.</p>	7.31
<p>Table Limitations:</p> <ul style="list-style-type: none"> • Databases cannot have tables whose primary access method is a user-defined access method. • Databases cannot have typed tables. • Tables cannot have any user-defined type columns. • Tables cannot have named row types with default values. 	7.31
<p>Index Limitations:</p> <ul style="list-style-type: none"> • You cannot revert detached indexes to Version 7.31. To enable reversion to Version 7.31, retain the Version 7.31 attached index behavior by setting the environment variable DEFAULT_ATTACH in the application environment. • All indexes must be B-tree indexes with a total key length less than or equal to 255. • Tables cannot have any functional or Virtual Index Interface (VII) indexes. • Semi-detached indexes must not be in the databases. • Indexes created with Dynamic Server 10.00 and an opclass that supports nearest-neighbor search cannot be reverted to the earlier database server. 	7.31
<p>Extensibility Features: Databases cannot use any extensibility features, including user-defined access methods, user-defined types, aggregates, routine languages, language authorizations, trace messages, trace message classes, operator classes, errors, type and casts.</p>	7.31

Check and Configure Available Space for Reversion

For Dynamic Server Versions 11.50, 11.10, 10.00, and 9.40, **tblspace** **tblspace** pages can be allocated in non-root chunks. If the Dynamic Server Version 11.50 root chunk is full and **tblspace** **tblspace** pages were allocated in non-root chunks, make sure you have enough space in the root chunk of the target database server.

If the reversion is to a Version 9.30 or earlier database server, all **tblspace** reserved pages are written to the root chunk.

The default number of **tblspace** **tblspace** pages in Version 9.30 is 250 pages (with a size 2k or 4k each, depending on the operating system). To determine how many

pages were allocated and where they were allocated, run **oncheck -pe** and look for the word TBLSpace. This space must be available on the device where the root chunk will be located.

For information on space requirements for Dynamic Server Version 11.50, see “Check and Configure Available Space” on page 3-2.

Save Copies of the Current Configuration Files

Save copies of the ONCONFIG and **concsm.cfg** files for when you migrate to Dynamic Server Version 11.50 again.

Dynamic Server uses the **concsm.cfg** file to configure CSMs.

Save System Catalog Information

If your current database server instance uses secure-auditing masks or external spaces, and you want to preserve the associated catalog information, you must unload these system catalog tables before you revert.

Run the following command to unload the system catalog tables:

```
$INFORMIXDIR/etc/smi_unld
```

When the **smi_unld** utility finishes unloading the information, the utility displays instructions for reloading the information. Save these instructions. After you complete the reversion and bring up your database server, you can reload the data that you preserved. Follow the instructions given with the **smi_unld** utility for reloading the information. Typically, you run the following command:

```
$INFORMIXDIR/etc/smi_load $INFORMIXDIR/etc/
```

Verify the Integrity of the Version 11.50 Data

Verify the integrity of your Version 11.50 data, if you did not do this after you migrated.

To verify the integrity of your Version 11.50 data, run the following commands:

```
oncheck -cI database_name
oncheck -cD database_name
oncheck -cr
oncheck -cc database_name
```

If the **oncheck** utility finds any problems, the utility prompts you to respond to corrective action that it can perform. If you respond Yes to the suggested corrective action, run the **oncheck** command again to make sure the problem has been fixed.

The **oncheck** utility cannot fix data that has become corrupt. If the **oncheck** utility is unable to fix a corruption problem, you might need to contact Technical Support before you proceed.

You will also need to verify the integrity of your data after you revert.

Back Up Dynamic Server Version 11.50

Before you begin the reversion, make a complete level-0 backup of Dynamic Server Version 11.50.

For more information, see the *IBM Informix Backup and Restore Guide*.

Remove Version 11.50 Features

Before you revert, remove all features that your older database server does not support.

For a list of features that you need to remove before reversion, see “Ascertain that Reversion Is Possible and Identify Reversion Requirements” on page 6-2.

Remove New BladeManager Extensions

When you run BladeManager against a database, you automatically create extensions because BladeManager registers its utility DataBlade module, which adds extensions to the database. If you need to revert from Version 11.50, and you ran BladeManager against a database, you must remove the new BladeManager extensions.

To remove the BladeManager extensions, run this command:

```
unprep database name
```

Disable HDR Before Reversion

If you use High-Availability Data Replication (HDR), you must disable it before reversion from Dynamic Server Version 11.50.

To disable HDR, set the primary database server to standard with the following command on the primary database server:

```
onmode -d standard
```

You do not need to perform the whole reversion process on the secondary database server. The secondary server is reverted automatically when you start HDR.

For more information on HDR, see the *IBM Informix Dynamic Server Administrator's Guide*.

Reverting from Dynamic Server Version 11.50

After preparing to revert, run the reversion utility and prepare to use the original database server.

Prerequisites: Before you revert:

- Complete the steps in “Preparing to Revert from Dynamic Server Version 11.50” on page 6-1. This includes determining if reversion is possible and preparing your database for reversion.

To revert from Dynamic Server Version 11.50, complete the following steps. Click the links for more information on each step.

To revert from Dynamic Server Version 11.50

1. Run the reversion utility (**onmode -b**). See “Run the Reversion Utility” on page 6-11.
2. Restore your original configuration parameters. See “Restore Original Configuration Parameters” on page 6-12.
3. Reset your environment variables. See “Restore Original Environment Variables” on page 6-12.

4. If your Dynamic Server Version 11.50 instance used Communications Support Module (CSMs), remove any CSM settings. See “Remove Any Communications Support Module Settings” on page 6-12.
5. If any Java UDRs were compiled using Java Development Kit (JDK) Version 5.0, recompile those UDRs with the earlier JDK version. For details, see “Recompile Any Java UDRs that Were Compiled Using JDK 5.0” on page 6-12.
6. Reinstall and start the target database server. See “Reinstall and Start the Earlier Database Server” on page 6-12.
7. Update statistics. See “Update Statistics After Reverting” on page 6-13.
8. Verify the integrity of the reverted data. See “Verify the Integrity of the Reverted Data” on page 6-13.
9. Back up the target database server. See “Back Up the Database Server” on page 6-13.
10. Return the target database server to online mode. See “Return the Database Server to Online Mode” on page 6-13.
11. If you use high-availability data replication (HDR), prepare to restart HDR. See “Prepare to Restart HDR” on page 6-14.

If you are reverting to a version that is earlier than Version 11.10, the database server automatically drops the **sysadmin** database.

Attention: When you revert to a previous version of the database server, do not reinitialize the database server by using the **-i** command-line parameter. Using the **-i** parameter for reversion would reinitialize the root dbspace, which would destroy your databases.

Run the Reversion Utility

After preparing to revert, run the reversion utility, using an **onmode -b** command.

Important: You must revert to the version of Dynamic Server that was your source database before you migrated. If you revert to a different version of the server, you will corrupt data.

Dynamic Server Version 11.50 must be running when you run the reversion utility. If the reversion utility detects and lists any remaining features that are specific to Dynamic Server Version 11.50, you must remove those features before reversion can complete.

For example, run:

- `onmode -b 11.10` to revert to Version 11.10.
- `onmode -b 10.00` to revert to Version 10.00.
- `onmode -b 9.4` to revert to any Version 9.4 release.
- `onmode -b 9.3` to revert to any Version 9.30.
- `onmode -b 9.2` to revert to Version 9.21.
- `onmode -b 7.3` to revert to Version 7.31.
- `onmode -b 7.2` to revert to Version 7.24

When you revert to the older version, Dynamic Server displays messages that tell you when reversion begins and ends.

When the reversion is complete, Dynamic Server Version 11.50 is offline. The reversion utility drops the Dynamic Server Version 11.50 system catalog tables and

restores compatibility so that you can access the data with the earlier database server. The reversion utility does not revert changes made to the layout of the data that do not affect compatibility.

For more information about the **onmode -b** command, see Chapter 12, “The onmode Utility Reversion Option,” on page 12-1.

Restore Original Configuration Parameters

Replace the Dynamic Server Version 11.50 ONCONFIG configuration file with the ONCONFIG file that you saved before you migrated. Alternatively, you can remove configuration parameters that the earlier database server does not support.

You might also need to adjust the values of existing configuration parameters.

For a list of new configuration parameters by server version, see Appendix B, “New Configuration Parameters,” on page B-1.

Restore Original Environment Variables

Reset the environment variables to values that are appropriate for the earlier database server.

Remove Any Communications Support Module Settings

If your Dynamic Server Version 11.50 instance used CSMs, edit the **sqlhosts** file to remove any **csm** option settings that are not supported in the older database server.

If you do not do this, the older database server will return an invalid **sqlhosts** options error.

You must also delete the **concsm.cfg** file if the older database server does not support CSMs.

Recompile Any Java UDRs that Were Compiled Using JDK 5.0

After you revert and before you start the earlier sever, recompile JAVA UDRs that were compiled using Version 5.0 of the Java Development Kit (JDK) with a JDK version that is earlier than or equal to the version included with the earlier server.

What you do depends on whether your application uses external JAR and class files or JAR files installed on the server:

- If your application uses external JAR and class files (for example, JAR and class files that are listed in JVPCLASSPATH), just recompile the files.
- If your application uses JAR files installed in the server (for example, through the **install_jar()** support function), then you must remove the old JAR file (using **remove_jar()** support function) and re-install the re-compiled JAR file in the database.

Reinstall and Start the Earlier Database Server

Reinstall and configure the earlier database server.

Refer to the instructions in your *IBM Informix Installation Guide* and your *IBM Informix Dynamic Server Administrator's Guide*.

Run the **oninit -s** command to start the earlier database server in quiescent mode.

Do not use the **oninit -i** command.

Update Statistics After Reverting

After you revert to a database server on a different operating system, run UPDATE STATISTICS to update the information that the target database server uses to plan efficient queries. This recommendation applies both to tables and to UDRs that perform queries.

An unqualified UPDATE STATISTICS statement that specifies no table and column scope clause and no resolution clause updates all tables and all UDRs that are written in SPL.

If the database has UDRs that are written in the C or Java languages, you must run UPDATE STATISTICS statements that explicitly include a FOR clause to specify each external language UDR.

Update Statistics on Some System Catalog Tables After Reverting

After a successful reversion, you must run UPDATE STATISTICS on some of the system catalog tables in your databases when the database server starts.

For reversion to a 7.31 database server from Dynamic Server Version 11.50, run UPDATE STATISTICS on the following system catalog tables in the 7.31 database server:

SYSBLOBS	SYSCOLAUTH	SYSFRAGMENTS	SYSSYNONYMS
SYSCOLUMNS		SYSINDEXES	SYSSYNTABLE
SYSCONSTRAINTS		SYSOBJSTATE	SYSTABAUTH SYSTABLES
SYSDEFAULTS		SYSOPCLSTR	SYSTRIGGERS SYSUSERS
SYSDISTRIB		SYSROCAUTH	
SYSFRAGAATH		SYSPROCEDURES	
		SYSROLEAUTH	

Verify the Integrity of the Reverted Data

Before you allow users to access the databases, check the integrity of the reverted data.

Follow the steps in “Verify the Integrity of the Data” on page 3-6.

Back Up the Database Server

After you complete the reversion, use ON-Bar or **ontape** to make a level-0 backup of the database server to which you reverted.

For more information about making backups, see your *IBM Informix Backup and Restore Guide*.

Important: Do not overwrite the tapes that you used to back up your source database server.

Return the Database Server to Online Mode

To bring the old database server online, run the **onmode -m** command.

Then users can access the data.

Prepare to Restart HDR

If you use High-Availability Data Replication (HDR), you need to install the earlier database server and any user-defined objects or DataBlade modules, and then copy the necessary supporting files (such as the ONCONFIG file) from the primary server. You do not need to perform reversion tasks on the secondary server,

The secondary server reverts automatically when you start High-Availability Data-replication (HDR).

For a list of requirements, see the section on configuring a system for HDR in the *IBM Informix Dynamic Server Administrator's Guide*.

You must start HDR from the backup version of it. This can take time, depending how you're using HDR and how much time you have. For information on starting HDR, see the *IBM Informix Dynamic Server Administrator's Guide*.

Part 3. Migration of Data Between Database Servers

Chapter 7. Migrating Database Servers to a New Operating System

This chapter contains procedures for moving data between database servers on different operating systems. You cannot migrate to an operating system on which Dynamic Server is not supported.

This chapter includes the following topics:

- “Choosing a Tool for Moving Data Before Migrating Between Operating Systems”
- “Adjusting Database Tables for File-System Variations” on page 7-2
- “Moving Data to a Database Server on a Different Operating System” on page 7-2
- “Moving Data Between Dynamic Server and Workgroup Edition on Different Operating Systems” on page 7-3

Choosing a Tool for Moving Data Before Migrating Between Operating Systems

If you are migrating between different operating systems, you must choose a method for exporting and importing data.

The tool that you choose for exporting and importing data depends on how much data you plan to move. All these methods deliver similar performance and enable you to modify the schema of the database. The tools that you can use include:

- The **dbexport** and **dbimport** utilities, which you can use to move an entire database
- The UNLOAD and LOAD statements, which move selected columns or tables (The LOAD statement does not change the data format.)
- The **dbload** utility, which you can use to change the data format
- The **onunload** utility, which unloads data in page-sized chunks, and the **onload** utility, which moves data to an identical database server on a computer of the same type
- The High-Performance Loader (HPL), which moves selected columns or tables or an entire database

For an overview of all of these data-migration tools, a comparison of tools, and information on which Dynamic Server versions do not support all of the tools, see “Data-Migration Tools” on page 2-1

For more information on data-migration tools, see:

- Chapter 8, “The dbexport and dbimport Utilities,” on page 8-1
- Chapter 9, “The dbload Utility,” on page 9-1
- Chapter 10, “The dbschema Utility,” on page 10-1
- Chapter 11, “The LOAD and UNLOAD Statements,” on page 11-1
- Chapter 13, “The onunload and onload Utilities,” on page 13-1
- “High-Performance Loader Performance Advantages for Large Databases” on page 2-4 and the *IBM Informix High-Performance Loader User’s Guide*

Adjusting Database Tables for File-System Variations

File-system limitations vary between NFS and non-NFS file systems. You might need to break up large tables when you migrate to a new operating system.

For example, if you have a 3 GB table, but your operating system allows only 2 GB files, break up your table into separate files before you migrate. For more information, see your *IBM Informix Administrator's Guide*.

An Informix storage space can reside on an NFS-mounted file system using regular operating-system files. For information about the NFS products you can use to NFS mount a storage space for an Informix database server, check product compatibility information.

Moving Data to a Database Server on a Different Operating System

You can move data between Informix database servers on UNIX or Linux and Windows.

To move data to a database server on a different operating system

1. Save a copy of the current configuration files.
2. Use ON-Bar, ON-Archive, or **ontape** to make a final level-0 backup. For more information, refer to your *IBM Informix Backup and Restore Guide*.
3. Choose one of the following sets of migration utilities to unload the databases:
 - **dbexport** and **dbimport**
 - UNLOAD, **dbschema**, and LOAD
 - UNLOAD, **dbschema**, and **dbload**
4. Bring the source database server offline.
5. Install and configure the target database server. If you are migrating to Windows, also install the administration tools.
6. Bring the target database server online.
7. Use **dbimport**, LOAD, or **dbload**, or external tables to load the databases into the target database server, depending on which utility you used to export the databases.
8. Make an initial level-0 backup of the target database server.
9. Run UPDATE STATISTICS to update the information that the target database server uses to plan efficient queries.

Adapting Your Programs for a Different Operating System

When you change to a different operating system, you must review your environment-dependent configuration parameters and environment variables.

Certain database server configuration parameters and environment variables are environment dependent:

- Dynamic Server Versions 11.50, 11.10, 10.00, 9.40, 9.30 and 9.21 support Enterprise Replication.
- Dynamic Server 7.31 and 7.30 supports Enterprise Replication and uses Version 3.0 of IBM Informix Enterprise Command Center (IECC).
- Workgroup Edition 7.31 and 7.30 on Windows supports GLS, ON-Bar, Enterprise Replication, and the Gateway products and uses Version 3.0 of IECC.
- Workgroup Edition 7.24 on UNIX supports GLS and uses Version 1.0 of IECC.

For details, see the information on configuration parameters in your *IBM Informix Dynamic Server Administrator's Guide* and the *IBM Informix Dynamic Server Administrator's Reference* and the information on environment variables in your *IBM Informix Dynamic Server Administrator's Guide* and the *IBM Informix Guide to SQL: Reference*.

Ensuring the Successful Creation of System Databases

The first time the database server is brought online, the **sysmaster**, **sysutils**, **sysuser**, and **sysadmin** databases are built. After moving to a database server on a different operating system, check the message log to ensure that the **sysmaster** and **sysutils** databases have been created successfully before you allow users to access the database server.

After you ensure that client users can access data on the database server, the migration process is complete.

Next you might want to seek ways to obtain maximum performance. For details on topics related to performance, see your *IBM Informix Performance Guide*.

Moving Data Between Dynamic Server and Workgroup Edition on Different Operating Systems

The UNLOAD statement lets you retrieve selected rows from a database and write those rows to a text file. If you want to move selected tables or columns instead of an entire database between Dynamic Server and Workgroup Edition, use the UNLOAD and LOAD statements in the DB-Access utility with the **dbschema** utility.

If you need to manipulate the data in the specified UNLOAD file before you load it into a new table, use a combination of the UNLOAD statement and the **dbschema** and **dbload** utilities.

For information on UNLOAD, LOAD, **dbload**, and **dbschema**, see Chapter 9, "The dbload Utility," on page 9-1 and Chapter 10, "The dbschema Utility," on page 10-1. For information on how to use DB-Access, see the *IBM Informix DB-Access User's Guide*.

For more information about using UNLOAD, dbschema, and LOAD to or from Version 7.31 or an earlier version or for information about moving data to or from Version 7.31 on a different operating system, see the Version 7.31 or earlier *Migration Guide*, installation information, and release notes.

Part 4. Data Migration Utilities

Chapter 8. The **dbexport** and **dbimport** Utilities

This chapter contains information on using the **dbexport** utility, which unloads an entire database into text files, and the **dbimport** utility, which creates a database from text files.

Overview of the **dbexport** and **dbimport** Utilities

The **dbexport** utility unloads a database into text files and creates a schema file. You can use the schema file with **dbimport** to re-create the database schema in another IBM Informix environment, and you can edit the schema file to modify the database that **dbimport** creates.

You might want to use the **dbexport** and **dbimport** utilities if you cannot use the **onunload** and **onload** utilities and you want to unload a database with or without its schema file to disk or tape.

The UNLOAD statement lets you manipulate the data as you unload it, but it requires that you unload to files on disk instead of to tape. If you unload to disk files, you might need to use UNIX, Linux, or Windows utilities to load those files onto tape.

The **dbexport** utility supports the following destination options:

- Unload a database and its schema file to disk
- Unload a database and its schema file to tape
- Unload the schema file to disk and unload the data to tape

The **dbimport** utility creates a database and loads it with data from text files on tape or disk. The input files consist of a schema file that is used to re-create the database and data files that contain the database data. Normally, you generate the input files with the **dbexport** utility, but you can use any properly formatted input files.

The **dbexport** and **dbimport** utilities support Dynamic Server 11.50, 11.10, 10.00, 9.40, 9.30, and 9.21 data types.

The **dbimport** utility supports the following options for a new IBM Informix database server:

- Create an ANSI-compliant database (includes unbuffered logging).
- Establish transaction logging for a database (unbuffered or buffered logging).
- Specify the dbspace where the database will reside.

You cannot use the **dbexport** and **dbimport** utilities on High-Availability Data Replication (HDR) secondary servers, remote standalone (RS) secondary servers, or shared disk (SD) secondary servers.

The **dbexport** Utility

The **dbexport** utility unloads an entire database into text files for later import into another database and creates a schema file. The **dbexport** utility supports Dynamic Server 11.50, 11.10, 10.00, 9.40, 9.30, and 9.21 extended data types.

Dates are stored in four-digit years. By default, **dbexport** exports dates in four-digit years unless the environment variable **DBDATE** is set to "mdy2" or to some other value that specifies abbreviated years. Use four-digit years, because data imported back into the database depends on either the **DBCENTURY** environment variable, if set, or the current century if **DBCENTURY** is not set.

Important: You must disable SELECT triggers before exporting a database with **dbexport**. The **dbexport** utility runs SELECT statements during export. The SELECT statement triggers can modify the database content.

If the database uses a non-default locale and the **GL_DATETIME** environment variable has a non-default setting, you must set the **USE_DTENV** environment variable to the value of 1 before you can process localized datetime values correctly with the **dbexport** Utility.

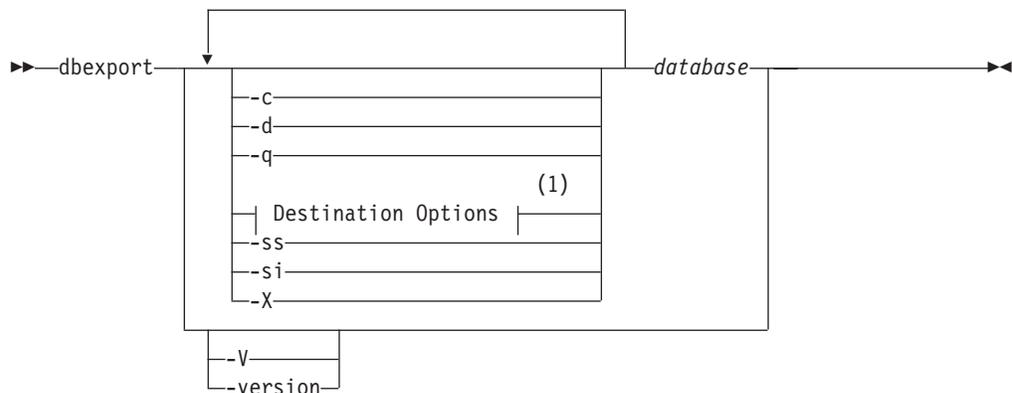
If the database contains label-based access control (LBAC) objects:

- You must have the DBSECADM role.
- You must have the necessary labels or exemptions before the **dbexport** utility can export all rows in protected tables.

For more information on LBAC objects, see the *IBM Informix Security Guide* and the *IBM Informix Guide to SQL: Syntax*.

Syntax of the dbexport Command

The **dbexport** command unloads a database into text files that you can later import into another database. The command also creates a schema file.



Notes:

- 1 See "dbexport Destination Options" on page 8-4

Element	Purpose	Key Considerations
-c	Makes dbexport complete exporting unless a fatal error occurs	References: For details on this option, see "dbexport Errors" on page 8-4.
-d	Makes dbexport export simple-large-object descriptors only, not simple-large-object data	References: For information about simple-large-object descriptors, see the <i>IBM Informix Optical Subsystem Guide</i> . Restrictions: Not supported by SE.

Element	Purpose	Key Considerations
-q	Suppresses the display of error messages, warnings, and generated SQL data-definition statements	None.
-ss	Generates database server-specific information for all tables in the specified database	References: For details on this option, see “dbexport Server-Specific Information” on page 8-4.
-si	Excludes the generation of index storage clauses for non-fragmented tables The -si option is available only with the -ss option.	References: For details on this option, see “dbexport Server-Specific Information” on page 8-4.
-X	Recognizes HEX binary data in character fields	None.
-V	Displays the software version number and the serial number	None.
-version	Extends the -V option to display additional information on the build operating system, build number, and build date	None.
<i>database</i>	Specifies the name of the database that you want to export	Additional Information: If your locale is set to use multibyte characters, you can use multibyte characters for the database name. References: If you want to use more than the simple name of the database, refer to the Database Name section of the <i>IBM Informix Guide to SQL: Syntax</i> .

You must have DBA privileges or log in as user **informix** to export a database.

Global Language Support: When the environment variables are set correctly, as described in the *IBM Informix GLS User's Guide*, **dbexport** can handle foreign characters in data and export the data from GLS databases. For more information, refer to “Database Renaming” on page 8-11.

You can use delimited identifiers with the **dbexport** utility. The utility detects database objects that are keywords, mixed case, or have special characters, and the utility places double quotes around them.

In addition to the data files and the schema file, **dbexport** creates a file of messages named **dbexport.out** in the current directory. This file contains error messages, warnings, and a display of the SQL data definition statements that it generates. The same material is also written to standard output unless you specify the **-q** option.

During export, the database is locked in exclusive mode. If **dbexport** cannot obtain an exclusive lock, it displays a diagnostic message and exits.

Tip: The **dbexport** utility can create files larger than 2 GB. To support such large files, make sure your operating system file-size limits are set sufficiently high. For example, on UNIX, set **ulimit** to unlimited.

Termination of the dbexport Utility

You can press the Interrupt key at any time to cancel **dbexport**.

The **dbexport** utility asks for confirmation before it terminates.

dbexport Errors

The **dbexport -c** option tells **dbexport** to complete exporting unless a fatal error occurs.

Even if you use the **-c** option, **dbexport** interrupts processing if one of the following fatal errors occurs:

- Unable to open the tape device specified
- Bad writes to the tape or disk
- Invalid command parameters
- Cannot open database or no system permission

dbexport Server-Specific Information

The **dbexport -ss** option generates server-specific information. This option specifies initial- and next-extent sizes, fragmentation information if the table is fragmented, the locking mode, the dbspace for a table, the blob space for any simple large objects, and the dbspace for any smart large objects.

The **dbexport -si** option, which is available only with the **-ss** option, does not generate index storage clauses for non-fragmented tables.

dbexport Destination Options

The **dbexport** utility supports disk and tape destination options.

Destination Options:

<code>--o <i>directory</i></code>
<code>--t <i>device</i> --b <i>blocksize</i> --s <i>tapesize</i></code>
<code>--f <i>pathname</i></code>

Element	Purpose	Key Considerations
-b <i>blocksize</i>	Specifies, in kilobytes, the block size of the tape device	None.
-f <i>pathname</i>	Specifies the name of the path where you want the schema file stored, if you are storing the data files on tape	Additional Information: The path name can be a complete path name or a file name. If only a file name is given, the file is stored in the current directory.
-o <i>directory</i>	Specifies the directory on disk in which dbexport creates the <i>database.exp</i> directory. This directory holds the data files and the schema file that dbexport creates for the <i>database</i> .	Restrictions: The directory specified as <i>directory name</i> must already exist.
-s <i>tapesize</i>	Specifies, in kilobytes, the amount of data that you can store on the tape	Additional Information: To write to the end of the tape, specify <i>tapesize</i> as 0. If you do not specify 0, the maximum <i>tapesize</i> is 2 097 151 KB.

Element	Purpose	Key Considerations
-t device	Specifies the path name of the tape device where you want the text files and, possibly, the schema file stored	The -t option does not allow you to specify a remote tape device.

When you write to disk, **dbexport** creates a subdirectory, *database.exp*, in the directory that the **-o** option specifies. The **dbexport** utility creates a file with the *.unl* extension for each table in the database. The schema file is written to the file *database.sql*. The *.unl* and *.sql* files are in the *database.exp* directory.

If you do not specify a destination for the data and schema files, the subdirectory *database.exp* is placed in the current working directory.

When you write the data files to tape, you can use the **-f** option to store the schema file to disk. You are not required to name the schema file *database.sql*. You can give it any name.

UNIX/Linux Only

For non-SE database servers on UNIX or Linux, the command is:

```
dbexport //finland/reports
```

The following command exports the database *stores_demo* to tape with a block size of 16 KB and a tape capacity of 24 000 KB. The command also writes the schema file to */tmp/stores_demo.imp*.

```
dbexport -t /dev/rmt0 -b 16 -s 24000 -f /tmp/stores_demo.imp
stores_demo
```

The following command exports the same *stores_demo* database to the directory named */work/exports/stores_demo.exp*. The resulting schema file is */work/exports/stores_demo.exp/stores_demo.sql*.

```
dbexport -o /work/exports stores_demo
```

Windows Only

For Windows, the following command exports the database *stores_demo* to tape with a block size of 16 KB and a tape capacity of 24 000 KB. The schema file is written to *C:\temp\stores_demo.imp*.

```
dbexport -t \\.\TAPE2 -b 16 -s 24000 -f
C:\temp\stores_demo.imp stores_demo
```

The following command exports the same *stores_demo* database to the directory named *D:\work\exports\stores_demo.exp*. The resulting schema file is *D:\work\exports\stores_demo.exp\stores_demo.sql*.

```
dbexport -o D:\work\exports stores_demo
```

Contents of the Schema File that dbexport Creates

The **dbexport** Utility creates a schema file. This file contains the SQL statements that you need to re-create the exported database.

You can edit the schema file to modify the schema of the database.

The schema file supports all data types for Dynamic Server 11.50, 11.10, 10.00, 9.40, 9.30, and 9.21.

If you use the **-ss** option, the schema file contains server-specific information, such as initial- and next-extent sizes, fragmentation information, lock mode, the **dbspace**

where each table resides, the blob space where each simple-large-object column resides, and the db space for smart large objects. The following information is not retained:

- Logging mode of the database
For information about logging modes, see the *IBM Informix Guide to SQL: Reference*.
- The starting values of SERIAL columns

The statements in the schema file that create tables, views, indexes, partition-fragmented tables and indexes, roles, and grant privileges do so with the name of the user who originally created the database. In this way, the original owner retains DBA privileges for the database and is the owner of all the tables, indexes, and views. In addition, the person who runs the **dbimport** command also has DBA privileges for the database.

The schema file that **dbexport** creates contains comments, enclosed in braces, with information about the number of rows, columns, and indexes in tables, and information about the unload files. The **dbimport** utility uses the information in these comments to load the database.

The number of rows must match in the unload file and the corresponding unload comment in the schema file. If you change the number of rows in the unload file but not the number of rows in the schema file, a mismatch occurs.

Attention: Do not delete any comments in the schema file, and do not change any existing comments or add any new comments. If you change or add comments, the **dbimport** utility might stop or produce unpredictable results.

If you delete rows from an unload file, update the comment in the schema file with the correct number of rows in the unload file. Then **dbimport** will be successful.

The dbimport Utility

The **dbimport** utility creates and populates a database from text files. You can use the schema file with **dbimport** to re-create the database schema in another Informix environment. You can also edit the schema file to modify the database that **dbimport** creates.

The **dbimport** utility supports extended data types in Dynamic Server 11.50, 11.10, 10.00, 9.40, 9.30, and 9.21.

Attention: When you import a database, use the same environment variable settings that were used when the database was created or you might get unexpected results. If any fragmentation expressions, check constraints, triggers, or user-defined routines were created with different settings than you use with **dbimport**, you cannot reproduce the database accurately with a single import.

If the date context during import is not the same as when these objects were created, you might get explicit errors, or you might not be able to find your data, or a check constraint might not work as expected. Many of these problems do not generate errors. The date context for an object includes the date the object was created, the values of the **DBCENTURY** and **DBDATE** environment variables, and some other environment variables. To avoid such problems with the date context, use four-digit dates in all cases.

Global Language Support: Similar problems might occur with environment variables that specify GLS locales, such as **DB_LOCALE**, **SERVER_LOCALE**, and **CLIENT_LOCALE**. For more information, see the *IBM Informix GLS User's Guide*.

If the database uses a non-default locale and the **GL_DATETIME** environment variable has a non-default setting, you must set the **USE_DTENV** environment variable to the value of 1 before you can process localized datetime values correctly with the **dbimport** utility.

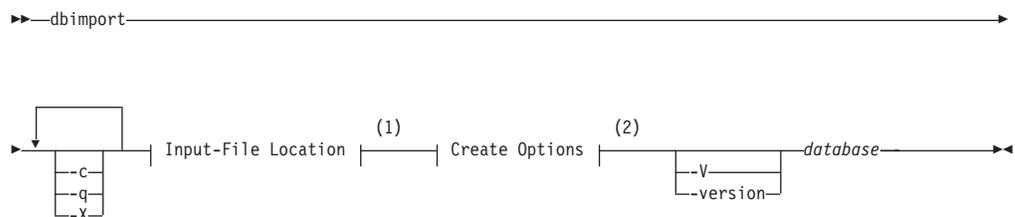
If the database contains label-based access control (LBAC) objects:

- You must have the DBSECADM role.
- You must have the necessary labels or exemptions before the **dbimport** utility can export or import all rows in protected tables.

For more information on LBAC objects, see the *IBM Informix Security Guide* and the *IBM Informix Guide to SQL: Syntax*.

Syntax of the dbimport Command

The **dbimport** command imports previously exported data into another database.



Notes:

- 1 See “dbimport Input-File Location Options” on page 8-8
- 2 See “dbimport Create Options” on page 8-10

Element	Purpose	Key Considerations
-c	Completes importing data even when certain nonfatal errors occur	References: For more information, see “dbimport Errors and Warnings” on page 8-8.
-q	Suppresses the display of error messages, warnings, and generated SQL data-definition statements	None.
-V	Displays the software version number and the serial number	None.
-version	Extends the -V option to display additional information on the build operating system, build number, and build date	None.
-X	Recognizes HEX binary data in character fields	None.
<i>database</i>	Specifies the name of the database to create	Additional Information: To use more than the simple name of the database, see the Database Names segment in the <i>IBM Informix Guide to SQL: Syntax</i> .

The **dbimport** utility can use files from the following location options:

- All input files are located on disk.

- All input files are located on tape.
- The schema file is located on disk, and the data files are on tape.

Important: Do not put comments into your input file. Comments might cause unpredictable results when the **dbimport** utility reads them.

The **dbimport** utility supports the following tasks for a new Informix database server (excluding SE):

- Create an ANSI-compliant database (includes unbuffered logging)
- Establish transaction logging for a database (unbuffered or buffered logging)
- Specify the dbspace where the database will reside

The user who runs **dbimport** is granted the DBA privilege on the newly created database. The **dbimport** process locks each table as it is being loaded and unlocks the table when the loading is complete.

Global Language Support: When the GLS environment variables are set correctly, as the *IBM Informix GLS User's Guide* describes, **dbimport** can import data into database versions that support GLS.

Termination of the dbimport Utility

You can press the Interrupt key at any time to cancel **dbimport**.

The **dbimport** utility asks for confirmation before it terminates.

dbimport Errors and Warnings

The **dbimport -c** option tells **dbexport** to complete exporting unless a fatal error occurs.

If you include the **-c** option in a **dbimport** command, **dbimport** ignores the following errors:

- A data row that contains too many columns
- Inability to put a lock on a table
- Inability to release a lock

Even if you use the **-c** option, **dbimport** interrupts processing if one of the following fatal errors occurs:

- Unable to open the tape device specified
- Bad writes to the tape or disk
- Invalid command parameters
- Cannot open database or no system permission
- Cannot convert the data

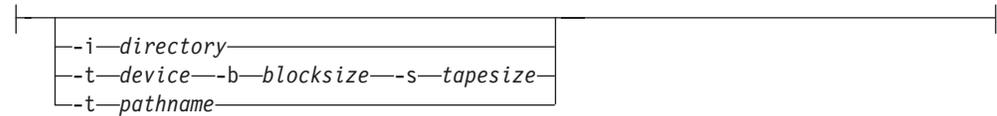
The **dbimport** utility creates a file of messages called **dbimport.out** in the current directory. This file contains any error messages and warnings that are related to **dbimport** processing. The same information is also written to the standard output unless you specify the **-q** option.

dbimport Input-File Location Options

The input-file location specifies the location of the *database.exp* directory, which contains the files that **dbimport** will import.

If you do not specify an input-file location, **dbimport** searches for data files in the directory *database.exp* under the current directory and for the schema file in *database.exp/database.sql*.

dbimport Input-File Location:



Element	Purpose	Key Considerations
-b <i>blocksize</i>	Specifies, in kilobytes, the block size of the tape device	If you are importing from tape, you must use the same block size that you used to export the database.
-f <i>pathname</i>	Specifies where dbimport can find the schema file to use as input to create the database when the data files are read from tape	Additional Information: If you use the -f option to export a database, you typically use the same path name that you specified in the dbexport command. If you specify only a file name, dbimport looks for the file in the .exp subdirectory of your current directory.
-i <i>directory</i>	Specifies the complete path name on disk of the <i>database.exp</i> directory, which holds the input data files and schema file that dbimport uses to create and load the new database. The directory name must be the same as the database name.	Additional Information: This directory must be the same directory that you specified with the dbexport -o option. If you change the directory name, you also rename your database.
-s <i>tapesize</i>	Specifies, in kilobytes, the amount of data that you can store on the tape	Additional Information: To read to the end of the tape, specify a tape size of 0. If you are importing from tape, you must use the same tape size that you used to export the database. If you do not specify 0 as the <i>tapesize</i> , then the maximum <i>tapesize</i> is 2 097 151 KB.
-t <i>device</i>	Specifies the path name of the tape device that holds the input files	The -t option does <i>not</i> allow you to specify a remote tape device.

Examples Showing Input File Location on UNIX or Linux

To import the **stores_demo** database from a tape with a block size of 16 KB and a capacity of 24 000 KB, issue this command:

```
dbimport -c -t /dev/rmt0 -b 16 -s 24000 -f
/tmp/stores_demo.imp stores_demo
```

The schema file is read from **/tmp/stores_demo.imp**.

To import the **stores_demo** database from the **stores_demo.exp** directory under the **/work/exports** directory, issue this command:

```
dbimport -c -i /work/exports stores_demo
```

The schema file is assumed to be **/work/exports/stores_demo.exp/stores_demo.sql**.

Examples Showing Input File Location on Windows

To import the **stores_demo** database from a tape with a block size of 16 KB and a capacity of 24 000 KB, issue this command:

```
dbimport -c -t \\.\TAPEDRIVE -b 16 -s 24000 -f
C:\temp\stores_demo.imp stores_demo
```

The schema file is read from **C:\temp\stores_demo.imp**.

To import the **stores_demo** database from the **stores_demo.exp** directory under the **D:\work\exports** directory, issue this command:

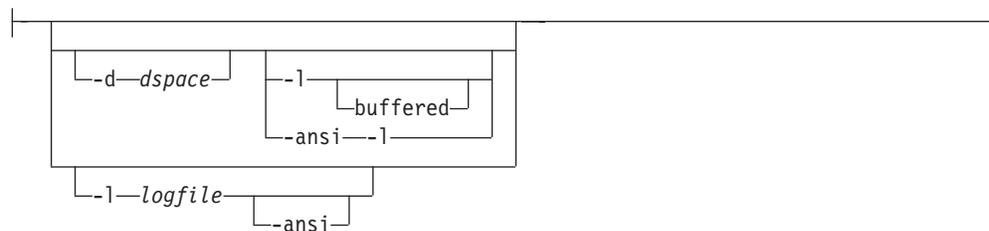
```
dbimport -c -i D:\work\exports stores_demo
```

The schema file is assumed to be **D:\work\exports\stores_demo.exp\stores_demo.sql**.

dbimport Create Options

The dbimport utility contains options for creating an ANSI-compliant database, specifying a dbspace for the database, and defining logging options.

Create Options:



Element	Purpose	Key Considerations
-ansi	Creates an ANSI-compliant database in which the ANSI rules for transaction logging are enabled	Additional Information: If you specify the -ansi option, you must also specify the -l logfile option. For more information about ANSI-compliant databases, see the <i>IBM Informix Guide to SQL: Reference</i> .
-d dspace	Specifies the dbspace where the database is created. The default dbspace location is the rootdbs .	Additional Information: For SE, the database is always in the current directory.
-l	Establishes unbuffered transaction logging for the imported database	References: For more information, see "Database-Logging Mode" on page 8-11.
-l buffered	Establishes buffered transaction logging for the imported database	References: For more information, see "Database-Logging Mode" on page 8-11.
-l logfile	Establishes transaction logging for the imported database and specifies the name of the transaction-log file	For SE, the <i>logfile</i> file name must be an absolute path name, or it must be in the current directory. References: For more information, see "Database-Logging Mode" on page 8-11.

If you created a table or index fragment containing partitions in Dynamic Server 11.50, 11.10, 10.00 or a later version, you must use syntax containing the partition name when importing a database that contains multiple partitions within a single dbspace. See the *IBM Informix Guide to SQL: Syntax* for syntax details.

Example Showing dbimport Create Options (UNIX or Linux)

To import the **stores_demo** database from the **/usr/informix/port/stores_demo.exp** directory, issue this command:

```
dbimport -c stores_demo -i /usr/informix/port -l
/usr/work/stores_demo.log -ansi
```

The new database is ANSI compliant, and the transaction-log file is specified as **stores_demo.log** in **/usr/work**.

Example Showing dbimport Create Options (Windows)

To import the **stores_demo** database from the **C:\USER\informix\port\stores_demo.exp** directory, issue this command:

```
dbimport -c stores_demo -i C:\USER\informix\port -l
C:\USER\work\stores_demo.log -ansi
```

The new database is ANSI compliant, and the transaction-log file is specified as **stores_demo.log** in **C:\USER\work**.

Database-Logging Mode

Because the logging mode is not retained in the schema file, you can specify logging information when you use **dbimport** to import a database.

You can specify any of the following options when you use **dbimport**:

- ANSI-compliant database with unbuffered logging
- Unbuffered logging
- Buffered logging

For more information, see “dbimport Create Options” on page 8-10.

The **-l** options are equivalent to the logging clauses of the CREATE DATABASE statement, as follows:

- The **-l** option is equivalent to the WITH LOG clause.
- The **-l buffered** option is equivalent to the WITH BUFFERED LOG.

For more information about the CREATE DATABASE statement, see the *IBM Informix Guide to SQL: Syntax*.

Database Renaming

The **dbimport** utility gives the new database the same name as the database that you exported. If you export a database to tape, you cannot change its name when you import it with **dbimport**. If you export a database to disk, you can change the database name.

The following examples show the steps to take to change the database name. In this example, assume that **dbexport** unloaded the database **stores_demo** into the directory **/work/exports/stores_demo.exp**. Thus, the data files (the **.unl** files) are stored in **/work/exports/stores_demo.exp**, and the schema file is **/work/exports/stores_demo.exp/stores_demo.sql**.

To change the database name to a new name on UNIX or Linux

1. Change the name of the **.exp** directory. That is, change **/work/exports/stores_demo.exp** to **/work/exports/newname.exp**.
2. Change the name of the schema file. That is, change **/work/exports/stores_demo.exp/stores_demo.sql** to **/work/exports/stores_demo.exp/newname.sql**. Do not change the names of the **.unl** files.
3. Import the database with the following command:

```
dbimport -i /work/exports newname
```

To change the database name to a new name on Windows

In the following example, assume that **dbexport** unloaded the database **stores_demo** into the directory **D:\work\exports\stores_demo.exp**. Thus, the data files (the **.unl** files) are stored in **D:\work\exports\stores_demo.exp**, and the schema file is **D:\work\exports\stores_demo.exp\stores_demo.sql**.

1. Change the name of the **.exp** directory. That is, change **D:\work\exports\stores_demo.exp** to **D:\work\exports\newname.exp**.
2. Change the name of the schema file. That is, change **D:\work\exports\stores_demo.exp\stores_demo.sql** to **D:\work\exports\stores_demo.exp\newname.sql**. Do not change the names of the **.unl** files.
3. Import the database with the following command:

```
dbimport -i D:\work\exports
```

Changing the Database Locale with dbimport

You can use **dbimport** to change the locale of a database.

To change the locale of a database

1. Set the **DB_LOCALE** environment variable to the name of the current database locale.
2. Run **dbexport** on the database.
3. Use the **DROP DATABASE** statement to drop the database that has the current locale name.
4. Set the **DB_LOCALE** environment variable to the desired database locale for the database.
5. Run **dbimport** to create a new database with the desired locale and import the data into this database.

Simple Large Objects (Version 9.21 or Later Versions)

When the **dbimport**, **dbexport**, and DB–Access utilities process simple-large-object data, they create temporary files for that data in a temporary directory.

Before you export or import data from tables that contain simple large objects, you must have one of the following items:

- A **\tmp** directory on your currently active drive
- The **DBTEMP** environment variable set to point to a directory that is available for temporary storage of the simple large objects

Windows Only

Windows sets the **TMP** and **TEMP** environment variables in the command prompt sessions, by default. However, if the **TMP**, **TEMP**, and **DBTEMP** environment variables are not set, **dbimport** places the temporary files for the simple large objects in the **\tmp** directory.

Attention: If a table has a CLOB or BLOB in a column, you cannot use **dbexport** to export the table to a tape. If a table has a user-defined type in a column, using **dbexport** to export the table to a tape might yield unpredictable results, depending on the export function of the user-defined type. Exported CLOB sizes are stored in hex format in the unload file.

Chapter 9. The dbload Utility

This chapter contains information on using the **dbload** utility, which loads data into databases or tables that IBM Informix products created.

Overview of the dbload Utility

The **dbload** utility loads data into databases or tables that IBM Informix products created. It transfers data from one or more text files into one or more existing tables.

This utility supports new data types in all versions of Dynamic Server since Version 9.21.

Prerequisites: If the database contains label-based access control (LBAC) objects, the **dbload** utility can load only those rows in which your security label dominates the column-security label or the row-security label. If the entire table is to be loaded, you must have the necessary LBAC credentials for writing all of the labeled rows and columns. For more information on LBAC objects, see the *IBM Informix Security Guide* and the *IBM Informix Guide to SQL: Syntax*.

You cannot use the **dbload** utility on High-Availability Data Replication (HDR) secondary servers, remote standalone (RS) secondary servers, or shared disk (SD) secondary servers.

When you use the **dbload** utility, you can manipulate a data file that you are loading or access a database while it is loading. When possible, use the LOAD statement, which is faster than **dbload**.

The **dbload** utility gives you a great deal of flexibility, but it is not as fast as the other methods, and you must prepare a command file to control the input. You can use **dbload** with data in a variety of formats.

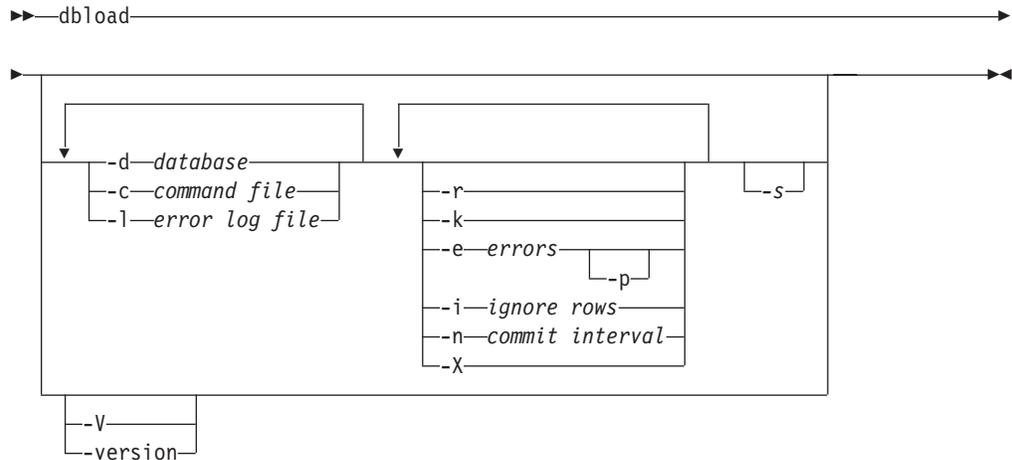
The **dbload** utility offers the following advantages over the LOAD statement:

- You can use **dbload** to load data from input files that were created with a variety of format arrangements. The **dbload** command file can accommodate data from entirely different database management systems.
- You can specify a starting point in the load by directing **dbload** to read but ignore *x* number of rows.
- You can specify a batch size so that after every *x* number of rows are inserted, the insert is committed.
- You can limit the number of bad rows read, beyond which **dbload** ends.

The cost of **dbload** flexibility is the time and effort spent creating the **dbload** command file, which is required for **dbload** operation. The input files are not specified as part of the **dbload** command line, and neither are the tables into which the data is inserted. This information is contained in the command file.

Syntax of the dbload Command

The **dbload** command loads data into databases or tables.



Element	Purpose	Key Considerations
-c <i>command file</i>	Specifies the file name or path name of a dbload command file	References: For information about building the command file, see “Command File for the dbload Utility” on page 9-5.
-d <i>database</i>	Specifies the name of the database to receive the data	Additional Information: If you want to use more than the simple name of the database, see the Database Name section of the <i>IBM Informix Guide to SQL: Syntax</i> .
-e <i>errors</i>	Specifies the number of bad rows that dbload reads before terminating. The default value for <i>errors</i> is 10.	References: For more information, see “Bad-Row Limit During a Load Operation” on page 9-4.
-i <i>ignore rows</i>	Specifies the number of rows to ignore in the input file	References: For more information, see “Rows to Ignore During a Load Operation” on page 9-4.
-k	Instructs dbload to lock the tables listed in the command file in exclusive mode during the load operation	References: For more information, see “Table Locking During a Load Operation” on page 9-3. You cannot use the -k option with the -r option because the -r option specifies that no tables are locked during the load operation.
-l <i>error log file</i>	Specifies the file name or path name of an error log file	If you specify an existing file, its contents are overwritten. If you specify a file that does not exist, dbload creates the file. Additional Information: The error log file stores diagnostic information and any input file rows that dbload cannot insert into the database.
-n <i>commit interval</i>	Specifies the commit interval in number of rows The default interval is 100 rows.	Additional Information: If your database supports transactions, dbload commits a transaction after the specified number of new rows is read and inserted. A message appears after each commit. References: For information about transactions, see the <i>IBM Informix Guide to SQL: Tutorial</i> .

Element	Purpose	Key Considerations
-p	Prompts for instructions if the number of bad rows exceeds the limit	References: For more information, see “Bad-Row Limit During a Load Operation” on page 9-4.
-r	Prevents dbload from locking the tables during a load, thus enabling other users to update data in the table during the load	Additional Information: For more information, see “Table Locking During a Load Operation.” You cannot use the -r option with the -k option because the -r option specifies that the tables are not locked during the load operation while the -k option specifies that the tables are locked in exclusive mode.
-s	Checks the syntax of the statements in the command file without inserting data	Additional Information: The standard output displays the command file with any errors marked where they are found.
-V	Displays the software version number and the serial number	None.
-version	Extends the -V option to display additional information on the build operating system, build number, and build date	None.
-X	Recognizes HEX binary data in character fields	None.

Tip: If you specify part (but not all) of the required information, **dbload** prompts you for additional specifications. The database name, command file, and error log file are all required. If you are missing all three options, you receive an error message.

dbload Command Example

The following command loads data into the **stores_demo** database in the **turku** directory on a database server called **finland**:

```
dbload -d //finland/turku/stores_demo -c commands -l errlog
```

Table Locking During a Load Operation

The **dbload -k** option overrides the default table lock mode during the load operation. The **-k** option instructs **dbload** to lock the tables in exclusive mode rather than shared mode.

If you do not specify the **-k** option, the tables specified in the command file are locked in shared mode. When tables are locked in shared mode, the database server still has to acquire exclusive row or page locks when it inserts rows into the table.

When you specify the **-k** option, the database server places an exclusive lock on the entire table. The **-k** option increases performance for large loads because the database server does not have to acquire exclusive locks on rows or pages as it inserts rows during the load operation.

If you do not specify the **-r** option, the tables specified in the command file are locked during loading so that other users cannot update data in the table. Table locking reduces the number of locks needed during the load but reduces concurrency. If you are planning to load a large number of rows, use table locking and load during nonpeak hours.

Rows to Ignore During a Load Operation

The **dbload -i** option specifies the number of new-line characters in the input file that **dbload** ignores before **dbload** begins to process data.

This option is useful if your most recent **dbload** session ended prematurely.

For example, if **dbload** ends after it inserts 240 lines of input, you can begin to load again at line 241 if you set *number rows ignore* to 240.

The **-i** option is also useful if header information in the input file precedes the data records.

Bad-Row Limit During a Load Operation

The **dbload -e** option lets you specify how many bad rows to allow before **dbload** terminates.

If you set *errors* to a positive integer, **dbload** terminates when it reads (*errors* + 1) bad rows. If you set *errors* to zero, **dbload** terminates when it reads the first bad row.

If **dbload** exceeds the bad-row limit and the **-p** option is specified, **dbload** prompts you for instructions before it terminates. The prompt asks whether you want to roll back or to commit all rows that were inserted since the last transaction.

If **dbload** exceeds the bad-row limit and the **-p** option is not specified, **dbload** commits all rows that were inserted since the last transaction.

Termination of the dbload Utility

If you press the Interrupt key, **dbload** terminates and discards any new rows that were inserted but not yet committed to the database (if the database has transactions).

Name and Object Guidelines for the dbload Utility

You must follow guidelines for specifying network names and handling simple large objects, indexes, and delimited identifiers when you use the **dbload** utility.

Table 9-1. Name and Object Guidelines for the dbload Utility

Objects	Guideline
Network names	If you are on a network, include the database server name and directory path with the database name to specify a database on another database server.
Simple large objects	You can load simple large objects with the dbload utility as long as the simple large objects are in text files.
Indexes	The presence of indexes greatly affects the speed with which the dbload utility loads data. For best performance, drop any indexes on the tables that receive the data before you run dbload . You can create new indexes after dbload has finished.

Table 9-1. Name and Object Guidelines for the dbload Utility (continued)

Objects	Guideline
Delimited identifiers	<p>You can use delimited identifiers with the dbload utility. The utility detects database objects that are keywords, mixed case, or have special characters, and places double quotes around them.</p> <p>If your most recent dbload session ended prematurely, specify the starting line number in the command-line syntax to resume loading with the next record in the file.</p>

Command File for the dbload Utility

Before you use **dbload**, you must create a command file that names the input data files and the tables that receive the data. The command file maps fields from one or more input files into columns of one or more tables within your database.

The command file contains only FILE and INSERT statements. Each FILE statement names an input data file. The FILE statement also defines the data fields from the input file that are inserted into the table. Each INSERT statement names a table to receive the data. The INSERT statement also defines how **dbload** places the data that is described in the FILE statement into the table columns.

Within the command file, the FILE statement can appear in these forms:

- Delimiter form
- Character-position form

The FILE statement has a size limit of 4,096 bytes.

Use the delimiter form of the FILE statement when every field in the input data row uses the same delimiter and every row ends with a new-line character. This format is typical of data rows with variable-length fields. You can also use the delimiter form of the FILE statement with fixed-length fields as long as the data rows meet the delimiter and new-line requirements. The delimiter form of the FILE and INSERT statements is easier to use than the character-position form.

Use the character-position form of the FILE statement when you cannot rely on delimiters and you must identify the input data fields by character position within the input row. For example, use this form to indicate that the first input data field begins at character position 1 and continues until character position 20. You can also use this form if you must translate a character string into a null value. For example, if your input data file uses a sequence of blanks to indicate a null value, you must use this form if you want to instruct **dbload** to substitute null at every occurrence of the blank-character string.

You can use both forms of the FILE statement in a single command file. For clarity, however, the two forms are described separately in sections that follow.

Delimiter Form of the FILE and INSERT Statements

The FILE and INSERT statements that define information for the **dbload** utility can appear in a delimiter form.

The following example of a **dbload** command file illustrates a simple delimiter form of the FILE and INSERT statements. The example is based on the **stores_demo** database. An UNLOAD statement created the three input data files, **stock.unl**, **customer.unl**, and **manufact.unl**.

```
FILE stock.unl DELIMITER '|' 6;
INSERT INTO stock;
FILE customer.unl DELIMITER '|' 10;
INSERT INTO customer;
FILE manufact.unl DELIMITER '|' 3;
INSERT INTO manufact;
```

To see the **.unl** input data files, refer to the directory **\$INFORMIXDIR/demo/prod_name** (UNIX or Linux) or **%INFORMIXDIR%\demo\prod_name** (Windows).

Syntax for the Delimiter Form

The syntax for the delimiter form specifies the field delimiter, the input file, and the number of fields in each row of data.

The following diagram shows the syntax of the delimiter FILE statement.

►►—FILE—*filename*—DELIMITER—'*c*'—*nfields*—►►

Element	Purpose	Key Considerations
<i>c</i>	Specifies the character as the field delimiter for the specific input file	If the delimiter specified by <i>c</i> appears as a literal character anywhere in the input file, the character must be preceded with a backslash (\) in the input file. For example, if the value of <i>c</i> is specified as a square bracket (I) , you must place a backslash before any literal square bracket that appears in the input file. Similarly, you must precede any backslash that appears in the input file with an additional backslash.
<i>filename</i>	Specifies the input file	None.
<i>nfields</i>	Indicates the number of fields in each data row	None.

The **dbload** utility assigns the sequential names **f01**, **f02**, **f03**, and so on to fields in the input file. You cannot see these names, but if you refer to these fields to specify a value list in an associated INSERT statement, you must use the **f01**, **f02**, **f03** format. For details, refer to “How to Write a dbload Command File in Delimiter Form” on page 9-8.

Two consecutive delimiters define a null field. As a precaution, you can place a delimiter immediately before the new-line character that marks the end of each data row. If the last field of a data row has data, you must use a delimiter. If you omit this delimiter, an error results whenever the last field of a data row is not empty.

Inserted data types correspond to the explicit or default column list. If the data field width is different from its corresponding character column width, the data is made to fit. That is, inserted values are padded with blanks if the data is not wide enough for the column or truncated if the data is too wide for the column.

If the number of columns named is fewer than the number of columns in the table, **dbload** inserts the default value that was specified when the table was created for the unnamed columns. If no default value is specified, **dbload** attempts to insert a

null value. If the attempt violates a not null restriction or a unique constraint, the insert fails, and an error message is returned.

If the INSERT statement omits the column names, the default INSERT specification is every column in the named table. If the INSERT statement omits the VALUES clause, the default INSERT specification is every field of the previous FILE statement.

An error results if the number of column names listed (or implied by default) does not match the number of values listed (or implied by default).

The syntax of **dbload** INSERT statements resembles INSERT statements in SQL, except that in **dbload**, INSERT statements cannot incorporate SELECT statements.

Do not use the CURRENT, TODAY, and USER keywords of the INSERT INTO statement in a **dbload** command file; they are not supported in the **dbload** command file. These keywords are supported in SQL only.

For example, the following **dbload** command is not supported:

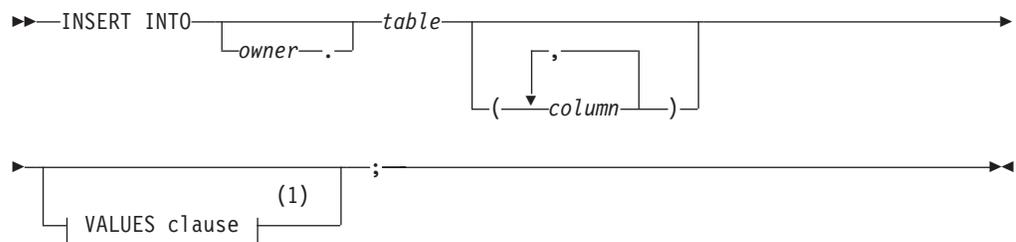
```
FILE "testtb12.un1" DELIMITER '|' 1;
    INSERT INTO testtb1
        (testuser, testtime, testfield)
    VALUES
        ('kae', CURRENT, f01);
```

Load the existing data first and then write an SQL query to insert or update the data with the current time, date, or user login. You could write the following SQL statement:

```
INSERT INTO testtb1
    (testuser, testtime, testfield)
VALUES
    ('kae', CURRENT, f01);
```

The CURRENT keyword returns the system date and time. The TODAY keyword returns the system date. The USER keyword returns the user login name.

The following diagram shows the syntax of the **dbload** INSERT statement for delimiter form.



Notes:

- 1 See *IBM Informix Guide to SQL: Syntax*

Element	Purpose	Key Considerations
<i>column</i>	Specifies the column that receives the new data	None.
<i>owner.</i>	Specifies the user name of the table owner	None.
<i>table</i>	Specifies the table that receives the new data	None.

Users who run **dbload** with this command file must have the Insert privilege on the named table.

How to Write a dbload Command File in Delimiter Form

Command files must contain required elements, including delimiters.

The FILE statement in the following example describes the **stock.unl** data rows as composed of six fields, each separated by a vertical bar (|) as the delimiter.

```
FILE stock.unl DELIMITER '|' 6;  
INSERT INTO stock;
```

Two consecutive delimiters define a null field. As a precaution, you can place a delimiter immediately before the new-line character that marks the end of each data row. If the last field of a data row has data, you must use a delimiter. If you omit this delimiter, an error results.

Compare the FILE statement with the data rows in the following example, which appear in the input file **stock.unl**. (Because the last field is not followed by a delimiter, an error results if any data row ends with an empty field.)

```
1|SMT|baseball gloves|450.00|case|10 gloves/case  
2|HRO|baseball|126.00|case|24/case  
3|SHK|baseball bat|240.00|case|12/case
```

The example INSERT statement contains only the required elements. Because the column list is omitted, the INSERT statement implies that values are to be inserted into every field in the **stock** table. Because the VALUES clause is omitted, the INSERT statement implies that the input values for every field are defined in the most recent FILE statement. This INSERT statement is valid because the **stock** table contains six fields, which is the same number of values that the FILE statement defines.

The following example shows the first data row that is inserted into **stock** from this INSERT statement.

Field	Column	Value
f01	stock_num	1
f02	manu_code	SMT
f03	description	baseball gloves
f04	unit_price	450.00
f05	unit	case
f06	unit_descr	10 gloves/case

The FILE and INSERT statement in the following example illustrates a more complex INSERT statement syntax:

```
FILE stock.unl DELIMITER '|' 6;  
INSERT INTO new_stock (col1, col2, col3, col5, col6)  
VALUES (f01, f03, f02, f05, 'autographed');
```

In this example, the VALUES clause uses the field names that **dbload** assigns automatically. You must reference the automatically assigned field names with the letter **f** followed by a number: **f01**, **f02**, **f10**, **f100**, **f999**, **f1000**, and so on. All other formats are incorrect.

Tip: The first nine fields must include a zero: f01, f02, ..., f09.

The user changed the column names, the order of the data, and the meaning of **col6** in the new **stock** table. Because the fourth column in **new_stock** (**col4**) is not named in the column list, the new data row contains a null value in the **col4** position (assuming that the column permits null values). If no default is specified for **col4**, the inserted value is null.

The following table shows the first data row that is inserted into **new_stock** from this INSERT statement.

Column	Value
col1	1
col2	baseball gloves
col3	SMT
col4	null
col5	case
col6	autographed

Character-Position Form of the FILE and INSERT Statements

The FILE and INSERT statements that define information for the **dbload** utility can appear in a character-position form.

The examples in this topic are based on an input data file, **cust_loc_data**, which contains the last four columns (**city**, **state**, **zipcode**, and **phone**) of the **customer** table. Fields in the input file are padded with blanks to create data rows in which the location of data fields and the number of characters are the same across all rows. The definitions for these fields are CHAR(15), CHAR(2), CHAR(5), and CHAR(12), respectively. Figure 9-1 displays the character positions and five example data rows from the **cust_loc_data** file.

	12		3
	1234567890123456789012345678901234		
Sunnyvale	CA94086408-789-8075		
Denver	C080219303-936-7731		
Blue Island	NY60406312-944-5691		
Brighton	MA02135617-232-4159		
Tempe	AZ85253xxx-xxx-xxxx		

Figure 9-1. A Sample Data File

The following example of a **dbload** command file illustrates the character-position form of the FILE and INSERT statements. The example includes two new tables, **cust_address** and **cust_sort**, to receive the data. For the purpose of this example, **cust_address** contains four columns, the second of which is omitted from the column list. The **cust_sort** table contains two columns.

```
FILE cust_loc_data
(city 1-15,
state 16-17,
area_cd 23-25 NULL = 'xxx',
phone 23-34 NULL = 'xxx-xxx-xxxx',
zip 18-22,
state_area 16-17 : 23-25);
```

```

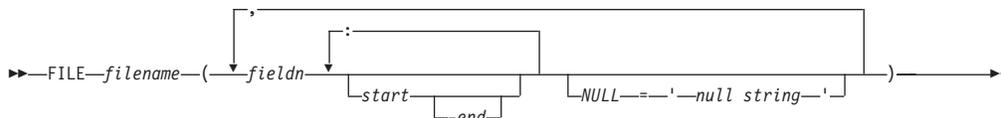
INSERT INTO cust_address (col1, col3, col4)
VALUES (city, state, zip);
INSERT INTO cust_sort
VALUES (area_cd, zip);

```

Syntax for the Character-Position Form

The syntax for the character-position form specifies information that includes the character position within a data row that starts a range of character positions and the character position that ends a range of character positions.

The following diagram shows the syntax of the character-position FILE statement.



Element	Purpose	Key Considerations
<i>-end</i>	Indicates the character position within a data row that ends a range of character positions	A hyphen must precede the <i>end</i> value.
<i>fieldn</i>	Assigns a name to the data field that you are defining with the range of character positions	None.
<i>filename</i>	Specifies the name of the input file	None.
<i>null string</i>	Specifies the data value for which dbload must substitute a null value	Must be a quoted string.
<i>start</i>	Indicates the character position within a data row that starts a range of character positions. If you specify <i>start</i> without <i>end</i> , it represents a single character.	None.

You can repeat the same character position in a data-field definition or in different fields.

The *null string* scope of reference is the data field for which you define it. You can define an explicit null string for each field that allows null entries.

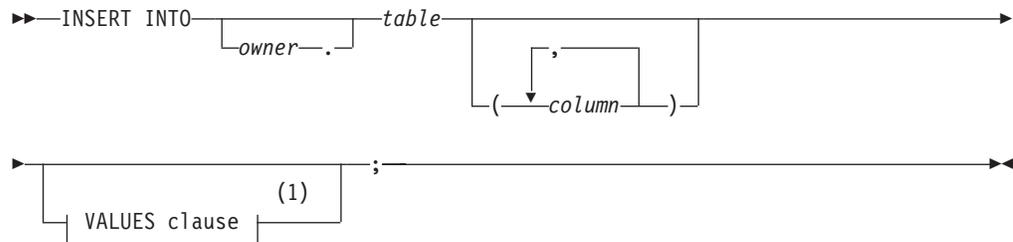
Inserted data types correspond to the explicit or default column list. If the data-field width is different from its corresponding character column, inserted values are padded with blanks if the column is wider, or inserted values are truncated if the field is wider.

If the number of columns named is fewer than the number of columns in the table, **dbload** inserts the default value that is specified for the unnamed columns. If no default value is specified, **dbload** attempts to insert a null value. If the attempt violates a not-null restriction or a unique constraint, the insert fails, and an error message is returned.

If the INSERT statement omits the column names, the default INSERT specification is every column in the named table. If the INSERT statement omits the VALUES clause, the default INSERT specification is every field of the previous FILE statement.

An error results if the number of column names listed (or implied by default) does not match the number of values listed (or implied by default).

The syntax of **dbload** INSERT statements resembles INSERT statements in SQL, except that in **dbload**, INSERT statements cannot incorporate SELECT statements. The following diagram shows the syntax of the **dbload** INSERT statement for character-position form.



Notes:

- 1 See *IBM Informix Guide to SQL: Syntax*

Element	Purpose	Key Considerations
<i>column</i>	Specifies the column that receives the new data	None.
<i>owner.</i>	Specifies the user name of the table owner	None.
<i>table</i>	Specifies the table that receives the new data	None.

The syntax for character-position form is identical to the syntax for delimiter form.

The user who runs **dbload** with this command file must have the Insert privilege on the named table.

How to Write a dbload Command File in Character-Position Form

Command files must define data fields and use character positions to define the length of each field.

The FILE statement in the following example defines six data fields from the **cust_loc_data** table data rows.

```
FILE cust_loc_data
  (city 1-15,
   state 16-17,
   area_cd 23-25 NULL = 'xxx',
   phone 23-34 NULL = 'xxx-xxx-xxxx',
   zip 18-22,
   state_area 16-17 : 23-25);
INSERT INTO cust_address (col1, col3, col4)
VALUES (city, state, zip);
```

The statement names the fields and uses character positions to define the length of each field. Compare the FILE statement in the preceding example with the data rows in the following figure.



Figure 9-2. A Sample Data File

The FILE statement defines the following data fields, which are derived from the data rows in the sample data file.

Column	Values from Data Row 1	Values from Data Row 2
city	Sunnyvale+++++	Tempe+++++++
state	CA	AZ
area_cd	408	null
phone	408-789-8075	null
zip	94086	85253
state_area	CA408	AZxxx

The null strings that are defined for the **phone** and **area_cd** fields generate the null values in those columns, but they do not affect the values that are stored in the **state_area** column.

The INSERT statement uses the field names and values that are derived from the FILE statement as the value-list input. Consider the following INSERT statement:

```
INSERT INTO cust_address (col1, col3, col4)
VALUES (city, state, zip);
```

The INSERT statement uses the data in the sample data file and the FILE statement to put the following information into the **cust_address** table.

Column	Values from Data Row 1	Values from Data Row 2
col1	Sunnyvale+++++	Tempe+++++++
col2	null	null
col3	CA	AZ
col4	94086	85253

Because the second column (**col2**) in **cust_address** is not named, the new data row contains a null (assuming that the column permits nulls).

Consider the following INSERT statement:

```
INSERT INTO cust_sort
VALUES (area_cd, zip);
```

This INSERT statement inserts the following data rows into the **cust_sort** table.

Column	Values from Data Row 1	Values from Data Row 2
col1	408	null
col2	94086	85253

Because no column list is provided, **dbload** reads the names of all the columns in **cust_sort** from the system catalog. (You cannot insert data into a temporary table because temporary tables are not entered into the system catalog.) Field names from the previous FILE statement specify the values to load into each column. You do not need one FILE statement for each INSERT statement.

Command File to Load Complex Data Types (Version 9.21 or Later Versions)

You can create **dbload** command files that load columns containing complex data types into tables.

You can use **dbload** with the following data types:

- A BLOB or CLOB
- A SET inside a ROW type

The **dbload** utility does not work with the following data types:

- A CLOB or BLOB inside a ROW type
- A ROW type inside a SET

Important: All the load utilities (**dbexport**, **dbimport**, **dbload**, **onload**, **onunload**, and **onxfer**) rely on an export and import function. If you do not define this function when you write a user-defined data type, you cannot use these utilities.

Loading a new data type inside another data type can cause problems if the representation of the data contains handles. If a string represents the data, you must be able to load it.

You can use **dbload** with named row types, unnamed row types, sets, and lists. For information, see:

- “Using the dbload Utility with Named Row Types”
- “Using the dbload Utility with Unnamed Row Types” on page 9-14
- “Using the dbload Utility with Collection Data Types” on page 9-15

Using the dbload Utility with Named Row Types

The procedure for how to use **dbload** with named row types is somewhat different than for other complex data types, because named row types are actually user-defined data types. You can follow the procedure steps for any user-defined data type.

Suppose you have a table named **person** that contains one column with a named row type. Also suppose that the **person_t** named row type contains six fields: **name**, **address**, **city**, **state**, **zip**, and **bdate**.

The following syntax shows how to create the named row type and the table used in this example:

```
CREATE ROW TYPE person_t
(
    name VARCHAR(30) NOT NULL,
    address VARCHAR(20),
    city VARCHAR(20),
    state CHAR(2),
    zip VARCHAR(9),
    bdate DATE
);
CREATE TABLE person OF TYPE person_t;
```

To load data for a named row type

1. Use the UNLOAD statement to unload the table to an input file. In this example, the input file sees the named row type as six separate fields:

```
Brown, James|13 First St.|San Francisco|CA|94070|01/04/1940|
Karen Smith|1820 Elm Ave #100|Fremont|CA|94502|01/13/1983|
```

2. Use the **dbschema** utility to capture the schema of the table and the row type. You must use the **dbschema -u** option to pick up the named row type.

```
dbschema -d stores_demo -u person_t > schema.sql
dbschema -d stores_demo -t person > schema.sql
```

3. Use DB–Access to re-create the **person** table in the new database.

For detailed steps, see “Use dbschema Output as DB–Access Input” on page 10-12.

4. Create the **dbload** command file. This **dbload** command file inserts two rows into the **person** table in the new database.

```
FILE person.unl DELIMITER '|' 6;
INSERT INTO person;
```

This **dbload** example shows how to insert new data rows into the **person** table. The number of rows in the INSERT statement and the **dbload** command file must match:

```
FILE person.unl DELIMITER '|' 6;
INSERT INTO person
VALUES ('Jones, Richard', '95 East Ave.',
        'Philadelphia', 'PA',
        '19115',
        '03/15/97');
```

5. Run the **dbload** command:

```
dbload -d newdb -c uds_command -l errlog
```

Tip: To find the number of fields in an unloaded table that contains a named row type, count the number of fields between each vertical bar (|) delimiter.

Using the dbload Utility with Unnamed Row Types

You can use the **dbload** utility with unnamed row types, which are created with the ROW constructor and define the type of a column or field.

In the following example, the **devtest** table contains two columns with unnamed row types, **s_name** and **s_address**. The **s_name** column contains three fields: **f_name**, **m_init**, and **l_name**. The **s_address** column contains four fields: **street**, **city**, **state**, and **zip**.

```
CREATE TABLE devtest
(
s_name ROW(f_name varchar(20), m_init char(1), l_name varchar(20)
not null),
s_address ROW(street varchar(20), city varchar(20), state char(20),
zip varchar(9)
);
```

The data from the **devtest** table is unloaded into the **devtest.unl** file. Each data row contains two delimited fields, one for each unnamed row type. The ROW constructor precedes each unnamed row type, as follows:

```
ROW('Jim','K','Johnson')|ROW('10 Grove St.','Eldorado','CA','94108')|
ROW('Maria','E','Martinez')|ROW('2387 West Wilton
Ave.','Hershey','PA','17033')|
```

This **dbload** example shows how to insert data that contains unnamed row types into the **devtest** table. Put double quotes around each unnamed row type or the insert will not work.

```
FILE devtest.unl DELIMITER '|' 2;
INSERT INTO devtest (s_name, s_address)
VALUES ("row('Stephen', 'M', 'Wu')",
"row('1200 Grand Ave.', 'Richmond', 'OR', '97200')");
```

Using the dbload Utility with Collection Data Types

You can use the **dbload** utility with collection data types such as SET, LIST, and MULTISSET.

SET Data Type Example

The SET data type is an unordered collection type that stores unique elements. The number of elements in a SET data type can vary, but no nulls are allowed.

The following statement creates a table in which the **children** column is defined as a SET:

```
CREATE TABLE employee
(
  name char(30),
  address char(40),
  children SET (varchar(30) NOT NULL)
);
```

The data from the **employee** table is unloaded into the **employee.unl** file. Each data row contains four delimited fields. The first set contains three elements (**Karen, Lauren, and Andrea**), whereas the second set contains four elements. The SET constructor precedes each SET data row.

```
Muriel|5555 SW Merry
Sailing Dr.|02/06/1926|SET{'Karen','Lauren','Andrea'}|
Larry|1234 Indian Lane|07/31/1927|SET{'Martha',
'Melissa','Craig','Larry'}|
```

This **dbload** example shows how to insert data that contains SET data types into the **employee** table in the new database. Put double quotes around each SET data type or the insert does not work.

```
FILE employee.unl DELIMITER '|' 4;
INSERT INTO employee
VALUES ('Marvin', '10734 Pardee', '06/17/27',
"SET{'Joe', 'Ann'}");
```

LIST Data Type Example

The LIST data type is a collection type that stores ordered, non-unique elements; that is, it allows duplicate element values.

The following statement creates a table in which the **month_sales** column is defined as a LIST:

```
CREATE TABLE sales_person
(
  name CHAR(30),
  month_sales LIST(MONEY NOT NULL)
);
```

The data from the **sales_person** table is unloaded into the **sales.unl** file. Each data row contains two delimited fields, as follows:

```
Jane Doe|LIST{'4.00','20.45','000.99'}|
Big Earner|LIST{'0000.00','00000.00','999.99'}|
```

This **dbload** example shows how to insert data that contains LIST data types into the **sales_person** table in the new database. Put double quotes around each LIST data type or the insert does not work.

```
FILE sales_person.unl DELIMITER '|' 2;  
INSERT INTO sales_person  
VALUES ('Jenny Chow', "{587900, 600000}");
```

You can load multisets in a similar manner.

Chapter 10. The `dbschema` Utility

This chapter contains information about the `dbschema` utility and how to use it.

Overview of the `dbschema` Utility

The `dbschema` utility prints the SQL statements that are necessary to replicate a specified table, view, or database. It also shows the distributions that the `UPDATE STATISTICS` statement creates.

You can use the `dbschema` utility for the following purposes:

- To display the SQL statements (the *schema*) that are required to replicate a database or a specific table, view, or procedure
- To display the schema for the Information Schema views
- To display the distribution information that is stored for one or more tables in the database
- To display information on user-defined data types and row types

After you obtain the schema of a database, you can redirect the `dbschema` output to a file that you can use with DB-Access.

Attention: 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.

Object Modes and Violation Detection

The output from the `dbschema` utility shows object modes and supports violation detection.

The `dbschema` output shows:

- The names of not-null constraints after the not-null specifications.
You can use the output of the utility as input to create another database. If the same names were not used for not-null constraints in both databases, problems could result.
- The object mode of objects that are in the disabled state. These objects can be constraints, triggers, or indexes.
- The object mode of objects that are in the filtering state. These objects can be constraints or unique indexes.
- The violations and diagnostics tables that are associated with a base table (if violations and diagnostics tables were started for the base table).

For more information about object modes and violation detection, see the `SET`, `START VIOLATIONS TABLE`, and `STOP VIOLATIONS TABLE` statements in the *IBM Informix Guide to SQL: Syntax*.

Guidelines for Using `dbschema`

You can use delimited identifiers with the `dbschema` utility. The utility detects database objects that are keywords, mixed case, or that have special characters, and the utility places double quotation marks around those keywords.

Element	Purpose	Additional Information
-d database	Specifies the database or co-server to which the schema applies. The <i>database</i> can be on a remote database server.	References: If you want to qualify the name of the <i>database</i> , see the "Database Name" section of the <i>IBM Informix Guide to SQL: Syntax</i> .
<i>filename</i>	Specifies the name of the file that will contain the dbschema output	If you omit a file name, dbschema sends output to the screen. If you specify a file name, dbschema creates a file named <i>filename</i> to contain the dbschema output.
-hd	Displays the distribution as data values	References: For more information, see "Distribution Information for Tables" on page 10-10.
-it	Displays the isolation type.	None.
-l	Displays the lock level.	None.
-seq sequence	Generates the DDL statement to define the specified <i>sequence</i> object	References: For more information, see "Sequence Creation" on page 10-6.
-ss	Generates server-specific information	This option is ignored if no table schema is generated. References: For more information, see "dbschema Server-Specific Information" on page 10-4.
-si	Excludes the generation of index storage clauses for non-fragmented tables This option is available only with the -ss option.	References: For more information, see "dbschema Server-Specific Information" on page 10-4.
-u	Prints the definitions of functions, casts, and user-defined data types	References: For more information, see "User-Defined and Complex Data Types (Version 9.21 or Later Versions)" on page 10-5.
-ui	Prints the definitions of user-defined data types, including type inheritance	References: For more information, see "User-Defined and Complex Data Types (Version 9.21 or Later Versions)" on page 10-5.
-V	Extends the -V option to display additional information on the build operating system, build number, and build date	None.
-version	Displays the software version number and the serial number	
-x	Expands dbslice names into dbspace name lists in -ss output	If the dbspace contains multiple partitions, dbspace partition names appear in the output.

You must be the DBA or have the Connect or Resource privilege for the database before you can run **dbschema** on it.

Database Schema Creation

You can create the schema for an entire database or for a portion of the database.

The options for **dbschema** allow you to perform the following actions:

- Display CREATE SYNONYM statements by owner, for a specific table or for the entire database.
- Display the CREATE TABLE, CREATE VIEW, CREATE FUNCTION, or CREATE PROCEDURE statement for a specific table or for the entire database.
- Display all GRANT privilege statements that affect a specified user or that affect all users for a database or a specific table. The user can be either a user name or role name.

- Starting with Dynamic Server Version 9.20, display user-defined and row data types with or without type inheritance.
- Starting with Dynamic Server Version 9.20, display the CREATE SEQUENCE statement defining the specified *sequence* object, or defining all sequence objects in the database.

When you use **dbschema** and specify only the database name, it is equivalent to using **dbschema** with all its options (except for the **-hd** and **-ss** options). In addition, if Information Schema views were created for the database, this schema is shown. For example, the following two commands are equivalent:

```
dbschema -d stores_demo
dbschema -s all -p all -t all -f all -d stores_demo
```

SERIAL fields included in CREATE TABLE statements that **dbschema** displays do not specify a starting value. New SERIAL fields created with the schema file have a starting value of 1, regardless of their starting value in the original database. If this value is not acceptable, you must modify the schema file.

Creating Schemas for Databases Across a UNIX or Linux Network

You can specify a database on any accessible non-SE Informix database server with the **-d** database syntax.

The following command displays the schema for the **stores_demo** database on the **finland** database server on the UNIX or Linux system console:

```
dbschema -d //finland/stores_demo
```

Changing the Owner of an Object

The **dbschema** utility uses the *owner.object* convention when it generates any CREATE TABLE, CREATE INDEX, CREATE SYNONYM, CREATE VIEW, CREATE SEQUENCE, CREATE PROCEDURE, CREATE FUNCTION, or GRANT statement, and when it reproduces any unique, referential, or check constraint. As a result, if you use the **dbschema** output to create a new object (table, index, view, procedure, constraint, sequence, or synonym), the owner of the original object owns the new object. If you want to change the owner of the new object, you must edit the **dbschema** output before you run it as an SQL script.

You can use the output of **dbschema** to create a new function if you also specify the path name to a file in which compile-time warnings are stored. This path name is displayed in the **dbschema** output.

For more information about the CREATE TABLE, CREATE INDEX, CREATE SYNONYM, CREATE VIEW, CREATE SEQUENCE, CREATE PROCEDURE, CREATE FUNCTION, and GRANT statements, see the *IBM Informix Guide to SQL: Syntax*.

dbschema Server-Specific Information

The **dbschema -ss** option generates server-specific information. In all Informix database servers except SE, the **-ss** option always generates the lock mode, extent sizes, and the dbspace name if the dbspace name is different from the database dbspace. In addition, if tables are fragmented, the **-ss** option displays information about the fragmentation strategy.

When you specify the **dbschema -ss** option, the output also displays any GRANT FRAGMENT statements that are issued for a particular user or in the entire schema.

The **-si option**, which is available only with the **-ss** option, excludes the generation of index storage clauses for non-fragmented tables.

The **-x** option expands dbslice names into dbspace name lists in **-ss** output.

Important: Use the **dbschema -ss** option to obtain information specific to a database server, including fragmentation and storage options.

If the dbspace contains multiple partitions, dbspace partition names appear in the output.

For information about fragment-level authority, see the GRANT FRAGMENT and REVOKE FRAGMENT statements in the *IBM Informix Guide to SQL: Syntax*.

User-Defined and Complex Data Types (Version 9.21 or Later Versions)

When you specify the **dbschema -u** option, the output displays the definitions of any user-defined and complex data types that the database contains. The suboption **i** lets you display the type inheritance.

The following command displays all the user-defined and complex data types for the **stork** database:

```
dbschema -d stork -u all
```

Output from **dbschema** that ran with the specified option **-u all** might appear as the following example shows:

```
create row type 'informix'.person_t
(
  name varchar(30, 10) not null,
  address varchar(20, 10),
  city varchar(20, 10),
  state char(2),
  zip integer,
  bdate date
);
create row type 'informix'.employee_t
(
  salary integer,
  manager varchar(30, 10)
) under person_t;
```

The following command displays the user-defined and complex data types, as well as their type inheritance for the **person_t** table in the **stork** database:

```
dbschema -d stork -ui person_t
```

Output from **dbschema** that ran with the option **-ui person_t** might appear as the following example shows:

```
create row type 'informix'.person_t
(
  name varchar(30, 10) not null,
  address varchar(20, 10),
  city varchar(20, 10),
  state char(2),
  zip integer,
```

```

        bdate date
    );
create row type 'informix'.employee_t
(
    salary integer,
    manager varchar(30, 10)
) under person_t;
create row type 'informix'.sales_rep_t
(
    rep_num integer,
    region_num integer,
    commission decimal(16),
    home_office boolean
) under employee_t;

```

Sequence Creation

The **dbschema -seq *sequence*** option generates information on sequence creation.

The following syntax diagram fragment shows sequence creation.



Element	Purpose	Key Considerations
-seq <i>sequence</i>	Displays the CREATE SEQUENCE statement defining <i>sequence</i>	None.
-seq all	Displays all CREATE SEQUENCE statements for the database	None.

Running **dbschema** with option **-seq** sequitur might produce this output:

```
CREATE SEQUENCE sequitur INCREMENT 10 START 100 NOCACHE CYCLE
```

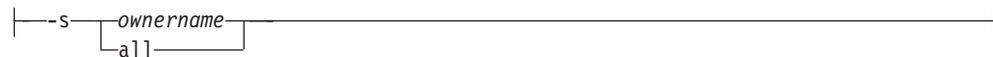
For more information about the CREATE SEQUENCE statement, see the *IBM Informix Guide to SQL: Syntax*.

Synonym Creation

The **dbschema -s** option generates information on synonym creation.

The following syntax diagram fragment shows the creation of synonyms.

Synonyms:



Element	Purpose	Key Considerations
-s <i>ownername</i>	Displays the CREATE SYNONYM statements owned by <i>ownername</i>	None.
-s all	Displays all CREATE SYNONYM statements for the database, table, or view specified	None.

Output from **dbschema** that ran with the specified option **-s alice** might appear as the following example shows:

```
CREATE SYNONYM 'alice'.cust FOR 'alice'.customer
```

For more information about the CREATE SYNONYM statement, see the *IBM Informix Guide to SQL: Syntax*.

Privileges

The **dbschema -p** option generates information on privileges.

The following syntax diagram fragment shows privileges information.

Privileges:



Element	Purpose	Key Considerations
-p user	Displays the GRANT statements that grant privileges to <i>user</i> , where <i>user</i> is a user name or role name. Specify only one user or role	You cannot specify a specific list of users with the -p option. You can specify either one user or role, or all users and roles.
-p all	Displays the GRANT statements for all users for the database, table, or view specified, or to all roles for the table specified	None.

The output also displays any GRANT FRAGMENT statements that are issued for a specified user or role or (with the **all** option) for the entire schema.

Granting Privileges

You can generate **dbschema** information on the grantor of a GRANT statement.

In the **dbschema** output, the AS keyword indicates the grantor of a GRANT statement. The following example output indicates that **norma** issued the GRANT statement:

```
GRANT ALL ON 'tom'.customer TO 'claire' AS 'norma'
```

When the GRANT and AS keywords appear in the **dbschema** output, you might need to grant privileges before you run the **dbschema** output as an SQL script. Referring to the previous example output line, the following conditions must be true before you can run the statement as part of a script:

- User **norma** must have the Connect privilege to the database.
- User **norma** must have all privileges WITH GRANT OPTION for the table **tom.customer**.

For more information about the GRANT, GRANT FRAGMENT, and REVOKE FRAGMENT statements, see the *IBM Informix Guide to SQL: Syntax*.

Displaying Privilege Information for a Role

You can generate **dbschema** information on the privileges that were granted for a particular role.

A *role* is a classification with privileges on database objects granted to the role. The DBA can assign the privileges of a related work task, such as an engineer, to a role

and then grant that role to users, instead of granting the same set of privileges to every user. After a role is created, the DBA can use the GRANT statement to grant the role to users or to other roles.

For example, issue the following **dbschema** command and to display privileges that were granted for the **calen** role.

```
sharky% dbschema -p calen -d stores_demo
```

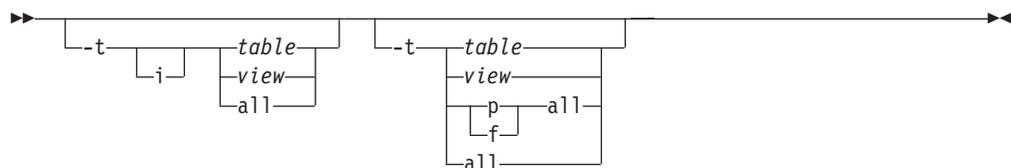
An example of information the **dbschema** utility displays is:

```
grant alter on table1 to 'calen'
```

Table, View, or Procedure Creation

You can generate **dbschema** information that shows the create of tables, views, and procedures.

The following syntax diagram shows the creation of tables, views and procedures.



Element	Purpose	Key Considerations
-f all	Limits the SQL statement output to those statements that are needed to replicate all functions and procedures	None.
-f function	Limits the SQL statement output to only those statements that are needed to replicate the specified function	None.
-f procedure	Limits the SQL statement output to only those statements that are needed to replicate the specified procedure	None.
-ff all	Limits the SQL statement output to those statements that are needed to replicate all functions	None.
-fp all	Limits the SQL statement output to those statements that are needed to replicate all procedures	None.
-t table	Limits the SQL statement output to only those statements that are needed to replicate the specified table	None.
-t view	Limits the SQL statement output to only those statements that are needed to replicate the specified view	None.
-t all	Includes in the SQL statement output all statements that are needed to replicate all tables and views	None.
-ti table	Includes in the SQL statement output all statements that are needed to replicate all table levels	None.
-ti all	Includes in the SQL statement output all statements that are needed to replicate all tables and views Functionally equivalent to -t all .	None.

For more information about the CREATE PROCEDURE and CREATE FUNCTION statements, see the *IBM Informix Guide to SQL: Syntax*.

Table Information

You can use the **dbschema -ss** option to retrieve information about fragmented tables, the lock mode, and extent sizes.

The following **dbschema** output shows the expressions specified for fragmented table.

```
{ TABLE "sallyc".t1 row size = 8 number of columns = 1 index size = 0 }
create table "sallyc".t1
(
c1 integer
) fragment by expression
(c1 < 100 ) in db1 ,
((c1 >= 100 ) AND (c1 < 200 ) ) in db2 ,
remainder in db4
extent size 16 next size 16 lock mode page;
revoke all on "sallyc".t1 from "public";
```

The following **dbschema** output shows information on partitions in partition-fragmented tables.

```
DBSCHEMA Schema Utility                                grant dba to "sqlqa";
{ TABLE "sqlqa".t1 row size = 24 number of columns = 2 index size = 13 }
create table "sqlqa".t1
(
c1 integer,
c2 char(20)
)
fragment by expression
partition part_1 (c1 = 10 ) in dbs1 ,
partition part_2 (c1 = 20 ) in dbs1 ,
partition part_3 (c1 = 30 ) in dbs1 ,
partition part_4 (c1 = 40 ) in dbs1 ,
partition part_5 (c1 = 50 ) in dbs1
extent size 16 next size 16 lock mode page;
```

Role Creation

You can use the **dbschema -ss** option to generate information on the creation of roles.

The following syntax diagram shows the creation of roles.

Roles:



Element	Purpose	Key Considerations
-r role	Displays the CREATE ROLE and GRANT statements that are needed to replicate and grant the specified role.	You cannot specify a list of users or roles with the -r option. You can specify either one role or all roles. SE does not support the -r option.
-r all	Displays all CREATE ROLE and GRANT statements that are needed to replicate and grant all roles.	None

The following **dbschema** command and output show that the role **calen** was created and was granted to **cathl**, **judith**, and **sallyc**:

```
sharky% dbschema -r calen -d stores_demo
```

```
DBSCHEMA Schema Utility
Software Serial Number RDS#N000000
create role calen;

grant calen to cathl with grant option;
grant calen to judith ;
grant calen to sallyc ;
```

Distribution Information for Tables

You can use the **dbschema -hd** option with the name of the table to retrieve the distribution information that is stored for a table in a database. If you specify the ALL keyword for the table name, the distributions for all the tables in the database are displayed.

During the **dbimport** operation, distribution information is created automatically for leading indexes on non-opaque columns. Run the UPDATE STATISTICS statement in MEDIUM or HIGH mode to create distribution information on tables that have the following types of indexes:

- Virtual Index Interface (VII) or function indexes
- Indexes on columns of user-defined data types
- Indexes on columns of built-in opaque data types (such as BOOLEAN or LVARCHAR)

Output from the **dbschema** utility shows distribution information if you used the SAMPLING SIZE keywords when UPDATE STATISTICS in MEDIUM or HIGH mode ran on the table.

For information about using the UPDATE STATISTICS statement, see the *IBM Informix Guide to SQL: Syntax*.

The output of **dbschema** for distributions is provided in the following parts:

- Distribution description
- Distribution information
- Overflow information

Each section of **dbschema** output is explained in the following sections. As an example, the discussion uses the following distribution for the fictional table called **invoices**. This table contains 165 rows, including duplicates.

You can generate the output for this discussion with a call to **dbschema** that is similar to the following example:

```
dbschema -hd invoices -d pubs_stores_demo
```

Example of Output Showing Distribution Information

The **dbschema** output can show the data distributions have been created for the specified table and the date when the UPDATE STATISTICS statement that generated the distributions ran.

The follow example of **dbschema** output shows distribution information.

```
Distribution for cath1.invoices.invoice_num
```

```
High Mode, 10.000000 Resolution
```

```
--- DISTRIBUTION ---
```

```
(
1: ( 16, 7, 11)
2: ( 16, 6, 17)
3: ( 16, 8, 25)
4: ( 16, 8, 38)
5: ( 16, 7, 52)
6: ( 16, 8, 73)
7: ( 16, 12, 95)
8: ( 16, 12, 139)
9: ( 16, 11, 182)
10: ( 10, 5, 200)
```

```
--- OVERFLOW ---
```

```
1: ( 5, 56)
2: ( 6, 63)
}
```

Description of the Distribution Information in the Example

The first part of the sample **dbschema** output describes which data distributions have been created for the specified table. The name of the table is stated in the following example:

```
Distribution for cath1.invoices.invoice_num
```

The output is for the **invoices** table, which is owned by user **cath1**. This data distribution describes the column **invoice_num**. If a table has distributions that are built on more than one column, **dbschema** lists the distributions for each column separately.

The **dbschema** output also shows the date when the **UPDATE STATISTICS** statement that generated the distributions ran. You can use this date to tell how outdated your distributions are.

The last line of the description portion of the output describes the mode (**MEDIUM** or **HIGH**) in which the distributions were created, and the resolution. If you create the distributions with medium mode, the confidence of the sample is also listed. For example, if the **UPDATE STATISTICS** statement runs in **HIGH** mode with a resolution of 10, the last line appears as the following example shows:

```
High Mode, 10.000000 Resolution
```

Distribution Information

The distribution information in **dbschema** output describes the bins that are created for the distribution, the range of values in the table and in each bin, and the number of distinct values in each bin.

Consider the following example:

```
(
1: ( 16, 7, 11)
2: ( 16, 6, 17)
3: ( 16, 8, 25)
4: ( 16, 8, 38)
5: ( 16, 7, 52)
6: ( 16, 8, 73)
```

```

7: ( 16, 12, 95)
8: ( 16, 12, 139)
9: ( 16, 11, 182)
10: ( 10, 5, 200)

```

The first value in the rightmost column is the smallest value in this column. In this example, it is 5.

The column on the left shows the bin number, in this case 1 through 10. The first number in parentheses shows how many values are in the bin. For this table, 10 percent of the total number of rows (165) is rounded down to 16. The first number is the same for all the bins except for the last. The last row might have a smaller value, indicating that it does not have as many row values. In this example, all the bins contain 16 rows except the last one, which contains 10.

The middle column within the parentheses indicates how many distinct values are contained in this bin. Thus, if there are 11 distinct values for a 16-value bin, it implies that one or more of those values are duplicated at least once.

The right column within the parentheses is the highest value in the bin. The highest value in the last bin is also the highest value in the table. For this example, the highest value in the last bin is 200.

Overflow Information

The last portion of the **dbschema** output shows values that have many duplicates.

The number of duplicates of indicated values must be greater than a critical amount that is determined as approximately 25 percent of the resolution times the number of rows. If left in the general distribution data, the duplicates would skew the distribution, so they are moved from the distribution to a separate list, as the following example shows:

```

--- OVERFLOW ---
1: ( 5, 56)
2: ( 6, 63)

```

For this example, the critical amount is $0.25 * 0.10 * 165$, or 4.125. Therefore, any value that is duplicated five or more times is listed in the overflow section. Two values in this distribution are duplicated five or more times in the table: the value 56 is duplicated five times, and the value 63 is duplicated six times.

Use dbschema Output as DB-Access Input

You can use the **dbschema** utility to get the schema of a database and redirect the **dbschema** output to a file. Later, you can import the file into DB-Access and use DB-Access to re-create the schema in a new database.

Inserting a Table into a Database Example

You can insert CREATE TABLE statements into the **dbschema** output file and use this output as DB-Access input.

The following example copies the CREATE TABLE statements for the customer table into the **dbschema** output file, **tab.sql**:

```
dbschema -d db -t customer > tab.sql
```

Remove the header information about **dbschema** from the output file, **tab.sql**, and then use DB–Access to re-create the table in another database, as follows:

```
dbaccess db1 tab.sql
```

Re-Creating the Schema of a Database

You can use **dbschema** and DB–Access to save the schema from a database and then re-create the schema in another database. A **dbschema** output file can contain the statements for creating an entire database.

To save a database schema and re-create the database

1. Use **dbschema** to save the schema to an output file, such as **db.sql**:

```
dbschema -d db > db.sql
```

You can also use the **-ss** option to generate server-specific information:

```
dbschema -d db -ss > db.sql
```

2. Remove the header information about **dbschema**, if any, from the output file.
3. Add a CREATE DATABASE statement at the beginning of the output file or use DB–Access to create a new database.
4. Use DB–Access to re-create the schema in a new database:

```
dbaccess - db.sql
```

When you use **db.sql** to create a database on a different database server, confirm that dbspaces exist.

The databases **db** and **testdb** differ in name but have the same schema.

Chapter 11. The LOAD and UNLOAD Statements

This chapter shows the syntax of the SQL UNLOAD and LOAD statements.

If the database contains label-based access control (LBAC) objects, you can load or unload only those rows in which your security label dominates the column-security label or the row-security label. If entire table is to be loaded or unloaded, you must have the necessary LBAC credentials for writing/reading all of the labeled rows and columns. For more information on LBAC objects, see the *IBM Informix Security Guide* and the *IBM Informix Guide to SQL: Syntax*.

Overview of the LOAD and UNLOAD Statements

The LOAD statement is moderately fast and easy to use, but it only accepts specified data formats. You usually use LOAD with data that is prepared with an UNLOAD statement.

You can use the UNLOAD statement in DB–Access to unload selected rows from a table into a text file.

To load tables, use LOAD or **dbload**. To manipulate a data file that you are loading or to access a database while it is loading, use the **dbload** utility. The cost of the flexibility is the time you spend creating the **dbload** command file and slower execution. When possible, use the LOAD statement, which is faster than **dbload**.

Syntax of the UNLOAD Statement

The UNLOAD statement in DB–Access unloads selected rows from a table into a text file.

```
►► UNLOAD TO 'filename' [ DELIMITER 'delimiter' ] SELECT Statement (1) ►►
```

Notes:

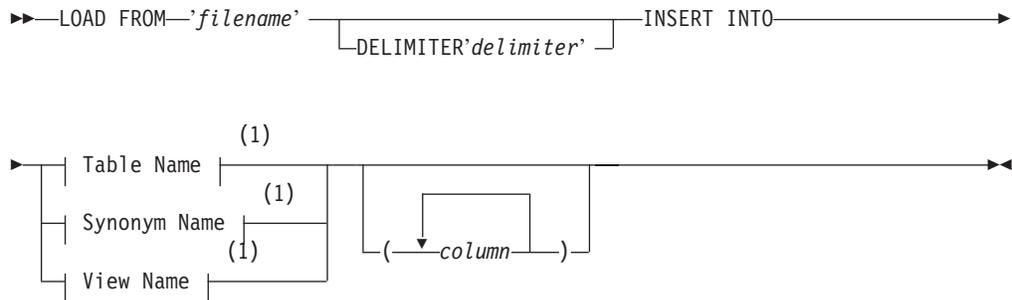
1 See the *IBM Informix Guide to SQL: Syntax*.

Element	Purpose	Key Considerations
<i>delimiter</i>	Character to use as delimiter	Requirements: See "Syntax for the Delimiter Form" on page 9-6
<i>filename</i>	Specifies the input file	None.

This syntax diagram is only for quick reference. For details about the syntax and use of the UNLOAD statement, see the *IBM Informix Guide to SQL: Syntax*.

Syntax of the LOAD Statement

The LOAD statement in DB–Access appends rows to an existing table of a database.



Notes:

- 1 See *IBM Informix Guide to SQL: Syntax*

Element	Purpose	Key Considerations
<i>column</i>	The name of a column to receive data from <i>filename</i>	Must be a column in the specified table or view.
<i>delimiter</i>	Character to use as delimiter	See "Syntax for the Delimiter Form" on page 9-6
<i>filename</i>	Specifies the input file	None.

This syntax diagram is only for quick reference. For details about the syntax and use of the LOAD statement, see the *IBM Informix Guide to SQL: Syntax*.

Load and Unload Statements for Locales That Support Multibyte Code Sets

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.

Load and Unload Statements for Non-Default Locales and GL_DATETIME Environment Variable

If the database uses a non-default locale and the GL_DATETIME environment variable has a non-default setting, you must set the USE_DTENV environment variable to the value of 1 before you can process localized datetime values correctly with the LOAD and UNLOAD statements, or with the **dbimport** and **dbexport** utilities.

Chapter 12. The onmode Utility Reversion Option

This chapter contains information about how to use the **-b** option of the **onmode** utility to revert to the older database server from which you converted.

Use of the onmode -b Command for Reversion

When you convert a database server, several modifications make the format of the databases incompatible with the older version. The **onmode** utility modifies the data in an Informix database so that the earlier version of the database server can access it.

The utility does not revert changes made to the layout of the data that do not affect compatibility.

You must revert the databases before users can access the data with the earlier database server version.

UNIX/Linux Only

You must be user **root** or user **informix** to run **onmode**.

Windows Only

You must be a member of the **Informix-Admin** group to run **onmode**.

For information about other **onmode** options, see your *IBM Informix Administrator's Guide*.

Preparation for Reversion

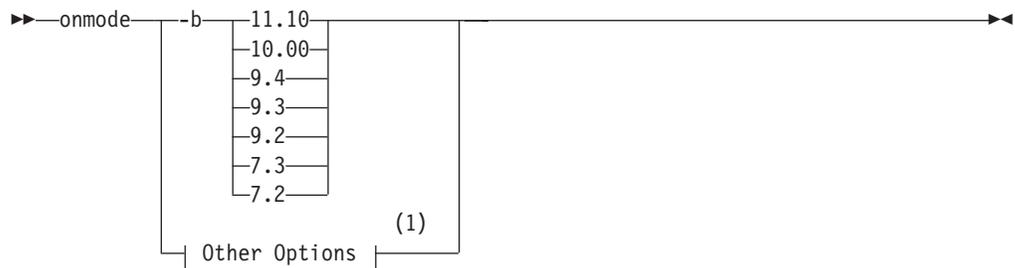
Before you use the **-b** option, notify users that you are going to bring the database server offline. The reversion utility forcibly removes all users and shuts down the database server.

The **-b** option includes an implicit **-yuk**.

Make sure that the **INFORMIXSERVER** environment variable is set to the correct database server.

Syntax of the onmode -b Command

The **onmode -b** command restores the databases to an earlier version in a format that is compatible with the earlier version.



Notes:

- 1 For all other onmode options, see your *IBM Informix Administrator's Reference*.

Element	Purpose
-b 11.10	Changes the database to the Version 11.10 format
-b 10.00	Changes the database to the Version 10.00 format
-b 9.4	Changes the database to the Version 9.40 format
-b 9.3	Changes the database to the Version 9.30 format
-b 9.2	Changes the database to the Version 9.2x format
-b 7.3	Changes the database to the Version 7.31 format
-b 7.2	Changes the database to the Version 7.24 format

Note:

1. To list the available options for your database server, type **onmode -b -**.
2. For more information, see "Reverting from Dynamic Server Version 11.50" on page 6-10.

Chapter 13. The **onunload** and **onload** Utilities

This chapter contains information about using the **onunload** and **onload** utilities to unload and load databases and tables from and into IBM Informix database servers.

Important: You can use the **onunload** and **onload** utilities with Dynamic Server 11.10, 10.00, 9.40, 9.30, or 9.21 if the databases contain only legacy data types and no extended data types. In addition, you cannot use these utilities with Dynamic Server versions that are earlier than version 7.24.

Overview of the **onunload** and **onload** Utilities

The **onunload** and **onload** utilities provide the fastest way to move data between computers that use the same database server on the same platform.

For example, your site purchases a more powerful UNIX computer to allow faster access for users. You need to transfer existing databases to the new database server on the new computer. Use **onunload** to unload data from the first database server and then use **onload** to load the data into the second database server. Both database servers must have the same version number, or they must have compatible version numbers. You can move an entire database or selected tables only, but you cannot modify the database schema.

The **onunload** utility can unload data more quickly than either **dbexport** or the UNLOAD statement because **onunload** copies the data in binary format and in page-sized units. The **onload** utility takes a tape or a file that the **onunload** utility creates and re-creates the database or the table. The **onunload** and **onload** utilities are faster than **dbimport**, **dbload**, or LOAD but are much less flexible and do not let you modify the database schema or move from one operating system or database server version to another.

You can use only **onunload** and **onload** if your answer to each of the following questions is *yes*. If your answer is *no*, you cannot use **onunload** and **onload**.

Use onunload and onload Only If Your Answer To Each Question Is Yes
Is the target database server on the same hardware platform?
Do you want to move to another database server of the same version?
Do you want to keep the existing database schema without modifying it?
Do you want to move an entire database or an entire table?
Do you want to move data between GLS and non-GLS databases?
Are the page images compatible?
Are the numeric representations the same?

Because the data is written in page-sized units, you can use **onunload** and **onload** only when certain conditions are met. For example, you cannot use **onunload** and **onload** to move data between UNIX or Linux and Windows because they use different page sizes. For example, the page size is 2 KB on some UNIX systems and 4 KB on Windows.

Requirements for Using **onload** and **onunload**

The **onload** and **onunload** utilities have limitations. You can only use these utilities to move data between database servers of the same version on the same operating system. You cannot modify the database schema, logging must be turned off, and the utilities can be difficult to use.

The **onload** and **onunload** utilities have the following requirements:

- The original database and the target database must be from the same version of the database server. You cannot use the **onload** and **onunload** utilities to move data from one version to another version.
- You cannot use **onload** and **onunload** to move data between different types of database servers.
- The **onload** command must have the same scope as the corresponding **onunload** command that unloaded the same table or tables that **onload** references. You cannot, for example, use **onunload** to unload an entire database, and then use **onload** to load only a subset of the tables from that database.
- Do not use **onload** and **onunload** to move data if the database contains extended data types. (Use the HPL instead to move the data.)
- Because the tape that **onload** reads contains binary data that is stored in disk-page-sized units, the computers where the original database resides (where you use **onunload**) and where the target database will reside (where you use **onload**) must have the same page size, the same representation of numeric data, the same byte alignment for structures and unions.
- You cannot use **onload** and **onunload** to move data between non-GLS and GLS locales.
- You cannot use **onload** and **onunload** on High-Availability Data Replication (HDR) secondary servers, remote standalone (RS) secondary servers, or shared disk (SD) secondary servers.

You can use **onunload** and **onload** to move data between databases if the NLS and GLS locales are identical. For example, if both the NLS and GLS tables were created with the same French locale, **onload** and **onunload** can move data. However, if user A has a French locale NLS table on server A and tries to load data into a German locale GLS table on server B, **onload** reports errors.

If the page sizes are different, **onload** fails. If the alignment or numeric data types on the two computers are different (for example, with the most significant byte as last instead of first, or different float-type representations), the contents of the data page could be misinterpreted.

How **onunload** and **onload** Work

The **onunload** utility, which unloads data from a database, writes a database or table into a file on tape or disk. The **onload** utility loads data that was created with the **onunload** command into the database server.

The **onunload** utility unloads the data in binary form in disk-page units, making this utility more efficient than **dbexport**.

You can use the **onunload** utility to move data between computers that have the same version of the database server.

Important: You cannot use the **onload** and **onunload** utilities to move data from one version of a database server to another or between different types of database

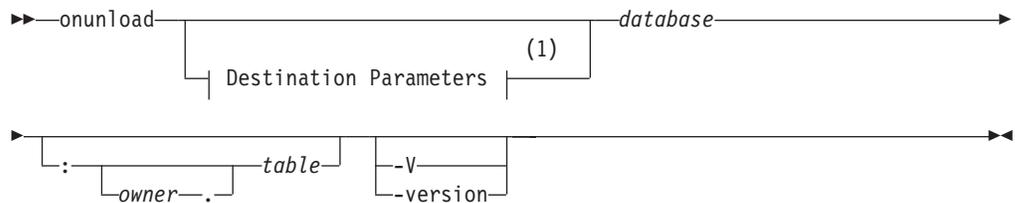
servers. In addition, the **onload** command must have the same scope as the corresponding **onunload** command that unloaded the same table or tables that **onload** references. You cannot, for example, use **onunload** to unload an entire database, and then use **onload** to load only a subset of the tables from that database.

The **onload** utility creates a database or table in a specified dbspace. The **onload** utility then loads it with data from an input tape or disk file that the **onunload** utility creates.

During the load, you can move simple large objects that are stored in a blob space to another blob space.

Syntax of the onunload Command

The **onunload** command unloads data from a database and writes a database or table into a file on tape or disk.



Notes:

- 1 See "onunload Destination Parameters" on page 13-4

Element	Purpose	Key Considerations
<i>database</i>	Specifies the name of a database	Additional Information: The database name cannot be qualified by a database server name (<i>database@dbservername</i>). References: Syntax must conform to the Identifier segment; see the <i>IBM Informix Guide to SQL: Syntax</i> .
<i>owner.</i>	Specifies the owner of the table	Additional Information: The owner name must not include invalid characters. References: For path name syntax, see your operating-system documentation.
<i>table</i>	Specifies the name of the table	Requirement: The table must exist. References: Syntax must conform to the Table Name segment; see the <i>IBM Informix Guide to SQL: Syntax</i> .

If you do not specify any destination parameter options, **onunload** uses the device that TAPEDEV specifies. The block size and tape size are the values specified as TAPEBLK and TAPESIZE, respectively. (For information about TAPEDEV, TAPEBLK, and TAPESIZE, see your *IBM Informix Dynamic Server Administrator's Reference*.)

The **-V** option displays the software version number and the serial number. The **-version** option extends the **-V** option to display additional information on the build operating system, build number, and build date.

onunload Destination Parameters

The **onunload** utility supports tape or file destination options.

The following syntax diagram fragment shows **onunload** destination parameters

Destination Parameters:



Notes:

- 1 Only one occurrence of each option allowed. More than one option can occur in a single invocation.

Element	Purpose	Key Considerations
-b <i>blocksize</i>	Specifies in kilobytes the block size of the tape device	Requirement: The <i>blocksize</i> must be an integer. Additional Information: This option overrides the default value in TAPEBLK or LTAPEBLK.
-l	Directs onunload to read the values for tape device, block size, and tape size from LTAPEDEV, LTAPEBLK, and LTAPESIZE, respectively	None.
-s <i>tapesize</i>	Specifies in kilobytes the amount of data that can be stored on the tape	Requirement: The <i>tapesize</i> must be an integer. To write to the end of the tape, specify a tape size of 0. If you do not specify 0, then the maximum <i>tapesize</i> is 2 097 151 KB. Additional Information: This option overrides the default value in TAPESIZE or LTAPESIZE.
-t <i>source</i>	Specifies the path name of the file on disk or of the tape device where the input tape is mounted	Additional Information: This option overrides the tape device specified by TAPEDEV or LTAPEDEV. The path name must be a valid path name.

Constraints That Affect onunload

When you use the **onunload** utility, you must be aware of constraints that affect how you load the data on the **onunload** tape.

The following constraints apply to **onunload**:

- You must load the data on the **onunload** tape into a database or table that your database server (excluding SE) manages.
- You can use **onunload** and **onload** with Dynamic Server 11.50, 11.10, 10.00, 9.40, 9.30, or 9.21 if the databases contain only legacy data types and no extended data types.
- You must load the tape that **onunload** writes onto a computer with the same page size and the same representation of numeric data as the original computer.

- You must read the file that **onunload** creates with the **onload** utility of the same version of your database server. You cannot use **onunload** and **onload** to move data from one version to another.
- When you unload a complete database, you cannot modify the ownership of database objects (such as tables, indexes, and views) until after you finish reloading the database.
- When you unload and load a table, **onunload** does not preserve access privileges, synonyms, views, constraints, triggers, or default values that were associated with the original tables. Before you run **onunload**, use the **dbschema** utility to obtain a listing of the access privileges, synonyms, views, constraints, triggers, and default values. After you finish loading the table, use **dbschema** to re-create the specific information for the table.

Privileges for Database or Table Unloading

To unload a database, you must have DBA privileges for the database or be user **informix**. To unload a table, you must either own the table, have DBA privileges for the database in which the table resides, or be user **informix**.

User **root** does not have special privileges with respect to **onunload** and **onload**.

Tables that are Unloaded with a Database

If you unload a database, all the tables in the database, including the system catalog tables, are unloaded.

All the triggers, SPL routines, defaults, constraints, and synonyms for all the tables in the database are also unloaded.

Data that is Unloaded with a Table

If you unload a table, **onunload** unloads the table data and information from the **sysables**, **systables**, **syscolumns**, **sysindexes**, and **sysblobs** system catalog tables.

When you unload a table, **onunload** does not unload information about constraints, triggers, or default values that are associated with a table. In addition, access privileges that are defined for the table and synonyms or views that are associated with the table are not unloaded.

Locking During Unload Operation

During the unload operation, the database or table is locked in shared mode. An error is returned if **onunload** cannot obtain a shared lock.

The **onload** utility creates a database or table in a specified dbspace (excluding SE). The **onload** utility then loads it with data from an input tape or disk file that the **onunload** utility creates.

Logging Mode

The **onunload** utility does not preserve the logging mode of a database. After you load the database with **onload**, you can make a database ANSI compliant or add logging.

For information about logging modes, refer to the *IBM Informix Guide to SQL: Syntax*.

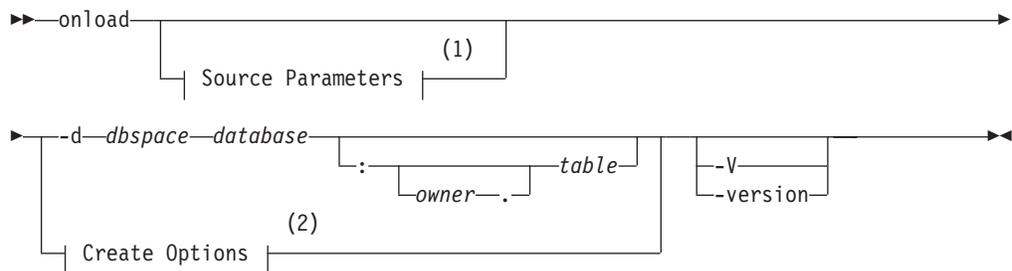
During the load, you can move simple large objects that are stored in a blob space to another blob space.

If you do not specify any source-parameter options, **onload** uses the device that is specified as TAPEDEV. The block size and tape size are the values that are specified as TAPEBLK and TAPESIZE, respectively. (For more information about TAPEDEV, TAPEBLK, and TAPESIZE, refer to your *IBM Informix Dynamic Server Administrator's Guide*.)

If you do not specify creation options, **onload** stores the database or table in the root db space.

Syntax of the onload Command

The **onload** command loads data that was created with the **onunload** command into the database server.



Notes:

- 1 See "onload Source Parameters" on page 13-7
- 2 See "onload Create Options" on page 13-7

Element	Purpose	Key Considerations
-d dbspace	Loads a database or table into the specified db space	The tape being loaded must contain the specified database or table.
database	Specifies the name of the database	The database name cannot include a database server name, such as <i>database@dbservername</i> . References: Syntax must conform to the Identifier segment; see the <i>IBM Informix Guide to SQL: Syntax</i> .
owner.	Specifies the owner of the table	The owner name must not include invalid characters. References: For path name syntax, refer to your operating-system documentation.
table	Specifies the name of the table	The table must exist. References: Syntax must conform to the Table Name segment; see the <i>IBM Informix Guide to SQL: Syntax</i> .

The **-V** option displays the software version number and the serial number. The **-version** option extends the **-V** option to display additional information on the build operating system, build number, and build date.

onload Source Parameters

The **onload** command includes options for specifying information about the tape or file source.

The following syntax diagram fragment shows **onload** source parameters.

Source Parameters:



Notes:

- 1 Only one occurrence of each option allowed. More than one option can occur in a single invocation.

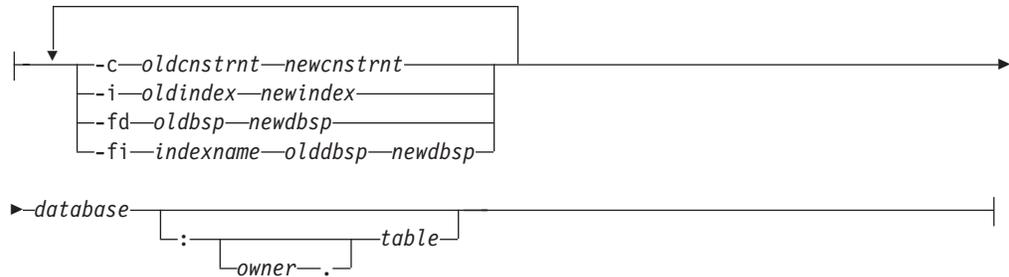
Element	Purpose	Key Considerations
-b <i>blocksize</i>	Specifies in kilobytes the block size of the tape device	<p>Requirements: Unsigned integer. Must specify the block size of the tape device.</p> <p>Additional Information: This option overrides the default value in TAPEBLK or LTAPEBLK.</p>
-l	Directs onload to read the values for tape device, block size, and tape size from the configuration parameters LTAPEDEV, LTAPEBLK, and LTAPESIZE, respectively	<p>Additional Information: If you specify -l, and then -b, -s, or -t, the value that you specify overrides the value in the configuration file.</p>
-s <i>tapesize</i>	Specifies in kilobytes the amount of data that the database server can store on the tape	<p>Requirements: Unsigned integer. To write to the end of the tape, specify a tape size of 0.</p> <p>If you do not specify 0, then the maximum <i>tapesize</i> is 2 097 151 KB.</p> <p>Additional Information: This option overrides the default value in TAPESIZE or LTAPESIZE.</p>
-t <i>source</i>	Specifies the path name of the file on disk or of the tape device where the input tape is mounted	<p>Must be a legitimate path name.</p> <p>Additional Information: This option overrides the tape device that TAPEDEV or LTAPEDEV specifies.</p> <p>References: For path name syntax, see your operating-system documentation.</p>

onload Create Options

The **onload** command included information that is used to recreate the database.

The following syntax diagram fragment shows **onload** create options.

Create Options:



Element	Purpose	Key Considerations
-c <i>oldcnstrnt</i> <i>newcnstrnt</i>	Directs onload to rename the specified constraint.	None.
-i <i>oldindex</i> <i>newindex</i>	Directs onload to rename the table index when it stores the index on disk.	Additional Information: Use the -i option to rename indexes during the load to avoid conflict with existing index names. References: Syntax must conform to the Identifier segment; see the <i>IBM Informix Guide to SQL: Syntax</i> .
-fd <i>olddbsp</i> <i>newdbsp</i>	Moves a data fragment from one dbspace to another.	The new dbspace must exist and must not already contain another data fragment for the table. Additional Information: This option is used with parallel data query (PDQ) and table fragmentation.
-fi <i>indexname</i> <i>olddbsp newdbsp</i>	Moves index fragments from one dbspace to another.	The new dbspace must exist and must not already contain another index fragment for the table. Additional Information: This option is used with PDQ and table fragmentation.
<i>database</i>	Specifies the name of the database	Requirement: The database name cannot include a database server name, such as <i>database@dbservername</i> . References: Syntax must conform to the Identifier segment; see the <i>IBM Informix Guide to SQL: Syntax</i> .
<i>owner.</i>	Specifies the owner of the table	Requirement: The owner name must not include invalid characters. References: For path name syntax, refer to your operating-system documentation.
<i>table</i>	Specifies the name of the table	Requirement: The table must not exist. References: Syntax must conform to the Table Name segment; see the <i>IBM Informix Guide to SQL: Syntax</i> .

If you do not specify any create options, the **onload** utility stores the database or table in the root dbspace.

You can use the **-c**, **-i**, **-fd**, and **-fi** options in any order and as often as necessary as long as you use unique pairs.

Constraints That Affect onload

The **onload** utility performs faster than the **dbimport**, **dbload**, or **LOAD** methods. In exchange for this higher performance, **onload** has certain constraints.

The **onload** utility has the following constraints:

- The **onload** utility only creates a new database or table; you must drop or rename an existing database or table of the same name before you run **onload**. During execution, the **onload** utility's prompt will ask you if you want to rename blobspaces.
- The **onload** utility places a shared lock on each of the tables in the database during the load. Although you cannot update a table row with the lock in place, the database is available for queries.
- When you load a complete database, the user who runs **onload** becomes the owner of the database.
- The **onload** utility creates a database without logging; you must initiate logging after **onload** loads the database.
- When you use **onload** to load a table into a logged database, you must turn off logging for the database during the operation.
- The **onload** utility does not preserve dbspace assignments of table creation or table fragmentation.

Logging During Loading

When you use **onload** to create tables from an **onunload** input tape, **onload** can only load information into a database without logging. Thus, before you load a table into an existing, logged database, you must end logging for the database.

You also might want to consider loading during off-peak hours. Otherwise, you might fill the logical-log files or consume excessive shared-memory resources. After you load the table, create a level-0 dbspace backup before you resume database logging.

When you use **onload** to create databases from an **onunload** input tape, the databases that result are not ANSI compliant and do not use transaction logging. You can make a database ANSI compliant or add logging after you load the database. (For more information about logging, refer to the *IBM Informix Guide to SQL: Reference*.)

The **onload** utility performs all its loading within a transaction. This feature allows the changes to be rolled back if an error occurs.

Movement of Simple Large Objects to a Blobspace

If you load a table that contains simple large objects stored in a blobspace, **onload** asks you if you want to move them to another blobspace.

If you respond yes, **onload** displays the blobspace name where the simple large objects were stored when the tape was created. It then asks you to enter the name of the blobspace where you want the simple large objects stored.

If you enter a valid blobspace name, **onload** moves all simple-large-object columns in the table to the new blobspace. Otherwise, **onload** prompts you again for a valid blobspace name.

Ownership and Privileges

When you load a new database, the user who runs **onload** becomes the owner. Ownership within the database (tables, views, and indexes) remains the same as when the database was unloaded to tape with **onunload**.

To load a table, you must have the Resource privilege on the database. When **onload** loads a new table, the user who runs **onload** becomes the owner unless you specify an owner in the table name. (You need the DBA privilege for the database to specify an owner in the table name.)

The **onunload** utility does not preserve synonyms or access privileges. To obtain a listing of defined synonyms or access privileges, use the **dbschema** utility, which Chapter 10, "The dbschema Utility," on page 10-1 describes, before you run **onunload**.

Exclusive Locking During Load Operation

During the load operation, **onload** places an exclusive lock on the new database or table.

Loading proceeds as a single transaction, and **onload** drops the new database or table if an error or system failure occurs.

Using onunload and onload

Use **onunload** and **onload** to move either a complete database or a table from one computer to another.

The following procedures tell you how to use the **onunload** and **onload** commands. For information on the syntax of these commands, see "Syntax of the onunload Command" on page 13-3 and "Syntax of the onload Command" on page 13-6.

To move a database from one computer to another

1. Make sure that the page size, numeric representations, and byte alignment on structures and unions are the same on both computers. (The page size is 2 KB on certain UNIX systems and 4 KB on Windows NT.) The page size is an Informix characteristic. For information about page size, see your *IBM Informix Administrator's Guide*. The numeric representation and the byte alignment are characteristics of your operating system. For information about numeric representation and byte alignment, refer to the manuals for your operating systems.
2. Decide where to store the unloaded data:
 - **On disk.** Create an empty file for **onunload** to hold the data. Make sure that you have write permission for the file.
 - **On tape.** Use the tape device and characteristics specified in the ONCONFIG configuration file by either TAPEDEV or LTAPEDEV or specify another tape device. Make sure that the tape device that you specify is available for **onunload**.
3. Run the **oncheck** utility to make sure that your database is consistent. For information about **oncheck**, see your *IBM Informix Administrator's Guide*.
4. Run the **onunload** utility to unload the data from the database.
5. If necessary, transfer the storage medium (tape or disk) to the new computer. If the two computers are on the same network, you can read or write the data remotely.
6. Run the **onload** utility to load the data into the new database.
7. Set the desired logging status for the new database. For information about logging status, see your *IBM Informix Administrator's Guide*.
8. If necessary, change the DBA privileges of the database.

9. If you want to restore the triggers, access privileges, SPL routines, defaults, constraints, and synonyms for the tables in the database, run the **dbschema** utility.
10. Create a level-0 backup of the new database.

To move a table from one computer to another

1. Make sure that the page size, numeric representations, and byte alignment on structures and unions are the same on both computers. (The page size is 2 KB on certain UNIX systems and 4 KB on Windows NT.)
2. Decide where to store the unloaded data.
3. Run the **oncheck** utility to make sure that your database is consistent.
4. If you want to save the triggers, access privileges, SPL routines, defaults, constraints, and synonyms for the table, run the **dbschema** utility.
5. Run the **onunload** utility.
6. If necessary, transfer the storage medium to the new computer.
7. If the table includes simple large objects that are stored in blobspaces, decide where to store the simple large objects. If necessary, create new blobspaces.
8. Turn off logging.
When you are loading a table, logging on the target database must be turned off. (When you are creating and loading an entire database, the logging status does not matter.)
9. Run the **onload** utility.
10. Create a level-0 backup of the modified database.
11. Turn logging back on, if you want logging.
12. If you want to restore the triggers, access privileges, SPL routines, defaults, constraints, and synonyms for the table, run the **dbschema** utility or create them manually.

To move a table from one dbspace to another dbspace on the same computer

1. Run the **onunload** utility to unload the table.
2. Turn off logging.
When you are loading a table, logging on the target database must be turned off.
3. Run the **onload** utility.
Specify a new table name and new dbspace name in the **onload** statement.
4. If the data loads successfully, delete the old table in the old dbspace and rename the new table to the old table name.
5. Create a level-0 backup of the modified database.
6. Turn logging back on, if you want logging.

Part 5. Appendixes

Appendix A. New Environment Variables

Each version of Dynamic Server contains new environment variables that could affect your installation. You might also need to adjust the values of existing environment variables.

For more information on environment variables, see the *IBM Informix Guide to SQL: Reference* and your *IBM Informix Dynamic Server Administrator's Guide*.

Table A-1 lists the new environment variables in various versions of Dynamic Server.

Table A-1. New Environment Variables

Version	Environment Variable	Description
11.50xC3	IFX_LOB_XFERSIZE	Provides error checking when transmitting large CLOB or BLOB data types from clients to the database server.
11.50	CDR_DISABLE_SPOOL	Prevents the generation of ATS and RIS files.
11.50	CDR_ATSRISNAME_DELIM	Sets the delimiter for the timestamp portion of the ATS and RIS file names.
11.50	IFX_NOT_STRICT_THOUS_SEP	Removes enforcement of the restriction that three digits must exist after the thousand separator.
11.10	IFX_AUTO_REPREPARE	Controls whether Dynamic Server automatically recompiles prepared objects and reoptimizes SPL routines that reference tables whose schemas change Enabling the IFX_AUTO_REPREPARE session environment variable can avoid many -710 errors, and can reduce the number of manual reprepare and reoptimize operations after the schema of a table is modified
11.10	IFX_NODBPROC	An environment variable that enables or prevents the execution of a sysdbopen() or sysdbclose() procedure
10.00xC4	BAR_SORT_DBS	A variable (used only in Version 10.00xC4 and later Version 10.00 fixpacks) for backup and restore operations when the scope is not the whole system.
10.0	IFX_EXTDIRECTIVES	A client-side external optimizer directive to use as a temporary solution to problems when you do not want to change SQL statements in queries
10.0	IFX_NO_TIMELIMIT_WARNING	Supports time-limited license
10.0	IFX_ONPLOAD_AUTO_UPGRADE	Automatically upgrades the onpload database the first time you start the HPL utility with the ipload or onpladm command after you migrate to a new database server version

Table A-1. New Environment Variables (continued)

Version	Environment Variable	Description
10.0	STDIO	A TAPEDEV configuration parameter variable that improves the speed of HDR setup
9.40	CDR_LOGDELTA	Determines when spooling of the Enterprise Replication queue occurs, based on the percentage of the logical log size. Use as directed by Technical Support.
9.40	CDR_PERFLOG	Enables Enterprise Replication queue tracing. Use as directed by Technical Support.
9.40	CDR_ROUTER	Determines whether intermediate processing for Enterprise Replication is allowed in a hierarchal topology. Use as directed by Technical Support.
9.40	CDR_RMSCALEFACT	Sets the maximum number of Enterprise Replication DataSync threads per CPU VP. Use as directed by Technical Support.
9.40	USETABLENAME	Disallows the use of a synonym of the table in certain SQL statements.
9.30	IFX_DEF_TABLE_LOCKMODE	Specifies the default lock mode for database tables.
9.21	JAR_TEMP_PATH	Specifies a non-default local file system location for temporary .jar files of the Java virtual machine.
9.21	JAVA_COMPILER	Disables JIT compilation.
9.21	JVM_MAX_HEAP_SIZE	Sets a non-default upper limit on the size of the heap for the Java virtual machine.
9.20	IFX_LONGID	Determines whether a given client application is capable of handling long identifiers.
9.20	IFX_UPDDESC	Allows the execution of a DESCRIBE of an UPDATE statement.
9.20	STMT_CACHE	Controls the use of the shared statement cache on a session.

In Dynamic Server 9.30, the environment variable **DELIMIDENT** must be set before a client starts to manipulate a table with an SQL DELETE statement that omits the FROM keyword.

Appendix B. New Configuration Parameters

Each version of Dynamic Server contains new configuration parameters that might affect your installation. If you need to revert to a prior version of the server, you must either replace the Dynamic Server Version 11.50 ONCONFIG configuration file with the ONCONFIG file that you used before you converted, or you must remove configuration parameters that the earlier database server does not support.

For a list of altered and removed configuration parameters, see Appendix C, “Configuration Parameters That Have Been Changed or Removed,” on page C-1

Table B-1 lists the new configuration parameters in various versions of Dynamic Server. All parameters are located in the ONCONFIG file, unless otherwise noted. For more information on the configuration parameters, see the *IBM Informix Dynamic Server Administrator's Reference* and the *IBM Informix Dynamic Server Administrator's Guide*.

Table B-1. New Configuration Parameters

Version	New Configuration Parameter	Description
11.50xC2	LIMITNUMSESSIONS	Defines the maximum number of sessions that you want connected to Dynamic Server. If you specify a maximum number, you can also specify whether you want Dynamic Server to print messages to the online.log file when the number of sessions approaches the maximum number.
11.50	FAILOVER_CALLBACK	Specifies the full path name of a script that the database server executes when the server transitions from a secondary server to a primary or standard server.
11.50	HA_ALIAS	When a secondary server connects to a primary server, specifies the name of a network alias to use if a failover occurs.
11.50	MSG_DATE	When enabled, adds a date to the front of each message in the online log.
11.50	SHMNOACCESS	Specifies a virtual memory address range to not use to attach shared memory.
11.50	SSL_KEYSTORE_FILE	On clients, specifies the fully qualified file name of the keystore that stores the certificates of all servers to which the client connects.
11.50	SSL_KEYSTORE_LABEL	On Dynamic Server, specifies the label of the server digital certificate used in the keystore database that stores Secure Sockets Layer (SSL) keys and digital certificates.
11.50	SSL_KEYSTORE_STH	On clients, specifies the fully qualified file name of the stash file containing the encrypted keystore password.
11.50	STORAGE_FULL_ALARM	Configures the frequency and severity of messages and alarms when storage spaces become full.

Table B-1. New Configuration Parameters (continued)

Version	New Configuration Parameter	Description
11.50	UPDATABLE_SECONDARY	Enables client applications to perform update, insert, and delete operations on a high-availability secondary server.
11.10	ADMIN_MODE_USERS	Enables user informix or a DBSA to give one or more specific users the ability to connect to the database server in administration mode through the onmode -j command, the oninit -U command, or the ADMIN_MODE_USERS configuration parameter.
11.10	ADMIN_USER_MODE_WITH_DBSA	Specifies whether user informix and the DBSA group users can connect to the database server while it is in administration mode.
11.10	AUTO_AIOVPS	Enables or disables the ability of the database server to automatically increase the number of AIO VPS and flusher threads when the server detects that AIO VPs are not keeping up with the I/O workload.
11.10	AUTO_CKPTS	Enables or disables automatic checkpoints when the database server starts.
11.10	AUTO_LRU_TUNING	Enables or disables automatic LRU tuning when the database server starts.
11.10	AUTO_REPREPARE	Controls whether Dynamic Server automatically re-optimizes SPL routines and re-prepares prepared objects after the schema of a table referenced by the SPL routine or by the prepared object has been significantly changed.
11.10	BACKUP_FILTER	Specifies the path name of a backup filter program and any options that ON-Bar uses.
11.10	BAR_PERFORMANCE	Controls the level of information in the ON-Bar Activity log.
11.10	DIRECT_IO	Controls the use of direct I/O for cooked files used for database space chunks.
11.10	DRDA_COMMBUFFSIZE	Sets the buffer size of the DRDA communications buffer.
11.10	ENCRYPT_HDR	Enables or disables HDR encryption.
11.10	ENCRYPT_SMX	Sets the level of encryption for high-availability secondary server configurations.
11.10	EXPLAIN_STAT	Enables or disables the inclusion of a Query Statistics section in the explain.out file that the SET EXPLAIN statement of SQL or the onmode -Y session_id command can display.
11.10	LOG_INDEX_BUILDS	Enables or disables index page logging Index page logging is required when using RS secondary servers.

Table B-1. New Configuration Parameters (continued)

Version	New Configuration Parameter	Description
11.10	MAX_FILL_DATA_PAGES	Enables the database server to insert more rows per page into tables with variable-length rows.
11.10	PLCY_HASHSIZE	Specifies the number of hash buckets in the cache that holds information about label-based access control (LBAC) credentials for users
11.10	PLCY_POOLSIZE	Specifies the maximum number of entries in each hash bucket of the security policy information cache.
11.10	RESTORE_FILTER	Specifies the path name of a restore filter program and any options that ON-Bar uses.
11.10	RTO_SERVER_RESTART	Sets the amount of time, in seconds, that Dynamic Server has to recover from a problem after you restart Dynamic Server and bring the server into online or quiescent mode.
11.10	SDS_ENABLE,	Enables the shared-disk (SD) secondary server function.
11.10	SDS_PAGING	Specifies the location of two files that act as buffer-paging files.
11.10	SDS_TEMPDBS	Specifies information that the SD secondary server uses to dynamically create temporary dbspaces when the SD secondary server starts.
11.10	SDS_TIMEOUT	Specifies the amount of time in seconds that the primary server waits for the SD secondary server to send a log-position acknowledgment.
11.10	SHMVIRT_ALLOCSEG	Specifies a threshold at which Dynamic Server allocates server memory, and specifies the alarm level activated if the server cannot allocate the new memory segment.
11.10	SQLTRACE	Controls the default behavior, such as the number of SQL statements to trace and the tracing mode, of the Query Drill-Down feature.
11.10	USELASTCOMMITTED	Specifies whether the database server uses the last committed version of the data when a lock occurs.
11.10	USRC_HASHSIZE	Specifies the number of hash buckets in the cache that holds information about LBAC credentials for users.
11.10	USRC_POOLSIZE	Specifies the maximum number of entries in each hash bucket of the cache that holds information about LBAC credentials for users.
11.10	TEMPTAB_NOLOG	Disables logging on temporary tables.

Table B-1. New Configuration Parameters (continued)

Version	New Configuration Parameter	Description
10.00xC6	VP_MEMORY_CACHE_KB	Enables a private memory cache that is associated with a CPU virtual processor and contains blocks of free memory.
10.00xC5	BAR_IXBAR_PATH	Specifies the path and name of the ixbar , the ON-Bar boot file.
10.00xC5	FASTPOLL	Enables fast polling of your network, if your operating-system platform supports fast polling.
10.00xC5	IFX_FOLDVIEW	Enables views to be folded into a parent query.
10.00xC4	DB_LIBRARY_PATH	Specifies a comma-separated list of valid directory prefix locations from which the database server can load external modules.
10.00xC4	SECURITY_LOCALCONNECTION	Lets you verify security on local connections by verifying that the ID of the local user who is running a program is the same ID of the user who is trying to access the database.
10.0	ALRM_ALL_EVENTS	Specifies whether ALARMPROGRAM runs for all events that are logged in the MSGPATH or only specified noteworthy events.
10.0	BUFFERPOOL	Specifies configuration information for a buffer pool for each different page size used by a dbspace.
10.0	CDR_SUPPRESS_ATSRISWARN	Enterprise Replication configuration parameter that specifies whether comma-separated error and warning numbers are suppressed from ATS and RIS files.
10.0	DRIDXAUTO	Determines how a secondary database server reacts to a high-availability data-replication failure.
10.0	DS_NONPDQ_QUERY_MEM	Increases the amount of sort memory that is available for a query that is not a PDQ query.
10.0	EXT_DIRECTIVES	An external optimizer directive that provides a temporary solution to problems when you do not want to change SQL statements in queries
10.0	IFX_EXTEND_ROLE	Enables a database server administrator (DBSA) to prevent unauthorized users from registering DataBlade modules or external user-defined routines (UDRs).
10.0	LISTEN_TIMEOUT	Sets the incomplete connection timeout period.
10.0	MAX_INCOMPLETE_CONNECTIONS	Restricts the number of incomplete requests for connections.

Table B-1. New Configuration Parameters (continued)

Version	New Configuration Parameter	Description
10.0	ONLIDX_MAXMEM	Limits the amount of memory that is allocated to the <i>preimage</i> log pool and to the <i>updater</i> log pool in shared memory. You can use this configuration parameter if you plan to complete other operations on a table column while executing the CREATE INDEX ONLINE statement on the column.
10.0	TBLTBLFIRST	Specifies the first extent size of tablespace tblspace in kilobytes.
10.0	TBLTBLNEXT	Specifies the next extent size of tablespace tblspace in kilobytes.
9.40	CDR_DBSPACE	Defines the default dbspace for the Enterprise Replication syscdr database.
9.40	CDR_ENV	Sets Enterprise Replication environment variables CDR_LOGDELTA, CDR_PERFLOG, CDR_ROUTER, and CDR_RMSCALEFACT.
9.40	CDR_MAX_ DYNAMIC_LOGS	Specifies the number of dynamic log file requests that Enterprise Replication can make in one server session.
9.40	ENCRYPT_CDR	Enables and sets the level of network encryption for Enterprise Replication.
9.40	ENCRYPT_CIPHERS	Specifies the ciphers to use for encryption for Enterprise Replication.
9.40	ENCRYPT_MAC	Specifies the level of message authentication coding to use for Enterprise Replication.
9.40	ENCRYPT_MACFILE	Specifies MAC key files for Enterprise Replication.
9.40	ENCRYPT_SWITCH	Defines the frequency at which ciphers and secret keys are re-negotiated for Enterprise Replication.
9.40	HPL_DYNAMIC_LIB_PATH	For the High-Performance Loader, sets the location of the shared-library file containing custom-code functions. Located in the plconfig file.
9.40	HPLAPIVERSION	For the High-Performance Loader, sets whether custom-code functions can use different input and output data lengths. Located in the plconfig file.
9.40	PLOG_OVERFLOW_PATH	Sets the location of the temporary space to extend the physical log during fast recovery.

Appendix C. Configuration Parameters That Have Been Changed or Removed

Dynamic Server Version 11.50 includes an improved **onconfig.std** file, with new default values for configuration parameters. In addition, some configuration parameters used with earlier versions of Dynamic Server have been changed or removed.

Configuration Parameter Changes in the Version 11.50 **onconfig.std** File

In the new **onconfig.std** file, comments and the parameters are listed separately and are grouped by functional areas. Some configuration parameters that specify sizes have now have higher values. Some configuration parameters that specify file locations now have more secure default locations under the **\$INFORMIXDIR** directory.

Deprecated configuration parameters were removed from the file.

Configuration Parameters that Have Been Added to the **onconfig.std** File

The following table lists the configuration parameters that were added to the **onconfig.std** file in Version 11.50 and their values.

Table C-1. Configuration Parameters Added to the onconfig.std File

Configuration Parameter	Value
ADMIN_USER_MODE_WITH_DBSA	none
BTSCANNER	num=1,priority=low,threshold=5000, rangesize=-1,alice=6,compression=default
BACKUP_FILTER	none
BAR_DEBUG	0
CDR_SUPPRESS_ATSRISWARN	none
DD_HASHMAX	10
DD_HASHSIZE	31
DEF_TABLE_LOCKMODE	page
DS_HASHSIZE	31
DS_POOLSIZE	127
ENCRYPT_CDR	none
ENCRYPT_CIPHERS	none
ENCRYPT_HDR	none
ENCRYPT_MAC	none
ENCRYPT_MACFILE	none
ENCRYPT_SMX	none
ENCRYPT_SWITCH	none
EXT_DIRECTIVES	0

Table C-1. Configuration Parameters Added to the onconfig.std File (continued)

Configuration Parameter	Value
FAILOVER_CALLBACK	none
FASTPOLL	1
HA_ALIAS	none
LOG_INDEX_BUILDS	none
MAX_INCOMPLETE_CONNECTIONS	1024
PC_HASHSIZE	31
PC_POOLSIZE	127
PLCY_HASHSIZE	127
PLCY_POOLSIZE	31
PLOG_OVERFLOW_PATH	UNIX: \$INFORMIXDIR/tmp Windows: none
REDIRECTED_WRITES	0
RESTORE_FILTER	none
S BSPACETEMP	none
SDS_ENABLE	none
SDS_PAGING	none
SDS_TEMPDBS	none
SDS_TIMEOUT	20
SECURITY_LOCALCONNECTION	none
SQLTRACE	Commented out: # SQLTRACE level=low,ntraces=1000,size=2,mode=global
SSL_KEYSTORE_LABEL	none
STMT_CACHE	0
STMT_CACHE_HITS	0
STMT_CACHE_NOLIMIT	0
STMT_CACHE_NUMPOOL	1
STMT_CACHE_SIZE	512
UNSECURE_ONSTAT	none
USRC_HASHSIZE	31
USRC_POOLSIZE	127
VPCLASS	cpu,num=1,noage Commented out: # VPCLASS aio,num=1 Commented out: #VPCLASS jvp,num=1

Configuration Parameters that Have New Default Values

The following table lists the configuration parameters that have new default values in the **onconfig.std** file.

Table C-2. Configuration Parameters with New Default Values in the onconfig.std File

Configuration Parameter	Previous Value	New Value
ADMIN_MODE_USERS	1	None
ALARMPROGRAM	UNIX: /usr/informix/etc/ alarmprogram.sh Windows: None	UNIX: \$INFORMIXDIR/etc/ alarmprogram.sh Windows: \$INFORMIXDIR\etc\ alarmprogram.bat
BAR_ACT_LOG	/usr/informix/bar_act.log	UNIX: \$INFORMIXDIR/ tmp/bar_act.log Windows: \$INFORMIXDIR\tmp\ bar_act.log
BAR_BSALIB_PATH	UNIX: \$INFORMIXDIR/lib/ libsad001.so Window: libbsa.dll	None
BAR_DEBUG_LOG	UNIX: /usr/informix/ bar_dbug.log Windows: bar_dbug.log	UNIX: \$INFORMIXDIR/ tmp/bar_dbug.log Windows: \$INFORMIXDIR\tmp\ bar_dbug.log
BUFFERPOOL	Operating systems with 2K page size: default,buffers=5000,lrus=8, lru_min_dirty=50, lru_max_dirty=60 size=2k,buffers=5000,lrus=8, lru_min_dirty=50, lru_max_dirty=60 Operating systems with 4K page size: default,buffers=1000,lrus=8, lru_min_dirty=50, lru_max_dirty=60 size=4k,buffers=1000,lrus=8, lru_min_dirty=50, lru_max_dirty=60	Operating systems with 2K page size: default,buffers=10000,lrus=8, lru_min_dirty=50.00, lru_max_dirty=60.50 size=2k,buffers=50000,lrus=8, lru_min_dirty=50, lru_max_dirty=60 Operating systems with 4K page size: default,buffers=10000,lrus=8, lru_min_dirty=50.00, lru_max_dirty=60.50 size=4k,buffers=10000,lrus=8, lru_min_dirty=50, lru_max_dirty=60
CLEANERS	1	8
CONSOLE	UNIX: /dev/console Windows: console.log	UNIX: \$INFORMIXDIR/ tmp/online.con Windows: online.con
DB_LIBRARY_PATH	commented out: # DB_LIBRARY_PATH \$INFORMIXDIR/extend	commented out: # DB_LIBRARY_PATH

Table C-2. Configuration Parameters with New Default Values in the onconfig.std File (continued)

Configuration Parameter	Previous Value	New Value
DRLOSTFOUND	UNIX: /usr/etc/dr.lostfound Windows: \tmp	UNIX: \$INFORMIXDIR/etc/dr.lostfound Windows: \$INFORMIXDIR\tmp
DUMPDIR	UNIX: /usr/informix/tmp Windows: INFORMIXDIR\tmp	UNIX: \$INFORMIXDIR/tmp Windows: \$INFORMIXDIR\tmp
EXPLAIN_STAT	0	1
LISTEN_TIMEOUT	10	60
LOCKS	2000	20000
LOGBUFF	32	64
LOGSIZE	2000	10000
LTAPEDEV	UNIX: /dev/tapedev Windows: \\.\TAPE1	UNIX: /dev/tapedev (same as previous value) Windows: NUL
MIRRORPATH	None	UNIX: \$INFORMIXDIR/tmp/demo_on.root_mirror Windows: none
MSGPATH	UNIX: /usr/informix/online.log Windows: online.log	UNIX: \$INFORMIXDIR/tmp/online.log Windows: online.log
NETTYPE	UNIX: none Windows: onsoctcp,drssoctcp,1,NET	UNIX: ipcshm,1,50,CPU Windows: none
PHYSBUFF	32	128
PHYSFILE	2000	50000
RA_PAGES	None	64
RA_THRESHOLD	None	16
ROOTPATH	UNIX: /dev/online_root Windows: None	UNIX: \$INFORMIXDIR/tmp/demo_on.rootdbs Windows: None
ROOTSIZE	30000	200000
SHMVIRT_ALLOCSEG	0	0,3
SHMVIRT_SIZE	8192	32656
SYSALARMPROGRAM	UNIX: /usr/informix/etc/evidence.sh Windows: INFORMIXIDR\etc\ evidence.bat	UNIX: \$INFORMIXDIR/etc/evidence.sh Windows: Commented out: # SYSALARMPROGRAM \$INFORMIXDIR\etc\ evidence.bat

Table C-2. Configuration Parameters with New Default Values in the onconfig.std File (continued)

Configuration Parameter	Previous Value	New Value
TAPEBLK	32	UNIX: 32 Windows: 16
TAPESIZE	10240	0

Configuration Parameters that Have Been Changed or Removed

The following table contains a list of other configuration parameters that have been changed or removed in Version 11.50.

Table C-3. Configuration Parameters that Have Been Changed or Removed

11.50	AFF_SPROC	Removed (AFF_NPROCS was previously removed.)
11.50	DUMPSHMEM	Has new options for controlling how much memory is written to a dump file.
11.50	JDKVERSION	Removed
11.50	JVPJAVAHOME	/usr/informix in the directory name of the configuration parameter is replaced with \$INFORMIXDIR . The value is now: \$INFORMIXDIR/extend/krakatoa/jre
11.50	JVPHOME	/usr/informix in the directory name of the configuration parameter is replaced with \$INFORMIXDIR . The value is now: \$INFORMIXDIR/extend/krakatoa
11.50	JVPPROFILE	/usr/informix in the directory name of the configuration parameter is replaced with \$INFORMIXDIR . The value is now: \$INFORMIXDIR/extend/krakatoa/ .jvpprops
11.50	JVPLOGFILE	/usr/informix in the directory name of the configuration parameter is replaced with \$INFORMIXDIR . The value is now: \$INFORMIXDIR/jvp.log
11.50	JVPCCLASSPATH	/usr/informix in the directory name of the configuration parameter is replaced with \$INFORMIXDIR . The value is now: \$INFORMIXDIR/extend/krakatoa/ krakatoa.jar:\$INFORMIXDIR/extend/ krakatoa/jdbc.jar
11.50	NOAGE	Removed
11.50	NUMCPUVPS	Removed
11.50	PHYSDBS	Removed

Configuration Parameters that Have Been Changed or Removed In Versions 9.30 through 11.10

The following table contains a list of configuration parameters that were changed or removed in prior versions of the server.

Table C-4. Configuration Parameters that Were Changed or Removed in Version 11.10, 10.0, 9.40, and 9.30

Version	Changed or Removed Configuration Parameter	Description of Change
11.10	FAST_RESTART_CKPT_FUZZYLOG	Removed. The RTO_SERVER_RESTART configuration parameter eliminates fuzzy checkpoints, using interval checkpoints instead.
11.10	FAST_RESTART_PHYSLOG	Removed.
11.10	NOFUZZYCKPT	Removed. The RTO_SERVER_RESTART configuration parameter eliminates fuzzy checkpoints, using interval checkpoints instead.
10.00xc6	SINGLE_USER_MODE_WITH_DBSA (renamed to ADMIN_USER_MODE_WITH_DBSA in Version 11.10)	Renamed. In Version 11.10, the name of this configuration parameter changed to ADMIN_USER_MODE_WITH_DBSA.
10.0	BUFFERS	Removed. Information now specified with the BUFFERPOOL configuration parameter.
10.0	LRUS	Removed. Information now specified with the BUFFERPOOL configuration parameter.
10.0	LRU_MAX_DIRTY	Removed. Information now specified with the BUFFERPOOL configuration parameter.
10.0	LRU_MIN_DIRTY	Removed. Information now specified with the BUFFERPOOL configuration parameter.
9.40	ALARMPROGRAM	Can be set to the alarmprogram.sh file to enable event alarms.
9.40	CDR_QDATA_SBSpace	Can accept up to 32 sbspaces.
9.40	CDR_QDATA_SBFLAGS	Removed. Enterprise Replication always uses the default log mode of the sbspace for spooling row data.
9.40	DBSERVERALIASES	Can accept up to 32 server alias values.
9.40	LTAPEBLK	New default value.
9.40	LTAPESIZE	Can accept a value of 0 to read or write to the end of the tape device.
9.40	LRU_MAX_DIRTY	Can accept a value of type INTEGER or FLOAT. (This configuration parameter was removed in Version 10.0.)

Table C-4. Configuration Parameters that Were Changed or Removed in Version 11.10, 10.0, 9.40, and 9.30 (continued)

Version	Changed or Removed Configuration Parameter	Description of Change
9.40	LRU_MIN_DIRTY	Can accept a value of type INTEGER or FLOAT. (This configuration parameter was removed in Version 10.0.)
9.40	OPTICAL_LIB_PATH	Is valid for both UNIX and Windows. Must be set to the location of the storage manager library.
9.40	TAPEBLK	New default value.
9.40	TAPESIZE	Can accept a value of 0 to read or write to the end of the tape device.
9.30	AFF_NPROCS	Removed; superseded by the VPCLASS configuration parameter.
9.30	AFF_SPROC	Superseded by the VPCLASS configuration parameter.
9.30	CDR_LOGBUFFERS	Removed.
9.30	CDR_LOGDELTA	Removed.
9.30	CDR_NIFRETRY	Removed.
9.30	CDR_NUMCONNECT	Removed.
9.30	JVPJAVAHOME	New default location for the JRE.
9.30	JVPJAVALIB	New default value that is platform dependent.
9.30	JVPJAVAVM	New default value that is platform dependent.
9.30	LBU_PRESERVE	Removed; configured an obsolete utility.
9.30	LOGSMAX	Removed.
9.30	NOAGE	Superseded by the VPCLASS configuration parameter.
9.30	NUMAIOVPS	Removed; superseded by the VPCLASS configuration parameter.
9.30	NUMCPUVPS	Superseded by the VPCLASS configuration parameter.

Appendix D. New Reserved Words

Each version of Dynamic Server supports new SQL keywords that are reserved words and might affect migration of your applications.

Although you can use almost any word as an SQL identifier, syntactic ambiguities might occur if you use an SQL reserved word. An ambiguous statement might not produce the results you want.

The following table shows a list of new SQL reserved words. For a complete list of SQL reserved words in IBM Informix Dynamic Server, see the *IBM Informix Guide to SQL: Syntax*.

Table D-1. New SQL Reserved Words

Dynamic Server Version	Reserved Words
Version 11.50xC3	SAVEPOINT
Version 11.50	BIGINT BIGSERIAL VERCOLS
Version 11.10	ADMIN AVOID_INDEX_SJ FINAL IDSSECURITYLABEL INDEX_SJ INSERTING REFERENCES SAMPLING SELECTING STATEMENT SYSDBCLOSE SYSDBOPEN TASK UPDATING USELASTCOMMITTED WITH In addition, the DBSECADM role is reserved for LBAC administrative work. Version 11.10 contains a new database, the sysadmin database. If your source database server contains a database named sysadmin , you must rename it.

Table D-1. New SQL Reserved Words (continued)

Dynamic Server Version	Reserved Words
Version 10.0	ACTIVE CURRENT_ROLE DEFAULT_ROLE DIRECTIVES ENCRYPTION HINT IGNORE INACTIVE INITCAP INLINE INOUT LIMIT LOAD ONLINE OPTCOMPIND PARTITION PASSWORD REUSE SAVE SKIP STORAGE TEMPLATE TEST TRUNCATE TYPEID TYPENAME TYPEOF UNLOAD XADATASOURCE XID
Version 9.40	COLLATION CROSS FULL INSTEAD RESTART RIGHT
Version 9.30	AVOID_EXECUTE AVOID_SUBQF USE_SUBQF
Version 9.21	AVOID_HASH AVOID_INDEX AVOID_NL RAW STANDARD USE_HASH USE_NL
Version 7.31	INNER JOIN LEFT LOCKS RETAIN

Appendix E. System Catalog and System Database Changes

New versions of Dynamic Server contain system catalog table changes and **sysmaster** database changes.

Changes for Dynamic Server 11.50

Dynamic Server Version 11.50 contains two new **sysmaster** database tables and new SMI tables.

The version 11.50xC3 **sysmaster** database contains this new table for the Change Data Capture API:

syscdc

The version 11.50 **sysmaster** database contains this new table:

sysstesappinfo

The following new SMI tables contain information about Enterprise Replication that you can use to monitor status and diagnose problems:

- The **syscdr_state** table contains information on whether Enterprise Replication, data capture, data apply, and the network between servers is active.
- The **syscdr_ddr** table contains information about the status of log capture and the proximity or status of transaction blocking (DDRBLOCK) or transaction spooling.
- The **syscdr_nif** table contains information about network connections and the flow of data between Enterprise Replication servers.
- The **syscdr_rcv** table contains information about transactions being applied on target servers and acknowledgements being sent from target servers.
- The **syscdr_atmdir** table contains information about the contents of the ATS directory.
- The **syscdr_risdir** table contains information about the contents of the RIS directory.
- The **syscdr_ats** table contains the first ten lines of content of each ATS file.
- The **syscdr_ris** table contains the first ten lines of content of each RIS file.
- The **syscdr_rqstamp** table contains information about which transaction is being added into each queue.
- The **syscdr_rqmhandle** table contains information about which transaction is being processed in each queue.

Changes for Dynamic Server 11.10

Dynamic Server Version 11.10 contains new **sysmaster** database tables and the new **sysadmin** database.

The version 11.10 **sysmaster** database contains these new tables:

syscheckpoint sysckptinfo

Version 11.10 includes a new database, **sysadmin**, which contains tables that store task properties. This database is dropped when you revert to earlier versions of Dynamic Server. If your source database server contains a **sysadmin** database, you must rename it.

Changes for Dynamic Server 10.0

Dynamic Server Version 10.0 contains new **sysmaster** database tables and new SMI tables.

The version 10.0 **sysmaster** database contains these new tables:

sysdirectives sysbufpool, a system-monitoring interface (SMI)

The following changes were made to other SMI tables:

- The **sysfragments** table contains a **Partition** column and the **Flags** column now tells you if the fragmentation scheme has partitions.
- The **sysusers** table contains a **defrole** column.
- The **sysams** table contains an **am_truncate** column.
- The **sysprocedures** table contains a **rtnparameters** column for information on INOUT parameters.
- The **syspaghdr** table has a **pg_pagesize** column.
- The **sysptnhdr** table has a **pagesize** column.
- The **sysptnhdr** table has a **bpoolindx** column that indicates which buffer pool the buffer is in.
- The **sysbufhdr** table has a **bufsize** column, which indicates the buffer page size.
- The **sysdbstab** and **syschktab** tables have **pagesize** columns.
- The views **syschunks** and **sysdbspaces** tables have a **pagesize** columns.
- The views **systabinfo** table has a **ti_pagesize** column.
- The views **systabpaghdrs** and **sysphyspaghdrs** tables have **pg_pagesize** columns.

In addition, tables added to the **syscdr** database are removed.

Changes for Dynamic Server 9.40

Dynamic Server Version 9.40 contains a new **sysmaster** database table. In addition, some tables contain a new **collation** column.

The following new system catalog table was added:

syssequences

A new **collation** column has been added to the following system catalog tables:

sysconstraints sysindices sysprocplan sysrigbody

Changes for Dynamic Server 9.30

Some **sysmaster** database tables were removed from Dynamic Server Version 9.30.

The following tables were deleted from the **sysmaster** database in Version 9.30:

arc_ae_view	arc_pendreq_view	arc_server arc_version
arc_db_file_view	arc_phys_dev arc_rep_table	arc_vol_lock_view
arc_dbspace arc_dbspace_set	arc_replicate	arc_volume_view arc_vset
arc_directory_view	arc_req_vset_view	arc_vset_user_view
arc_file_view	arc_request_view	arc_vset_view
arc_file_copy_view	arc_save_set_view	

Column-Width Changes in **sysmaster** Tables in Version 9.20 and Later Versions

Version 9.20 and later versions of Dynamic Server provide long identifiers. All identifiers in the system catalog tables and the **sysmaster** database reflect these new limits on identifier length.

The *IBM Informix Guide to SQL: Syntax* defines *identifiers*, which specify the names of database objects.

The column widths for identifiers that refer to database objects and other identifiers changed from CHAR(18) to VARCHAR(128,0) in the following system catalog tables:

sysaggregates sysams	sysfragauth sysfragments	sysroutinelangs
sysattrtypes sysblobs	sysindexes sysindices	sys synonyms sysstable
syscasts syscolattribs	sysobjstate sysopclasses	sysstabamdata systables
syscolumns sysconstraints	sysopclstr sysprocedures	sysracemsgs sys triggers
sysdomains		sysxdtypes

Identifiers changed from CHAR(18) to CHAR(128) in the following **sysmaster** database tables:

arc_dbspace arc_dbspace_set	syscrtadt sysdatabases	syslocks sysopendb sysprc
arc_phys_dev arc_rep_table	sysdbslocale sysdbspaces	sysproccache sysptprof
arc_replicate arc_server	sysdbspartn sysdbstab	sysssdblock sysssqlcurall
arc_version arc_vset	sysdic sysdiccache	sysssqlcurses sysssqlstat
arc_vset_view flags_text	sysdistcache sysdsc	sysstabnames systrans
syscftab sysconfig	sysextents sysextspaces	sysxtptab

Column widths for user login identifiers changed from CHAR(8) to CHAR(32) in some system catalog tables and **sysmaster** database tables. The following system catalog tables changed:

sysaggregates sysams	sysindices syslangauth	sys synonyms sysstable
syscasts syscolauth	sysobjstate sysopclasses	sysstabauth systables
sysconstraints sysdomains	sysopclstr sysprocauth	sys triggers sysusers
sysfragauth sysindexes	sysprocedures sysroleauth	sysxdtypeauth sysxdtypes

The following **sysmaster** database tables changed:

sysaudit sysdatabases	sysdiccache sysdistcache	sysrstcb sysyclst
sysdbspaces sysdbspartn	sysdsc sysextspaces sysprc	sysessions systabnames
sysdbstab sysdic	sysproccache	sysuserthreads

Columns that include path names or other values changed from CHAR(128) to CHAR(256) in the following **sysmaster** database tables:

sysadtinfo syschktab	syscrtadt sysdrbc sysdri	sysmchktab
syschunks		

The path for a physical device changed from CHAR(128) to CHAR(260) in the following **sysmaster** database table:

arc_phys_dev

Column widths changed from CHAR(20) to CHAR(128) for longer object names in the following **sysmaster** database tables:

sysdrbc sysdri

Column widths changed from CHAR(37) to CHAR(257) in the following **sysmaster** database tables:

sysdistcache sysdsc	sysprc sysproccache
----------------------------	----------------------------

Column widths changed from DECIMAL(16,0) to DECIMAL(32,0) in the following **sysmaster** database table:

sysesprof

The **tabauth** column of the **sysstabauth** system catalog table is now CHAR(9) instead of CHAR(8). The 9th character indicates the Under privilege.

Data Type Changes in **sysmaster** tables in Version 9.20 and Later Versions

The data type of some **sysmaster** database tables was changed in Dynamic Server Version 9.30.

The “Column-Width Changes in **sysmaster** Tables in Version 9.20 and Later Versions” on page E-3 topic lists columns that have changed from the CHAR data type to the VARCHAR data type. In addition, one or more columns changed from the SMALLINT data type to the integer data type in the following **sysmaster** database tables:

sysdbspaces sysdbstab	sysdic sysrstcb	sysdbblock
------------------------------	------------------------	-------------------

The CHAR data type changed to the STAT data type in the following system catalog table:

sysdistrib

Changes in Treatment of Null Values in sysmaster Tables in Version 9.30

Starting with Version 9.30, nulls are allowed in some **sysmaster** database tables, in which nulls were previously not allowed.

Nulls are allowed for some columns in the following **sysmaster** database tables:

arc_ae_view	arc_pendreq_view	arc_volume_view
arc_db_file_view	arc_req_vset_view	arc_vset_user_view
arc_directory_view	arc_request_view	arc_vset_view
arc_file_copy_view	arc_save_set_view	
arc_file_view	arc_vol_lock_view	

Other sysmaster Database Table and Column Changes in Version 9.30

Dynamic Server Version 9.30 contains a new **sysmaster** database table. In addition, some tables were removed and new columns were added to a few tables..

The following tables have been added to the **sysmaster** database:

logmessage	syscdrack_buf	syscdrctrl_txn	syscdrprog	syscdrrecv_txn	syscdrtx
syscdrack_txn	syscdrctrl_buf	syscdrq	syscdrrecv_buf		

The following **arc_change_log** table has been deleted from the **sysmaster** database:

One or more columns have been added to the following system catalog tables:

sysams	sysdistrib	sysprocedures	sysroutinelangs
--------	------------	---------------	-----------------

Several columns have been added to the **sysdbstab** system catalog table.

Remote Queries on System Catalog Tables Between Version 7.31 and Later Versions

Certain system catalog tables use data types that are not supported in Dynamic Server Version 7.31. Remote queries that issue a **SELECT *** statement on these system catalog tables from Version 7.31 to later versions will fail.

For example, the following queries that originate on Version 7.31 fail if you try to run them on a later version of the server:

```
SELECT * FROM dbname@remoteserver:sysindices;  
SELECT * FROM dbname@remoteserver:sysindexes;
```

Instead of using an asterisk as the Projection clause, specify the required column names explicitly. You cannot specify any columns that have user-defined types.

Difference in sysindexes between Version 7.31 and Later Versions

In Version 7.31, **sysindexes** is a table. In Dynamic Server Version 11.50, 11.10, 10.0, 9.40, 9.30, and 9.21, **sysindexes** is a view.

The ALTER TABLE statement fails for **sysindexes** because this statement is not valid for altering a view.

Appendix F. New and Changed Features

Each version of Dynamic Server contains many new and changed features.

For descriptions of these changes, see the *IBM Informix Dynamic Server Getting Started Guide*.

For information on new and changed features that affect migration, see “What’s New in Migration for Dynamic Server, Version 11.50” on page x.

If you are migrating from a version of Dynamic Server that is earlier than Version 11.10, you need to know that:

- ADMIN_USER_MODE_WITH_DBSA is the new name for the SINGLE_USER_MODE_WITH_DBSA configuration parameter
- Dynamic Server supports including \$INFORMIXDIR as the first path name value in path name specifications in the ONCONFIG file

Server Library Name Changes in Version 11.50

Dynamic Server Version 11.50 contains some new server library names.

The following database server library names have new names. These library files have a **.so** or **.dll** extension.

Library	Name for 9.14, 9.21, 9.30, or 9.40 Server	Name for 10.0 Server	Name for 11.10 or 11.50 Server
Optical	iosm09a	iosm10a	iosm11a
pload	ipldd09a	ipldd10a	ipldd11a
Simple password CSM	ispws09a	ispws10a	ispws11a
Encryption CSM	iencs09a	iencs10a	iencs11a

Appendix G. 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.

Licensees 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

A

- accessibility G-1
 - keyboard G-1
 - shortcut keys G-1
- Accessibility G-1
 - dotted decimal format of syntax diagrams G-1
 - syntax diagrams, reading in a screen reader G-1
- ADMIN_MODE_USERS configuration parameter B-2
- ADMIN_USER_MODE_WITH_DBSA configuration parameter B-2
- ALARMPROGRAM configuration parameter 5-4, 5-5
- ALL keyword
 - dbschema utility 10-2
- ALRM_ALL_EVENTS configuration parameter B-4
- ANSI joins 6-3
- AUTO_AIOVPS configuration parameter B-2
- AUTO_CKPTS configuration parameter B-2
- AUTO_LRU_TUNING configuration parameter B-2
- AUTO_REPREPARE configuration parameter B-2

B

- BACKUP_FILTER configuration parameter B-2
- Backups
 - after upgrading 5-9
 - before reverting 6-9
 - before upgrading to a new version 3-7
 - logical logs 5-5
 - ON-Bar utility 3-7, 6-13
 - ontape utility 3-7, 6-13
 - source database 3-7
- BAR_IXBAR_PATH configuration parameter B-4
- BAR_PERFORMANCE configuration parameter B-2
- Binary files, loading 13-1, 13-2
- BladeManager
 - installing and registering DataBlade modules 5-5, 5-9
 - removing new extensions before reversion 6-10
- Blobspaces
 - moving, with onunload and onload 13-9, 13-10
- boot90.sql and boot901.sql 6-1
- BUFFERPOOL configuration parameter B-4

C

- CDR_DBSPACE configuration parameter B-5
- CDR_ENV configuration parameter B-5
- CDR_MAX_DYNAMIC_LOGS configuration parameter B-5
- CDR_SUPPRESS_ATSRISWARN configuration parameter B-4
- Character-position form of FILE and INSERT statements 9-9
- Checking available space
 - before migration 3-3
- Checking database integrity 5-8
- Chunks
 - reverting reserve pages 6-8
- Closing
 - Dynamic Server 3-5, 3-8
- Column-level encryption 6-3
- Command file
 - dbload 9-5

- Communications Support Module
 - configuring after migration 5-5
 - removing if reverting 6-12
 - saving before reverting 6-9
- concsn.cfg file 5-5
 - creating entries after migration 5-5
 - removing if reverting 6-12
 - saving before reverting 6-9
- Configuration file
 - customizing after migration 5-4
 - replacing after reversion 6-12
 - saving before migration 3-4
 - saving before reverting 6-9
- Configuration parameters
 - added in Dynamic Server 10.0 B-1
 - added in Dynamic Server 11.10 B-1
 - added in Dynamic Server 9.40 B-1
 - ADMIN_MODE_USERS B-2
 - ADMIN_USER_MODE_WITH_DBSA B-2
 - ALARMPROGRAM 5-4, 5-5
 - ALRM_ALL_EVENTS B-4
 - AUTO_AIOVPS B-2
 - AUTO_CKPTS B-2
 - AUTO_LRU_TUNING B-2
 - AUTO_REPREPARE B-2
 - BACKUP_FILTER B-2
 - BAR_IXBAR_PATH B-4
 - BAR_PERFORMANCE B-2
 - BUFFERPOOL B-4
 - CDR_DBSPACE B-5
 - CDR_ENV B-5
 - CDR_MAX_DYNAMIC_LOGS B-5
 - CDR_SUPPRESS_ATSRISWARN B-4
 - changed in Dynamic Server 10.0 C-6
 - changed in Dynamic Server 11.10 C-6
 - changed in Dynamic Server 11.50 C-1
 - changed in Dynamic Server 9.30 C-6
 - changed in Dynamic Server 9.40 C-6
 - changed in the onconfig.std file C-1
 - DB_LIBRARY_PATH B-4
 - DIRECT_IO B-2
 - DRDA_COMMBUFFSIZE B-2
 - DRIDXAUTO B-4
 - DS_NONPDQ_QUERY_MEM B-4
 - ENCRYPT_CDR B-5
 - ENCRYPT_CIPHERS B-5
 - ENCRYPT_HDR B-2
 - ENCRYPT_MAC B-5
 - ENCRYPT_MACFILE B-5
 - ENCRYPT_SMX B-2
 - ENCRYPT_SWITCH B-5
 - EXPLAIN_STAT B-2
 - EXT_DIRECTIVES B-4
 - FAILOVER_CALLBACK B-1
 - FASTPOLL B-4
 - HA_ALIAS B-1
 - HPL_DYNAMIC_LIB_PATH 5-4, B-5
 - IFX_EXTEND_ROLE B-4
 - IFX_FOLDVIEW B-4
 - IFX_LISTEN_TIMEOUT B-4

Configuration parameters (continued)

LIMITNUMSESSIONS B-1
LOG_INDEX_BUILDS B-2
MAX_FILL_DATA_PAGES B-3
MAX_INCOMPLETE_CONNECTIONS B-4
MSG_DATE B-1
ONLIDX_MAXMEM B-5
OPTICAL_LIB_PATH 5-4
PLCY_HASHSIZE B-3
PLCY_POOLSIZE B-3
PLOG_OVERFLOW_PATH B-5
removed in Dynamic Server 10.0 C-6
removed in Dynamic Server 11.10 C-6
removed in Dynamic Server 11.50 C-1
removed in Dynamic Server 9.30 C-6
removed in Dynamic Server 9.40 C-6
RESTORE_FILTER B-3
ROOTOFFSET 5-4
ROOTPATH 5-4
ROOTSIZE 5-4
RTO_SERVER_RESTART B-3
SDS_ENABLE B-3
SDS_PAGING B-3
SDS_TEMPDBS B-3
SDS_TIMEOUT B-3
SECURITY_LOCALCONNECTION B-4
SHMNOACCESS B-1
SHMVIRT_ALLOCSEG B-3
SQLTRACE B-3
SSL_KEYSTORE_LABEL B-1
SSL_KEYSTORE_FILE B-1
SSL_KEYSTORE_STH B-1
STORAGE_FULL_ALARM B-1
TBLTBLFIRST B-5
TBLTBLNEXT B-5
TEMPTAB_NOLOG B-3
UPDATABLE_SECONDARY B-2
USELASTCOMMITTED B-3
USRC_HASHSIZE B-3
USRC_POOLSIZE B-3
VP_MEMORY_CACHE_KB B-4
conpload.sh script 5-6
conploadlegacy.sh script 5-6
Converting to Dynamic Server Express Edition 4-4

D

Data integrity 5-8
Data migration
 automatic 2-1
 Code-set conversion
 HPL 2-4
 constraints 2-1
 High-Performance Loader
 ipload utility 2-4
 loading data 2-4
 onpladm utility 2-4
 ipload utility 2-4
 issues to consider 2-1
 Loading
 data 2-4
 onpladm utility 2-4
 overview 2-1
 prerequisites 2-1
 tools 2-1, 2-4, 2-5
 Utilities
 onpladm 2-4

Data migration (continued)

 Utilities (continued)
 onpload 2-4
Data types
 user defined
 converting replication of from Dynamic Server 9.21 4-2
Database servers
 migrating 5-1, 5-6
 platforms and versions 1-5
 preparing for migration 3-1
 upgrading 5-1, 5-6
Databases
 ownership, set by onload 13-10
 verifying integrity 5-8
DataBlade modules
 installing 5-5
 registering 5-9
DB_LIBRARY_PATH configuration parameter B-4
DB-Access
 input from the dbschema utility 10-12
dbexport
 SELECT triggers, disabling 8-2
dbexport utility 8-1
 -c option 8-2, 8-4
 -d option 8-2
 -q option 8-3
 -si option 8-3, 8-4, 10-3
 -ss option 8-3, 8-4
 -V option 8-3
 -version option 8-3
 -X option 8-3
 defined 8-1, 8-2
 destination options 8-4
 Interrupt key 8-4
 schema output 8-5
 syntax 8-2
dbimport utility 8-1
 -c option 8-7, 8-8
 -l option 8-11
 -q option 8-7
 -V option 8-7
 -version option 8-7
 -X option 8-7
 create options 8-10
 database logging mode 8-11
 defined 8-1, 8-6
 importing from another computer 2-6
 input file location options 8-9
 Interrupt key 8-8
 locale, changing 8-12
 renaming a database 8-11
 syntax 8-7
 using with GLS 8-8
 using with NLS 8-12
dbload utility
 -c command file option 9-2
 -d database option 9-2
 -e errors option 9-2
 -e option 9-4
 -i ignore rows option 9-2
 -i option 9-4
 -k option 9-2
 -l error log file option 9-2
 -p option 9-2
 -r option 9-2, 9-3
 -s option 9-2
 -V option 9-2

- dbload utility (*continued*)
 - version option 9-2
 - X option 9-2
 - compared to LOAD 9-1
 - creating a command file 9-5
 - dbload utility
 - n commit interval option 9-2
 - FILE statement 9-5
 - guidelines for handling objects 9-4
 - ignoring rows 9-4
 - importing from another computer 2-6
 - INSERT statements 9-5
 - compared to SQL INSERT statement 9-11
 - using 9-6
 - Interrupt key 9-4
 - number errors to allow 9-4
 - overview 9-1
 - speed, increasing 9-4
 - syntax 9-2
 - table locking 9-3
 - writing a command file
 - in character-position form 9-11
 - in delimiter form 9-8
- dbschema utility
 - ss option 10-5
 - u option 10-5
 - create schema across a network 10-4
 - create schema for a database 10-3
 - distribution information 10-10
 - example of file for DB-Access 10-12
 - guidelines 10-2
 - output example 10-10
 - overview 10-1
 - owner conventions 10-4
 - privileges information 10-7
 - privileges information for a role 10-7
 - re-creating the schema 10-13
 - sequence schema 10-6
 - specifying a table, view, or procedure 10-8
 - synonym schema 10-6
 - syntax 10-2
 - syntax for role schema 10-9
 - table information 10-9
- dbspaces
 - moving tables to another dbspace 13-10
- Delimiter form of FILE and INSERT statements 9-6, 9-8
- Diagnostic information to gather 3-8
- DIRECT_IO configuration parameter B-2
- Directories
 - installation 5-2
- Disabilities, visual
 - reading syntax diagrams G-1
- disability G-1
- Disabling
 - High-Availability Data Replication 3-7
- Distributed queries
 - with ANSI joins 6-3
- Dotted decimal format of syntax diagrams G-1
- DRDA_COMMBUFFSIZE configuration parameter B-2
- DRIDXAUTO configuration parameter B-4
- DS_NONPDQ_QUERY_MEM configuration parameter B-4
- Dynamic Server
 - backing up after upgrading 5-9
 - backing up before reverting 6-9
 - closing before migration 3-8
 - DataBlade modules, installing 5-5
 - DataBlade modules, registering 5-9

- Dynamic Server (*continued*)
 - initializing after upgrading 5-5
 - installing 5-2
 - migration
 - space requirements 3-3
 - new version
 - performance tuning 5-9
 - platforms and versions 1-5
 - reverting 6-1
 - reverting from current version 6-1, 6-10
 - starting after upgrading 5-5

E

- ENCRYPT_CDR configuration parameter B-5
- ENCRYPT_CIPHERS configuration parameter B-5
- ENCRYPT_HDR configuration parameter B-2
- ENCRYPT_MAC configuration parameter B-5
- ENCRYPT_MACFILE configuration parameter B-5
- ENCRYPT_SMX configuration parameter B-2
- ENCRYPT_SWITCH configuration parameter B-5
- Enterprise Edition
 - migrating to 1-1
- Enterprise Gateway
 - using to import data 2-6
- Enterprise Gateway Manager 2-6
- Enterprise Replication 4-4
- Environment variables
 - DB_LOCALE 8-12
 - DBTEMP 8-12
 - Environment variables
 - new in Dynamic Server 9.20 A-1
 - IFX_ONPLOAD_AUTO_UPGRADE 5-6
 - INFORMIXSERVER 5-4
 - INFORMIXSQLHOSTS 5-4
 - new in Dynamic Server 10.0 A-1
 - new in Dynamic Server 11.10 A-1
 - new in Dynamic Server 11.50 A-1
 - new in Dynamic Server 9.21 A-1
 - new in Dynamic Server 9.30 A-1
 - new in Dynamic Server 9.40 A-1
 - ONCONFIG 5-4
 - PATH 5-4
 - resetting after reversion 6-12
 - TEMP 8-12
 - TMP 8-12
- EXPLAIN STAT configuration parameter B-2
- Express Edition
 - does not support Enterprise Replication 4-4
 - migrating to 1-1
- EXT_DIRECTIVES configuration parameter B-4
- Extents, checking 5-8
- Extracting schema information 7-1

F

- FAILOVER_CALLBACK configuration parameter B-1
- Fast recovery
 - initiating 3-6
- FASTPOLL configuration parameter B-4
- FILE statement
 - character-position form 9-9
 - delimiter form 9-6, 9-8
 - syntax for
 - character-position form 9-10
 - delimiter form 9-6

FILE statement (*continued*)
 with dbload 9-5
FIRST clause 6-3
Fix packs
 naming conventions 1-7

G

Global Language Support (GLS)
 dbimport utility 8-8
 using onload and onunload 13-2
GRANT statement
 role privileges 10-7

H

HA_ALIAS configuration parameter B-1
Hardware
 prerequisites for migrating 1-5
High-Availability Data Replication
 disabling 3-7
 disabling before reversion 6-10
 preparing after migration 5-9
HPL_DYNAMIC_LIB_PATH configuration parameter 5-4, B-5
HPLAPIVERSION configuration parameter B-5

I

IBM Informix Storage Manager 3-5
IFX_EXTEND_ROLE configuration parameter B-4
IFX_FOLDVIEW configuration parameter B-4
IFX_LISTEN_TIMEOUT configuration parameter B-4
IFX_ONPLOAD_AUTO_UPGRADE environment variable 5-6
Importing
 non-Informix data 2-6
in dbschema output 10-7
In-place migration 1-2, 1-3, 2-1
Index
 checking 5-8
 rebuilding 5-5
INFORMIXDIR directory 5-2
INFORMIXSERVER environment variable 5-4
INFORMIXSQLHOSTS environment variable 5-4
Initializing
 Dynamic Server after upgrading 5-5
INSERT statements
 character-position form 9-9
 delimiter form 9-6
 syntax for character-position form 9-10
 with dbload 9-5
Installation directory 5-2
Installing Dynamic Server 5-2

J

Java UDRs 6-12

K

Kernel parameters on UNIX or Linux 3-8
Keywords D-1

L

Level-0 backup 3-7, 5-9, 6-13
 after moving data 13-10
LIMIT keyword 6-3
LIMITNUMSESSIONS configuration parameter B-1
Linux
 migrating on 1-8
LOAD SQL statement
 for locales that support multibyte code sets 11-2
 for non-default GL_DATETIME environment
 variables 12-1
 for non-default locales 12-1
 overview 11-1
 syntax 11-2
Loading
 ASCII files 2-6
 binary data 13-1, 13-2
 data 2-1, 2-5, 2-6
Locking
 set by onload 13-10
LOG_INDEX_BUILDS B-2
Logical log
 backup 5-5
 out of space 5-5
 space required for Dynamic Server migration 3-3
LTAPEDEV configuration parameter
 onunload/onload 13-10

M

MAX_FILL_DATA_PAGES B-3
MAX_INCOMPLETE_CONNECTIONS configuration
 parameter B-4
Migrating
 between 32-bit and 64-bit database servers 3-9
 on the same computer 1-3
 to a different computer 1-4
 to Dynamic Server Enterprise Edition 1-1
 to Dynamic Server Express Edition 1-1
Migrating a database
 constraints 2-1
 issues to consider 2-1
 overview 2-1
 prerequisites 2-1
 to a new operating system 7-1
 to Workgroup Edition 7-3
 tools 2-1, 2-4, 2-5
Migrating with Enterprise Replication 4-1
Migration
 before you begin 3-1
 checklist 3-8
 diagnostic information that you need before
 upgrading 3-8
 hardware prerequisites 1-5
 in-place 1-2, 1-3
 monitoring status with online.log 5-3
 Non-in-place 1-2, 1-4
 on Linux 1-8
 on UNIX 1-7
 on Windows 1-8
 onload utility 13-2
 onunload utility 13-2
 paths 1-5
 planning for 1-2
 prerequisites 3-1
 process 1-1, 1-2

- Migration *(continued)*
 - skills you need 1-1
 - software prerequisites 1-5
 - time it takes 1-1
 - tools 1-3
 - with Enterprise Replication 4-1
- Mode
 - checking 3-7
- Monitoring migration status 5-3
- Moving data
 - automatically 2-1
 - blobspaces 13-10
 - constraints 2-1
 - overview 2-1
 - using dbexport and dbimport 8-1
 - using dbload 9-1
 - using distributed SQL 2-7
 - using onload and onunload 13-1
 - using onunload and onload 13-10
 - when changing operating systems 7-1, 7-2, 7-3
- MSG_DATE configuration parameter B-1

N

- Native Language Support (NLS)
 - populating with dbimport 8-12
- Non-in-place migration 1-4
- Non-Informix data, importing 2-6

O

- ON-Bar utility 3-4, 5-7
 - backing up
 - Dynamic Server 11.50 5-9
 - backing up Dynamic Server 3-7
- oncheck utility
 - cc database_name option 3-6, 5-8
 - cD database_name option 3-6, 5-8
 - ce option 3-6, 5-8
 - cI database_name option 3-6, 5-8
 - cr option 3-6, 5-8
 - cs sbspace_name option 5-8
 - cS sbspace_name option 5-8
 - rebuilding table indexes 5-5
 - verifying database integrity 3-6, 5-8
- ONCONFIG environment variable 5-4
- ONCONFIG file
 - changes in Dynamic Server 11.50 C-1
 - customizing after migration 5-4
- onconfig.std file
 - changes in Dynamic Server 11.50 C-1
- oninit utility
 - s option 3-6
 - starting Dynamic Server 5-5
- ONLIDX_MAXMEM configuration parameter B-5
- online.log 5-3
- onload and onunload utilities 13-10
- onload utility
 - constraints 13-9
 - constraints on use 13-2
 - create options 13-7
 - handling large objects in a blobspace 13-9
 - how it works 13-2
 - logging status 13-9
 - moving a database 13-10
 - moving a table 13-10

- onload utility *(continued)*
 - moving locales 13-2
 - moving to another dbspace 13-10
 - ownership and privileges 13-10
 - specifying source parameters 13-7
 - syntax 13-6
 - using between computers 13-1
- onmode utility
 - b option 6-11
 - d standard option 3-7
 - ky option 3-5
 - sy option 3-5
 - reverting from the current Dynamic Server version 6-11
 - shutting down 3-5
 - shutting down Dynamic Server 3-5
- onmode-b command 12-1
 - syntax 12-1
- onpload database
 - upgrading 5-6
- onstat utility 3-7
- ontape utility
 - a option 5-5
 - backing up
 - Dynamic Server 11.50 5-9
 - Dynamic Server source database server 3-7
- onunload utility
 - constraints on use 13-2, 13-4
 - destination parameters 13-4
 - how it works 13-2
 - locking 13-5
 - logging mode 13-5
 - moving a database 13-10
 - moving a table 13-10
 - moving locales 13-2
 - moving to another dbspace 13-10
 - ownership and privileges 13-5
 - syntax 13-3
 - unloading tables 13-5
 - using between computers 13-1
 - what is included with a
 - database 13-5
 - table 13-5
- Operating system
 - adjusting tables after changing operating systems 7-2
 - moving data to another one 7-1, 7-2, 7-3
 - reconfiguring 3-8
- Operating systems
 - for Version 11.50 1-5
 - not supported for Version 11.50 1-5
- OPTICAL_LIB_PATH configuration parameter 5-4
- ORDER BY clause 6-3
- Overview of migration process 1-1, 1-2, 1-3

P

- PATH environment variable 5-4
- Performance tuning
 - after upgrading 5-9
- Planning
 - data migration 2-1
- Platforms, moving data between
 - compatible computers 2-6
- PLCY_HASHSIZE B-3
- PLCY_POOLSIZ E B-3
- PLOG_OVERFLOW_PATH configuration parameter B-5
- Privileges 10-7
 - required for onunload 13-4

Q

Quiescent mode 3-7

R

Recompiling Java UDRs 6-12

Reconfiguring the operating system 3-8

Registering DataBlade modules 5-9

Reserve pages, checking 5-8

Reserved words D-1

RESTORE_FILTER configuration parameter B-3

Reversion utility 6-11

Reverting

before using the onmode -b command 12-1

chunk reserve pages 6-8

determining if it is possible 6-2

disabling High-Availability Data Replication 6-10

from Dynamic Server Version 11.50 6-1, 6-10

from Dynamic Server with Enterprise Replication 4-3

from Version 11.50 6-1

limitations 6-3

overview 6-1

removing Version 11.50 features 6-10

restrictions for 6-3

restrictions for reverting to prior versions 6-2

using the onmode -b command 12-1

ROOTOFFSET configuration parameter 5-4

ROOTPATH configuration parameter 5-4

ROOTSIZE configuration parameter 5-4

RTO_SERVER_RESTART configuration parameter B-3

Running the reversion utility 6-11

S

Schema

create across a network 10-4

create for a database 10-3

display with dbschema 10-1

Screen reader

reading syntax diagrams G-1

Scripts

concdr.bat 4-1

concdr.sh 4-1

conpload.sh 5-6

conploadlegacy.sh 5-6

revcdr.bat 4-3

revcdr.sh 4-3

SDS_ENABLE configuration parameter B-3

SDS_PAGING configuration parameter B-3

SDS_TEMPDBS configuration parameter B-3

SDS_TIMEOUT configuration parameter B-3

SECURITY_LOCALCONNECTION configuration parameter B-4

SELECT triggers, disabling with dbexport 8-2

SHMNOACCESS configuration parameter B-1

SHMVIRT_ALLOCSEG configuration parameter B-3

shortcut keys

keyboard G-1

Simple large objects

moving with onload 13-9, 13-10

SKIP keyword 6-3

sm_versions file 3-5

sm_versions.std file

renaming to sm_versions 5-7

smi_unld utility 6-9

Software

prerequisites for migrating 1-5

Space

checking availability before migration 3-3

for sysmaster database 3-3

SQL reserved words D-1

SQL statements

UPDATE STATISTICS

data distributions 10-10

sqlhosts file, UNIX

changing name or path 5-4

csn option 5-5

save a copy when migrating 3-4

SQLTRACE configuration parameter B-3

SSL_KEYSTORE_LABEL configuration parameter B-1

SSL_KEYSTORE_FILE configuration parameter B-1

SSL_KEYSTORE_STH configuration parameter B-1

Starting

Dynamic Server after upgrading 5-5

target database server after reversion 6-12

Storage manager 3-5

STORAGE_FULL_ALARM configuration parameter B-1

Syntax diagrams

reading in a screen reader G-1

sysadmin database E-1

sysbufpool table E-2

syscdc database E-1

sysdirectives table E-2

sysindexes

changes to E-6

sysmaster database

and logical logs 5-5

changes to E-1

space required for migration 3-3

sysessappinfo database E-1

System catalogs

boot scripts 6-1

changes to E-1, E-5

checking tables 5-8

T

TAPEDEV configuration parameter, with onunload and onload 13-10

TBLTBLFIRST configuration parameter B-5

TBLTBLNEXT configuration parameter B-5

tctermcap file

saving 3-4

TEMPTAB_NOLOG configuration parameter B-3

termcap file in \$INFORMIXDIR/etc 3-4

Transactions

checking for open ones 3-6

Truncate keyword 6-3

U

UNIX

migrating on 1-7

UNLOAD SQL statement

for locales that support multibyte code sets 11-2

for non-default GL_DATETIME environment

variables 12-1

for non-default locales 12-1

overview 11-1

syntax 11-1

UPDATABLE_SECONDARY configuration parameter B-2

- UPDATE STATISTICS statement
 - data distributions 10-10
 - using after migrating to Version 11.50 5-7
 - using after reversion 6-13
- Upgrading 1-3
 - migration paths 1-5
- USE_DTENV environment variable 8-2, 8-7, 12-1
- USELASTCOMMITTED configuration parameter B-3
- User-defined data types
 - converting replication of from Dynamic Server 9.21 4-2
- USRC_HASHSIZE configuration parameter B-3
- USRC_POOLSIZE configuration parameter B-3
- Utilities
 - Abbreviated years 8-2
 - dbexport 7-1, 8-1, 8-2
 - dbexport syntax 8-2
 - dbimport 7-1, 8-1, 8-6
 - dbimport syntax 8-7
 - dbload 9-1
 - dbschema 10-1
 - Environment variables
 - DBCENTURY 8-2
 - DBDATE 8-2
 - onload 7-1, 13-6
 - onload and onunload 13-1, 13-10
 - onunload 7-1, 13-2

V

- Verifying data integrity 5-8
- Visual disabilities
 - reading syntax diagrams G-1
- VP_MEMORY_CACHE_KB B-4

W

- Windows operating systems
 - migrating on 1-8



Printed in USA

SC23-7758-02



Spine information:

IBM Informix **Version 11.50**

IBM Informix Migration Guide

