



## **IBM Informix Migration Guide**





## **IBM Informix Migration Guide**

**Note:**

Before using this information and the product it supports, read the information in “Notices” on page A-1.

**Second Edition (August 2004)**

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, 2004. All rights reserved.

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

---

# Contents

<b>Introduction.</b>	<b>ix</b>
About This Manual	ix
Types of Users	x
Software Dependencies	x
Assumptions About Your Locale	xi
Demonstration Databases	xii
New Features in Database Servers	xii
Documentation Conventions	xii
Typographical Conventions	xiii
Other Conventions	xiii
Syntax Diagrams	xiv
Sample-Code Conventions	xvii
Additional Documentation	xviii
Related Reading	xx
Compliance with Industry Standards	xx
IBM Welcomes Your Comments	xx

---

## Part 1. Overview of Informix Migration

<b>Chapter 1. Database Server Migration</b>	<b>1-1</b>
IBM Informix Database Server Products	1-1
Migration on the Same Operating System.	1-4
Source and Target Database Servers on UNIX	1-4
Source and Target Database Servers on Linux	1-4
Source and Target Database Server on Windows XP	1-5
Source and Target Database Servers on Windows NT	1-5
Source and Target Database Servers on Windows 2000	1-5
Source and Target Database Servers on Windows 95	1-6
Migration on Different Operating Systems	1-6
Database Server Migration Paths.	1-7
<b>Chapter 2. Data Migration</b>	<b>2-1</b>
Migrating a Database or Selected Data.	2-2
Migrating Between Different Versions of a Database Server	2-2
Changing Database Servers, Operating Systems, or GLS Locales	2-2
Distributing a Client Application.	2-2
Importing Non-Informix Data.	2-3
Setting Environment Variables Before Using Utilities	2-3
Choosing Data-Migration Tools	2-3
Automatic Data Migration	2-5
The onunload and onload Utilities	2-5
The dbexport and dbimport Utilities	2-8
LOAD, UNLOAD, and dbload	2-9
The dbschema Utility	2-11

External Tables . . . . .	2-12
The High-Performance Loader . . . . .	2-12
Nonlogging Raw Tables (IDS 9.x) . . . . .	2-13
Movement of TEXT and BYTE Data . . . . .	2-13
Moving Data Between Computers and Dbspaces . . . . .	2-14
Importing Data from a Non-Informix Source . . . . .	2-14
Importing Data with IBM Informix Enterprise Gateway Products . . . . .	2-14

---

## Part 2. Migration to a Later Version of a Database Server

### Chapter 3. Preparing for Migration . . . . . 3-1

Preparing for Migration. . . . .	3-1
Migration Guidelines . . . . .	3-2
Changes in Dynamic Server 9.40, 9.30, and 9.2x. . . . .	3-2
Storage-Manager Validation and Installation . . . . .	3-27
Migrating Between 32-bit and 64-Bit Database Servers . . . . .	3-28

### Chapter 4. Migrating to Dynamic Server 9.40 with Enterprise Replication. . . . . 4-1

Converting to Dynamic Server 9.40 with Enterprise Replication . . . . .	4-1
Converting Replication of 9.2x User-Defined Data Types. . . . .	4-2
Reverting from Dynamic Server 9.40 with Enterprise Replication . . . . .	4-3

### Chapter 5. Converting to Dynamic Server 9.40 . . . . . 5-1

Converting to Version 9.40. . . . .	5-1
Check and Configure Available Space . . . . .	5-3
Save Copies of the Current Configuration Files . . . . .	5-5
Remove Outstanding In-Place Alters . . . . .	5-5
Close All Transactions and Shut Down the Source Database Server . . . . .	5-6
Check for Any Open Transactions . . . . .	5-7
Verify the Integrity of the Data . . . . .	5-8
Verify the Database Server Mode . . . . .	5-8
Disable High-Availability Data Replication . . . . .	5-8
Make a Final Backup of the Source Database Server . . . . .	5-9
Verify That the Source Database Server Is Offline . . . . .	5-9
On UNIX or Linux, Modify Kernel Parameters . . . . .	5-9
Install Dynamic Server 9.40 . . . . .	5-10
Set Environment Variables . . . . .	5-10
Upgrade the High-Performance Loader onpload Database . . . . .	5-11
Customize Configuration Files . . . . .	5-12
Add Any Communications Support Modules (UNIX/Linux) . . . . .	5-13
Install and Configure Any DataBlade Modules. . . . .	5-13
Initialize Dynamic Server 9.40 . . . . .	5-13
Monitor the Conversion Complete Status . . . . .	5-14
For ON-Bar, Rename the sm_versions.std File . . . . .	5-15
Update Statistics. . . . .	5-15
Verify the Integrity of the Data . . . . .	5-15
Make an Initial Backup of Dynamic Server 9.40 . . . . .	5-16
Tune Dynamic Server 9.40 for Performance . . . . .	5-16
Enable HDR . . . . .	5-16

<b>Chapter 6. Reverting from Dynamic Server 9.40.</b>	<b>6-1</b>
Reverting from Dynamic Server 9.40 to Versions 9.30, 9.2x, 7.3x, or 7.24.	6-1
Determine Whether Reversion Is Possible.	6-2
Check and Configure Available Space	6-6
Save Copies of the Current Configuration Files (UNIX/Linux).	6-6
Verify the Integrity of the Data	6-6
Back Up Dynamic Server 9.40.	6-6
Remove BladeManager Extensions	6-6
Remove Version 9.40 Features.	6-7
Disable HDR	6-7
Run the Reversion Utility	6-7
Modify Configuration Parameters	6-8
Reset Environment Variables	6-8
Remove Any Communications Support Module Settings (UNIX/Linux)	6-8
Reinstall and Start the Target Database Server	6-8
Update Statistics	6-8
Verify the Integrity of the Data	6-9
Back Up the Target Database Server	6-9
Return the Target Database Server to Online Mode	6-9
Enable HDR	6-9

---

## Part 3. Migration of Data Between Database Servers

<b>Chapter 7. Migrating Between Database Servers and Operating Systems.</b>	<b>7-1</b>
Choosing a Migration Method	7-1
Adjusting Database Tables for File-System Variations	7-2
Moving Data to a Database Server on a Different Operating System	7-2
Using the Migration Utilities	7-3
Adapting Your Programs for UNIX or Windows	7-3
Completing Migration	7-4
Moving Data Between Dynamic Server and Workgroup Edition on Different Operating Systems	7-4

---

## Part 4. Data Migration Utilities

<b>Chapter 8. The dbexport and dbimport Utilities</b>	<b>8-1</b>
Syntax of the dbexport Command	8-2
Syntax of the dbimport Command	8-7
Simple Large Objects (IDS 9.x)	8-14
Database Locale Changes	8-14
<b>Chapter 9. The dbload Utility</b>	<b>9-1</b>
Syntax of the dbload Command	9-1
Command File for dbload	9-5
Command File to Load Complex Data Types (IDS 9.x)	9-14
<b>Chapter 10. The dbschema Utility</b>	<b>10-1</b>
Syntax of the dbschema Command	10-2
Database Schema Creation	10-3
Server-Specific Information	10-5

User-Defined and Complex Data Types (IDS 9.x) . . . . .	10-5
Sequence Creation . . . . .	10-6
Synonym Creation . . . . .	10-7
Privileges . . . . .	10-7
Granting Privileges . . . . .	10-8
Displaying Privilege Information for a Role. . . . .	10-8
Table, View, or Procedure Creation. . . . .	10-8
Table Information . . . . .	10-9
Role Creation . . . . .	10-10
Distribution Information for Tables . . . . .	10-10
Example Output . . . . .	10-11
Distribution Description . . . . .	10-11
Distribution Information. . . . .	10-12
Overflow Information . . . . .	10-13
DB-Access Input from dbschema Output . . . . .	10-13
Inserting a Table into a Database Example. . . . .	10-13
Re-Creating the Schema of a Database . . . . .	10-13
 <b>Chapter 11. The LOAD and UNLOAD Statements . . . . .</b>	 <b>11-1</b>
Syntax of the UNLOAD Statement. . . . .	11-1
Syntax of the LOAD Statement . . . . .	11-2
 <b>Chapter 12. The onmode Utility . . . . .</b>	 <b>12-1</b>
Use of the onmode -b Command for Reversion . . . . .	12-1
Preparation for Reversion. . . . .	12-2
Syntax of the onmode -b Command . . . . .	12-2
 <b>Chapter 13. The onunload and onload Utilities. . . . .</b>	 <b>13-1</b>
How onunload and onload Work . . . . .	13-1
Syntax of the onunload Command. . . . .	13-2
Destination Parameters . . . . .	13-3
Constraints That Affect onunload . . . . .	13-3
Database or Table Unloading . . . . .	13-4
Logging Mode . . . . .	13-5
Locking During Unload Operation. . . . .	13-5
Syntax of the onload Command . . . . .	13-5
Source Parameters . . . . .	13-6
Create Options . . . . .	13-7
Constraints That Affect onload . . . . .	13-8
Constraints That Affect onload and onunload . . . . .	13-9
Restrictions That Affect onload and onunload . . . . .	13-10
Logging During Loading . . . . .	13-10
Movement of Simple Large Objects to a BlobSpace . . . . .	13-11
Ownership and Privileges . . . . .	13-11
Exclusive Locking During Load Operation. . . . .	13-11
Steps for Using onunload and onload . . . . .	13-11

---

## Part 5. Appendixes



**Notices . . . . . A-1**

**Index . . . . . X-1**



---

# Introduction

About This Manual . . . . .	ix
Types of Users . . . . .	x
Software Dependencies . . . . .	x
Assumptions About Your Locale . . . . .	xi
Demonstration Databases . . . . .	xii
New Features in Database Servers . . . . .	xii
Documentation Conventions . . . . .	xii
Typographical Conventions . . . . .	xiii
Other Conventions . . . . .	xiii
Comments . . . . .	xiii
Feature, Product, and Platform . . . . .	xiv
Syntax Diagrams . . . . .	xiv
How to Read a Command-Line Diagram . . . . .	xvi
Keywords and Punctuation . . . . .	xvii
Identifiers and Names . . . . .	xvii
Sample-Code Conventions . . . . .	xvii
Additional Documentation . . . . .	xviii
Related Reading . . . . .	xx
Compliance with Industry Standards . . . . .	xx
IBM Welcomes Your Comments . . . . .	xx

---

## In This Introduction

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

---

## About This Manual

This manual describes the migration procedures to convert to Version 9.4 of the IBM Informix database server, revert to older database servers, and move data between databases, database servers, and computers. This manual 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.

For information on migrating to other IBM Informix database server versions, see the *Migration Guide* that is included in the documentation set for that version of the server.

*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.

## Types of Users

This manual is for the following users:

- Database users
- Database administrators
- Database server administrators
- System administrators
- Database-application programmers
- Backup operators
- Performance engineers

To understand this manual, you need to have the following background:

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

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

## Software Dependencies

The table below lists IBM Informix database servers with their short database server names.

*Table 1. IBM Informix 9.x, 7.x, and 5.x Database Servers*

Short Database Server Name	Complete Database Server Name
Dynamic Server 9.40	IBM Informix Dynamic Server, Version 9.40
Dynamic Server 9.30	IBM Informix Dynamic Server, Version 9.30
Dynamic Server 9.21	IBM Informix Dynamic Server, Version 9.21
Workgroup Edition 9.21	IBM Informix Dynamic Server, Version 9.21
Dynamic Server 9.20	IBM Informix Dynamic Server, Version 9.20
Dynamic Server 7.31	IBM Informix Dynamic Server, Version 7.31
Workgroup Edition 7.31	IBM Informix Dynamic Server, Version 7.31

*Table 1. IBM Informix 9.x, 7.x, and 5.x Database Servers*

Short Database Server Name	Complete Database Server Name
Dynamic Server, Linux Edition 7.31	IBM Informix Dynamic Server, Linux Edition, Version 7.31
Dynamic Server 7.30	IBM Informix Dynamic Server, Version 7.30
Workgroup Edition 7.30	IBM Informix Dynamic Server, Version 7.30
Dynamic Server, Linux Edition 7.30	IBM Informix Dynamic Server, Version 7.30
Dynamic Server 7.24	IBM Informix Dynamic Server, Version 7.24
Workgroup Edition 7.24	IBM Informix Dynamic Server, WorkGroup Edition, Version 7.24
OnLine 5.1x	IBM Informix OnLine, Version 5.1x
SE 7.25	IBM Informix SE, Version 7.25
SE 7.24	IBM Informix SE, Version 7.24
SE 7.23	IBM Informix SE, Version 7.23
SE 7.22	IBM Informix SE, Version 7.22
SE 5.1x	IBM Informix SE, Version 5.1x

In this guide, a short database server name followed by a version number refers only to that version, or set of versions, of the database server. A version number without an “x” refers to a single version of the database server; for example, Dynamic Server 9.21 refers only to IBM Informix Dynamic Server, Version 9.21. A version number with an “x” can refer to more than one version of the database server; for example, Dynamic Server 9.x refers to all Dynamic Server Versions 9.20, 9.21, 9.30, and 9.40.

The migration can be between operating systems. Each of the Informix database servers can run on one or more of the following operating systems:

- UNIX
- Linux
- Windows 2000
- Windows NT
- Windows XP
- Windows 95

## 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 manual 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 é, è, 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 need to 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*.

## Demonstration Databases

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

- The **stores\_demo** database illustrates a relational schema with information about a fictitious wholesale sporting-goods distributor. Many examples in IBM Informix manuals are based on the **stores\_demo** database.
- The **superstores\_demo** database illustrates an object-relational schema. The **superstores\_demo** database contains examples of extended data types, type and table inheritance, and user-defined routines.

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

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

---

## New Features in Database Servers

For a comprehensive list of new database server features, see your *IBM Informix: Getting Started Guide* and your release notes. This manual includes information about new features in the following sections:

- “Changes in Dynamic Server 9.40, 9.30, and 9.2x” on page 3-2

---

## Documentation Conventions

This section describes the conventions that this manual uses. These conventions make it easier to gather information from this and other volumes in the documentation set.

The following conventions are discussed:

- Typographical conventions
- Other conventions
- Syntax diagrams
- Command-line conventions
- Sample-code conventions

## Typographical Conventions

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

Convention	Meaning
KEYWORD	All primary elements in a programming language statement (keywords) appear in uppercase letters in a serif font.
italics <i>italics</i> <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 <b>boldface</b>	Names of program entities (such as classes, events, and tables), environment variables, file and pathnames, and interface elements (such as icons, menu items, and buttons) appear in boldface.
monospace <i>monospace</i>	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 <b>Tools</b> > <b>Options</b> ” means choose the <b>Options</b> item from the <b>Tools</b> menu.

**Tip:** When you are instructed to “enter” characters or to “execute” a command, immediately press RETURN after the entry. When you are instructed to “type” the text or to “press” other keys, no RETURN is required.

## Other Conventions

Throughout the documentation, you will find text that is identified by several different types of markup.

### Comments

Comments identify three types of information, as shown in the following examples:

**Warning:** Identifies paragraphs that contain vital instructions, cautions, or critical information

**Important:** Identifies paragraphs that contain significant information about the feature or operation that is being described

**Tip:** Identifies paragraphs that offer additional details or shortcuts for the functionality that is being described

### **Feature, Product, and Platform**

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

<b>Dynamic Server</b>
Identifies information that is specific to IBM Informix Dynamic Server
<b>End of Dynamic Server</b>

<b>UNIX Only</b>
Identifies information that is specific to UNIX platforms
<b>End of UNIX Only</b>

<b>Windows Only</b>
Identifies information that is specific to the Windows environment
<b>End of Windows Only</b>

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:

#### **Table Sorting (Windows Only)**

### **Syntax Diagrams**

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

**Note:** Starting with Version 9.4, syntax diagrams have been reformatted to conform to the IBM standard.







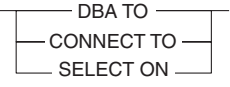
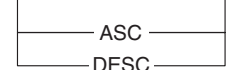
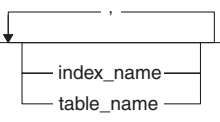


Syntax diagrams depicting SQL and command-line statements have changed in the following ways:

- The symbols at the beginning and end of statements are now double arrows instead of a vertical line at the end.



- The symbols at the beginning and end of syntax segment diagrams are now vertical lines instead of arrows.
- How many time a loop can be repeated is now explained in a diagram footnote instead of a number in a gate symbol.
- Syntax statements that are longer than one line now continue on the next line instead of looping down with a continuous line.
- Product or condition-specific paths are now explained in diagram footnotes instead of icons.

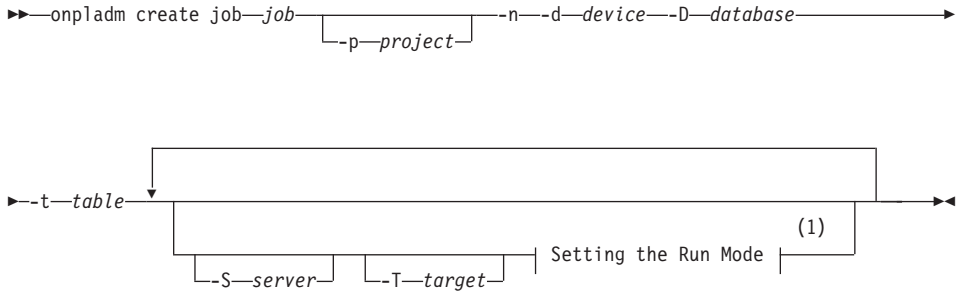
The following table describes syntax diagram components.

Component	Meaning
	Statement begins.
	Statement continues on next line. Syntax elements other than complete statements end with this symbol.
	Statement continues from previous line. Syntax elements other than complete statements begin with this symbol.
	Statement ends.
	Required item.
	Optional item.
	Required item with choice. One and only one item must be present.
	Optional item with choice. If a default value exists, it is printed above the line.
	Optional items. Several items are allowed; a comma must precede each repetition.
	Start of syntax segment
	End of syntax segment

## How to Read a Command-Line Diagram

The following command-line diagram uses some of the elements listed in the table in the previous section.

### Creating a No-Conversion Job




#### Notes:

- 1 See page 17-4

The second line in this diagram has a segment named “Setting the Run Mode,” which according to the diagram footnote, is on page 17-4. This segment is shown in the following segment diagram (the diagram uses segment start and end components).

### Setting the Run Mode:



To construct a command correctly, start at the top left with the command. Follow the diagram to the right, including the elements that you want. The elements in the diagram are 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.

Your diagram is complete.

### 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.

### Sample-Code Conventions

Examples of SQL code occur throughout this manual. 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 DB–Access, you must delimit multiple statements with semicolons. 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.

**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 manual for your product.

---

## Additional Documentation

IBM Informix Dynamic Server documentation is provided in a variety of formats:

- **Online manuals.** The documentation CD in your media pack allows you to print the product documentation. You can obtain the same online manuals at the IBM Informix Online Documentation site at <http://www.ibm.com/software/data/informix/pubs/library>.
- **Online help.** This facility provides context-sensitive help, an error message reference, language syntax, and more.
- **Documentation notes, release notes, and machine notes.** Documentation notes, which contain additions and corrections to the manuals, and release notes and machine notes are located in the directory where the product is installed.

Please examine these files because they contain vital information about application and performance issues.

The following table describes these files.

UNIX/Linux Only
-----------------

On UNIX or Linux, the following online files are in the <b>\$INFORMIXDIR/release/en_us/0333</b> directory.
--

Online File	Purpose
<b>ids_migrate_docnotes_9.40.html</b>	The documentation notes file for your version of this manual describes topics that are not covered in the manual or that were modified since publication.
<b>ids_unix_release_notes_9.40.html</b> <b>ids_unix_release_notes_9.40.txt</b>	The release notes file (in HTML or text format) describes feature differences from earlier versions of IBM Informix products and how these differences might affect current products. This file also contains information about known problems and their workarounds.
<b>ids_machine_notes_9.40.txt</b>	The machine notes file describes special actions that you must take to configure and use IBM Informix products on your computer.

End of UNIX/Linux Only

#### Windows Only

On Windows, the following items are in the **Informix** folder. To display this folder, choose **Start > Programs > Informix > Documentation Notes** or **Release Notes** from the taskbar.

Program Group Item	Description
<b>Documentation Notes</b>	The documentation notes file for your version of this manual describes topics that are not covered in the manual or that were modified since publication.
<b>Release Notes</b>	The release notes file describes feature differences from earlier versions of IBM Informix products and how these differences might affect current products. This file also contains information about known problems and their workarounds.

Machine notes do not apply to Windows.

End of Windows Only

- **Error message files.** IBM Informix software products provide ASCII files that contain error messages and their corrective actions.

---

#### UNIX/Linux Only

To read error messages and corrective actions on UNIX or Linux, you can use the **finderr** command to display the error messages online.

---

End of UNIX/Linux Only

---

#### Windows Only

To read error messages and corrective actions on Windows, use the **Informix Error Messages** utility. To display this utility, choose **Start > Programs > Informix** from the taskbar.

---

End of Windows Only

---

## Related Reading

For a list of publications that provide an introduction to database servers and operating-system platforms, refer to your *IBM Informix: Getting Started Guide*.

---

## Compliance with Industry Standards

The American National Standards Institute (ANSI) has established a set of industry standards for 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 Informix database servers comply with the SQL-92 Intermediate and Full Level and X/Open SQL CAE (common applications environment) standards.

---

## IBM Welcomes Your Comments

We want to know about any corrections or clarifications that you would find useful in our manuals that would help us with future versions. Include the following information:

- The name and version of the manual that you are using
- Any comments that you have about the manual
- Your name, address, and phone number

Send electronic mail to us at the following address:

docinf@us.ibm.com

This address is reserved for reporting errors and omissions in our documentation. For immediate help with a technical problem, contact Customer Services.

We appreciate your suggestions.





---

## **Part 1. Overview of Informix Migration**



---

# Chapter 1. Database Server Migration

IBM Informix Database Server Products . . . . .	1-1
Migration on the Same Operating System. . . . .	1-4
Source and Target Database Servers on UNIX . . . . .	1-4
Source and Target Database Servers on Linux . . . . .	1-4
Source and Target Database Server on Windows XP . . . . .	1-5
Source and Target Database Servers on Windows NT . . . . .	1-5
Source and Target Database Servers on Windows 2000 . . . . .	1-5
Source and Target Database Servers on Windows 95 . . . . .	1-6
Migration on Different Operating Systems . . . . .	1-6
Database Server Migration Paths. . . . .	1-7

---

## In This Chapter

This chapter provides an overview of database server migration, including:

- A list of all IBM Informix database server products and their operating systems
- Information about migration on the same operating system
- Information about migration on different operating systems
- Database server migration paths from previous IBM informix database servers to Version 9.4 of the IBM Informix Dynamic Server

For information about how to move data between database servers on different operating systems, see Chapter 7, “Migrating Between Database Servers and Operating Systems,” on page 7-1

---

## IBM Informix Database Server Products

Table 1-1 lists all Version 9.x, 8.x, 7.x, and 5.1x IBM Informix database servers with the operating systems on which you can run them. Table 1-1 uses short names for the database servers. For the complete database server names that correspond to these short names, see “Software Dependencies” on page x of the Introduction.

*Table 1-1. Operating Systems on Which to Run Informix Database Servers*

Database Server	Operating Systems
Dynamic Server 9.40	UNIX Linux Windows 2000 Windows XP

*Table 1-1. Operating Systems on Which to Run Informix Database Servers (continued)*

<b>Database Server</b>	<b>Operating Systems</b>
Dynamic Server 9.40	UNIX Linux Windows 2000 Windows XP
Dynamic Server 9.30	UNIX Linux Windows 2000 Windows NT
Dynamic Server 9.21	UNIX Linux Windows 2000 Windows 95
Dynamic Server 9.20	UNIX Linux Windows 2000 Windows NT
Extended Parallel Server 8.40	UNIX
Extended Parallel Server 8.32	UNIX
Extended Parallel Server 8.31	UNIX
Extended Parallel Server 8.30	UNIX
Dynamic Server with AD and XP Options 8.21	UNIX Linux Windows NT
Dynamic Server 7.31	UNIX Windows 2000 Windows NT Windows 95
Workgroup Edition 7.31	UNIX Windows 2000 Windows NT Windows 95
Dynamic Server, Linux Edition 7.31	Linux UNIX
Dynamic Server 7.30	UNIX Linux Windows NT Windows 95
Workgroup Edition 7.30	UNIX Windows NT Windows 95

*Table 1-1. Operating Systems on Which to Run Informix Database Servers (continued)*

<b>Database Server</b>	<b>Operating Systems</b>
Dynamic Server, Linux Edition 7.30	Linux UNIX
Dynamic Server 7.24	UNIX Windows NT Windows 95
Workgroup Edition 7.24	UNIX Windows NT Windows 95
OnLine 5.1x	UNIX Linux
SE 7.25	UNIX
SE 7.24	UNIX Linux Windows NT
SE 7.23	UNIX Windows NT
SE 7.22	UNIX Windows NT
SE 5.1x	UNIX

If you have Dynamic Server Version 9.30, 9.2x, 7.3x, or 7.24, you can migrate directly to Dynamic Server Version 9.40, the current database server.

You can also migrate to the current database server from the following database servers after you migrate to an intermediary database server:

- IBM Informix Universal Server, Version 9.14 (Universal Server 9.14)
- IBM Informix OnLine Dynamic Server, Version 7.22 through Version 7.23 (OnLine Dynamic Server 7.22 through 7.23)

For information on migration paths, see “Database Server Migration Paths” on page 1-7.

To become more familiar with the Informix client-server environment, read *IBM Informix: Getting Started with Database Server Products*. This book contains information on the differences between Informix database servers plus information on network and server configurations. Also, read the *IBM Informix: Getting Started Guide* for your database server.

## Migration on the Same Operating System

This section provides tables that show migration paths between database servers on the same operating system.

### Source and Target Database Servers on UNIX

The table below lists the source and target versions for migration of database servers on UNIX.

Table 1-2. Migrating Between Database Servers on UNIX

Source Version	Target Version												
	9.40	9.30	9.20-21	9.14	8.40	8.31-32	8.30	8.21.UD4	8.21	7.30-31	7.25	7.24	5.1x
9.40	x	x	x	x						x		x	
9.30	x	x	x	x						x		x	
9.21	x	x	x	x						x		x	
9.20	x	x	x	x						x		x	
9.14		x	x	x						x		x	
8.40					x	x	x	x		x		x	
8.32					x	x	x	x		x		x	
8.31					x	x	x	x		x		x	
8.30					x	x	x	x		x		x	
8.21.UD4							x	x	x			x	
8.21								x	x			x	
7.31	x	x	x	x	x	x	x			x		x	x
7.30	x	x	x	x	x	x	x			x		x	x
7.25											x	x	x
7.24	x	x	x	x	x	x	x	x		x	x	x	x
7.23			x	x						x	x	x	x
7.22			x	x						x	x	x	x
5.1x.										x	x	x	x

### Source and Target Database Servers on Linux

The table below lists the source and target versions for migration of database servers on Linux.

Table 1-3. Migrating Between Database Servers on Linux

Source Version	Target Version			
	9.40.U	9.30.U	7.31.U	7.30.U
9.40.U	x	x	x	x
9.30.U	x	x	x	x
9.21.U	x	x	x	x
9.20.U	x	x	x	x
7.31.U	x	x	x	x
7.30.U	x	x	x	x

### Source and Target Database Server on Windows XP

Currently, only Dynamic Server Versions 9.4 is available on Windows XP.

### Source and Target Database Servers on Windows NT

**Important:** Dynamic Server Version 9.4 is not available on Windows NT.

Table 1-4 lists the source and target versions for migration of previous database server versions on Windows NT.

Table 1-4. Migrating Between Database Servers on Windows NT

Source Version	Target Version						
	9.30.T	8.21.T	7.31.T	7.30.T	7.24.T	7.23.T	7.22.T
9.30.T	x		x	x	x		
9.21.T	x		x	x	x	x	x
9.20.T	x		x	x	x	x	x
9.14.T	x		x	x	x	x	x
8.21.T		x			x	x	x
7.31.T	x	x	x	x	x	x	x
7.30.T	x	x	x	x	x	x	x
7.24.T	x	x	x	x	x	x	x
7.23.T		x	x	x	x	x	x
7.22.T		x	x	x	x	x	x

### Source and Target Database Servers on Windows 2000

The table below lists the source and target versions for migration of Informix database servers on Windows 2000.

Table 1-5. Migrating Between Database Servers on Windows 2000

Source Version	Target Version		
	9.40.T	9.30.T	7.31.T
9.40.T	x	x	x
9.30.T	x	x	x
9.21.T	x	x	x
7.31.T	x	x	x

## Source and Target Database Servers on Windows 95

Table 1-6 lists the source and target versions for migration of Informix database servers on Windows 95.

Table 1-6. Migrating Between Database Servers on Windows 95

Source Version	Target Version		
	7.31.T	7.30.T	7.24.T
7.31.T	x	x	x
7.30.T	x	x	x
7.24.T	x	x	x

## Migration on Different Operating Systems

If you are using Dynamic Server on Windows NT and migrating to Dynamic Server Version 9.40, you also need to migrate to Windows 2000 or Windows XP.

For specific information on migrating between operating systems, see Chapter 7, “Migrating Between Database Servers and Operating Systems,” on page 7-1.

Table 1-7 lists the operating systems supported by each version of *Dynamic Server*.



Table 1-7. Migrating Between Operating Systems

Server Version	Operating System					
	UNIX	Linux	Windows XP	Windows NT	Windows 2000	Windows 95
9.40	x	x	x		x	
9.30	x	x		x	x	
9.21	x	x		x	x	
9.20	x	x		x		
9.14	x			x		
8.40	x					
8.31	x					
8.30	x					
8.21	x			x		
7.31	x	x		x	x	x
7.30	x	x		x		x
7.25	x					
7.24	x			x		x
7.23	x			x		
7.22	x			x		
5.1x	x					

## Database Server Migration Paths

Table 1-8 lists migration paths for converting from an existing source database server to a newer target database server. The **Reference** column shows where to find information about how to convert.

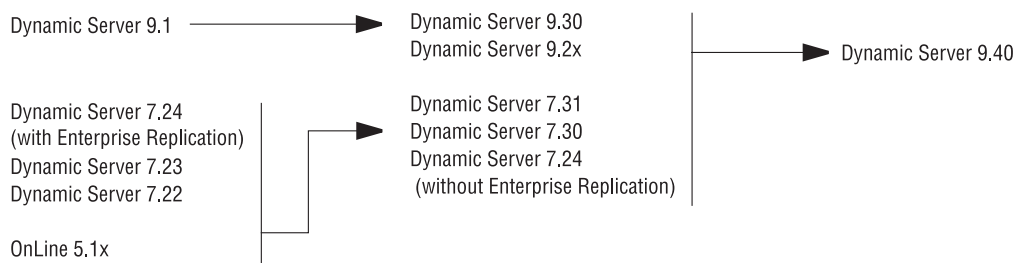
Table 1-8. Converting to a Newer Database Server

Source Database Server	Target Database Server	Reference
Dynamic Server 9.30 Dynamic Server 9.20 or Dynamic Server 9.21	Dynamic Server 9.40	Chapter 3, "Preparing for Migration," on page 3-1 and Chapter 5, "Converting to Dynamic Server 9.40," on page 5-1
Universal Server 9.14	Dynamic Server 9.30	The Version 9.30 Migration Guide

*Table 1-8. Converting to a Newer Database Server (continued)*

Source Database Server	Target Database Server	Reference
Dynamic Server 7.31 Dynamic Server, Linux Edition 7.31 Workgroup Edition 7.31 Dynamic Server 7.30 Dynamic Server, Linux Edition 7.30 Workgroup Edition 7.30 Dynamic Server 7.24 Workgroup Edition 7.24	Dynamic Server 9.40	Chapter 3, "Preparing for Migration," on page 3-1 and Chapter 5, "Converting to Dynamic Server 9.40," on page 5-1
OnLine Dynamic Server 7.23 OnLine Dynamic Server 7.22 OnLine 5.1x	Dynamic Server 7.31	The Version 7.31 Migration Guide

Figure 1-1 shows the paths for converting from these older database servers to Dynamic Server Version 9.40.



*Figure 1-1. Migration Paths to Dynamic Server 9.40*

For information on migrating directly to Dynamic Server Version 9.40 from a 9.30, 9.2x, 7.3x, or 7.24 database server, see Chapter 3, "Preparing for Migration," on page 3-1 and Chapter 5, "Converting to Dynamic Server 9.40," on page 5-1.

If you need to migrate to another version of Dynamic Server before you migrate to Version 9.40, see the Migration Guide that is included in the documentation set for that database server. For example, if you need to migrate to Version 7.31 before migrating to Version 9.40, see the Version 7.31 *IBM Informix: Migration Guide* for information on migrating to Version 7.31. You must follow the steps in that guide before you follow the instructions in the current guide for migrating to Version 9.40.

Table 1-9 shows the database servers to which you can revert from Dynamic Server 9.40. After you install Version 9.40, you cannot revert from Version 9.40 to any IBM Informix database server that is not shown in the list. The **Page** column shows where to find information in this manual about how to revert.

*Table 1-9. Reverting to an Older Database Server*

Source	Target	Page
Dynamic Server 9.40	Dynamic Server 9.30	Chapter 6, “Reverting from Dynamic Server 9.40,” on page 6-1
	Dynamic Server 9.2x	
	Dynamic Server 7.3x	
	Dynamic Server, Linux Edition Version 7.3x	
	Workgroup Edition 7.3x	
	Dynamic Server 7.24	
	Dynamic Server, Linux Edition Version 7.3x	
	Workgroup Edition 7.24	



---

## Chapter 2. Data Migration

Migrating a Database or Selected Data . . . . .	2-2
Migrating Between Different Versions of a Database Server . . . . .	2-2
Changing Database Servers, Operating Systems, or GLS Locales . . . . .	2-2
Distributing a Client Application. . . . .	2-2
Importing Non-Informix Data. . . . .	2-3
Setting Environment Variables Before Using Utilities . . . . .	2-3
Choosing Data-Migration Tools . . . . .	2-3
Automatic Data Migration . . . . .	2-5
The onunload and onload Utilities . . . . .	2-5
Constraints on Using onunload and onload . . . . .	2-6
Restrictions on Using onunload and onload . . . . .	2-7
The dbexport and dbimport Utilities . . . . .	2-8
Destination Options . . . . .	2-9
Location Options . . . . .	2-9
Database Server Options . . . . .	2-9
LOAD, UNLOAD, and dbload . . . . .	2-9
The dbschema Utility . . . . .	2-11
Guidelines for Using dbschema . . . . .	2-11
DB-Access Input from dbschema Output. . . . .	2-11
Object Modes and Violation Detection . . . . .	2-11
External Tables . . . . .	2-12
The High-Performance Loader . . . . .	2-12
HPL Tools . . . . .	2-12
Performance Advantage of the HPL . . . . .	2-12
Nonlogging Raw Tables (IDS 9.x) . . . . .	2-13
Movement of TEXT and BYTE Data . . . . .	2-13
Moving Data Between Computers and Dbspaces . . . . .	2-14
Importing Data from a Non-Informix Source . . . . .	2-14
Importing Data with IBM Informix Enterprise Gateway Products . . . . .	2-14

---

### In This Chapter

This chapter provides an overview of data migration and compares the IBM Informix migration utilities. This chapter covers the following topics:

- Migrating a database or selected data
- Setting environment variables before using utilities
- Choosing the most effective data migration tool
- Moving data between computers and dbspaces

---

## Migrating a Database or Selected Data

You might need to perform data migration on a database or selected data to complete any of the following tasks:

- Migrating between different versions of a database server
- Changing database servers, operating systems, or GLS locales
- Distributing a client application
- Importing non-Informix data

### Migrating Between Different Versions of a Database Server

When you convert to a later version of a database server or revert to an earlier version, you need to consider the following data migration issues:

- Changes in the configuration parameters and environment variables
- Amount of memory and dbospace space required
- Organization of the data

Later chapters discuss these issues in detail.

### Changing Database Servers, Operating Systems, or GLS Locales

You might want to move between database servers, to a different operating system, or to a different GLS locale. You also might 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 to a different operating system, see Chapter 7, “Migrating Between Database Servers and Operating Systems,” on page 7-1.

---

#### Global Language Support

---

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

---

#### End of Global Language Support

---

### Distributing a Client Application

After you convert 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. You might need to recompile or modify client applications, or update the **sqlhosts** file or registry key.

Verify that the client-application version you use is compatible with your database server version. 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.

## Importing Non-Informix Data

You can use the **dbimport** and **dbload** utilities, the High-Performance Loader (HPL), the IBM Informix Enterprise Gateway products, or external tables to import data from non-Informix sources.

---

## Setting Environment Variables Before Using Utilities

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*.

---

## Choosing Data-Migration Tools

This section provides brief descriptions of the tools, utilities, and SQL statements that you can use to move data from one database to another.

When you migrate to the current version of the database server, you can use these data-migration tools:

- The **onunload** and **onload** utilities (only between database servers of the same version)
- The **dbexport** and **dbimport** utilities
- UNLOAD and LOAD statements and the **dbload** utility
- The **dbschema** utility
- The High-Performance Loader (HPL)
- Nonlogging raw tables

For more information about the utilities and statements listed above, see the chapters in the last section of this guide.

In some previous versions of the database server, you could use IBM Informix Enterprise Command Center (IECC) to move data from one database to another.

If you need to migrate to another version of Dynamic Server before you migrate to Version 9.4 and you want information on the migration tools and utilities that can be used for migration to an intermediary version of Dynamic Server, see the Migration Guide that is included in the documentation set for that database server.

The table below contains a list of the tools you can use for each IBM Informix database server. An asterisk (\*) after the utility name identifies a tool that was used in a previous version of Dynamic Server, but is no longer used.

*Table 2-1. Utilities for Moving Data*

	<b>IDS 9.40</b>	<b>IDS .30</b>	<b>IDS 7.3x or 7.24</b>	<b>SE 7.2x and 5.1x</b>	<b>OL 5.1x</b>
<b>dbexport/ dbimport</b>	x	x	x	x	x
<b>dbload</b>	x	x	x	x	x
<b>HPL</b>	x	x	x		
<b>onunload/ onload</b>	x	x	x		
<b>UNLOAD/ LOAD</b>	x	x	x	x	x
<b>IECC*</b>			x		

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 table below summarizes the characteristics of the methods for loading data and the constraints and advantages of each method.

*Table 2-2. Comparison of Tools for Loading Data*

	<b>dbexport/ dbimport</b>	<b>dbload</b>	<b>HPL</b>	<b>onunload/ onload</b>	<b>UNLOAD/ LOAD</b>
<b>Granularity of Data</b>	Database only	Partial or complete table	Partial or complete table	Table or database	Partial or complete table
<b>Performance</b>	Moderate	Slow	Fast	Fast	Moderate
<b>Source of Data</b>	Usually produced by <b>dbexport</b>	Any data in the format specified by the input file	Any ASCII or COBOL data. User can create custom read capabilities.	Must be produced by <b>on-unload</b>	Any data in the specified format, usually produced by UNLOAD
<b>Database Schema</b>	Can modify	Can modify	Can modify	Cannot modify	Can modify
<b>Location of Data</b>	Disk or tape	Disk only	Disk, tape, or pipe	Disk or tape	Disk only
<b>Type of File</b>	Text	Text	Text	Binary	Text
<b>Logging Status</b>	Logging optional	Logging optional	Logging optional	Logging must be turned off	Logging optional



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

	<b>dbexport/ dbimport</b>	<b>dbload</b>	<b>HPL</b>	<b>onunload/ onload</b>	<b>UNLOAD/ LOAD</b>
<b>Move Data Between Operating Systems?</b>	Yes	Yes	Yes, or from a non-Informix database	No	Yes
<b>Ease of Use</b>	Moderate	Moderate	Most difficult	More difficult	Easiest

The following sections provide guidelines on how to choose the appropriate migration tool or tools.

### Automatic Data Migration

Data migration is automatic. Automatic data migration means that when you migrate from one database server to another, you do not need to use any data migration tools to move the data. The data migrates automatically from the source database server to the target database server after you bring up the target database server.

Data migration can also be automatic when you move between different versions of a database server in the same operating system.

### The onunload and onload Utilities

The **onunload** and **onload** utilities provide the fastest way to move data, but they do not let you modify the database schema or move from one operating system or database server version to another. The **onunload** utility unloads data from the specified database or table onto a tape or a file on disk in disk-page-sized units, making this utility more efficient than **dbexport**. 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.

Because the data is written in page-sized units, you can use **onunload** and **onload** only when certain constraints are met. For example, you cannot use **onunload** and **onload** to move data between UNIX and Windows. You can, however, use the **onunload** and **onload** utilities to move data between computers that use the same database server on the same platform. Figure 2-1 summarizes the questions that you need to ask before you use **onunload** and **onload**.

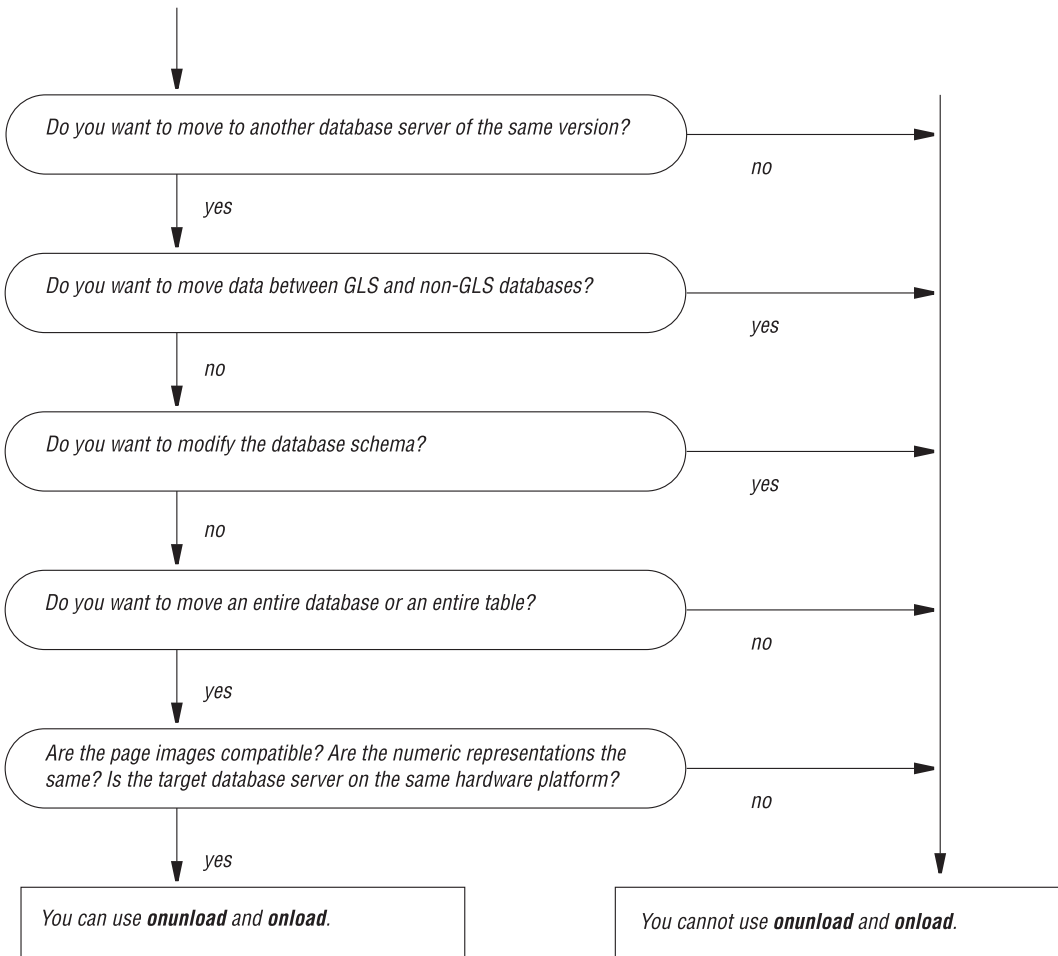


Figure 2-1. Criteria for Using the **onunload** and **onload** Utilities

### Constraints on Using **onunload** and **onload**

The **onunload** and **onload** utilities are the fastest way to unload and load data, but you can use them only when all the following criteria are the same for the source and target computers:

- Page size
- Representation of numeric data
- Byte alignment for structures and unions
- Informix database server version

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 two kilobytes on some UNIX systems and four kilobytes on Windows.

You can use the **onunload** and **onload** utilities to unload from and load data into the following database servers between computers that have the same operating system:

- Dynamic Server 9.40, 9.30, or 9.2x (only databases that do not contain extended data types)
- Dynamic Server 7.3x or 7.24
- Dynamic Server, Linux Edition 7.3x
- Workgroup Edition 7.3x or 7.24

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.

### Restrictions on Using **onunload** and **onload**

The **onunload** and **onload** utilities have the following restrictions:

Do not use **onunload** and **onload** to move data between two Dynamic Server 9.40, 9.30, or 9.2x databases if either database contains an extended data type.

Use the HPL instead to move the data.

#### Global Language Support

- You cannot use **onunload** and **onload** to move data between non-GLS and GLS locales.

#### End of Global Language Support

You can use **onunload** and **onload** to move data between databases if the NLS and GLS locales are identical. For example, 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** and **onunload** report errors. However, if both the NLS and GLS tables were created with the same French locale, **onload** and **onunload** would work.

The tape that **onload** reads contains binary data that is stored in disk-page-sized units. For this reason, 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 following characteristics:

- The same page size
- The same representation of numeric data
- The same byte alignment for structures and unions

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 last instead of first or different float-type representations), the contents of the data page could be misinterpreted.

**Important:** You cannot use the **onload** and **onunload** utilities to move data from one version of a database server to another. You also cannot use these utilities to move data between different types of database servers.

For additional constraints and restrictions, see Chapter 13, “The onunload and onload Utilities,” on page 13-1.

## The dbexport and dbimport Utilities

If you cannot use **onunload** and **onload**, you have three other methods to choose from:

- The **dbload** utility (to load data)
- The **dbexport** and **dbimport** utilities
- The UNLOAD and LOAD SQL statements

All these methods enable you to modify the database schema. The **dbexport** and **dbimport** utilities provide some flexibility, but you must move an entire database.

If you cannot use **onunload** and **onload** to export and import data, you can unload your data to text files. You can use the **dbexport** utility to unload data to tape from any of the following database servers:

- Dynamic Server 9.40, 9.30, 9.2x, 7.3x, or 7.24
- Dynamic Server, Linux Edition 7.3x
- Workgroup Edition 7.3x or 7.24
- SE

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 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 Informix environment. You can edit the schema file to modify the database that **dbimport** creates. The **dbexport** utility supports Dynamic Server 9.40, 9.30, and 9.2x data types.

### Destination Options

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. 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 **dbimport** utility supports new data types in Dynamic Server 9.40, 9.30, and 9.2x.

### Location Options

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 located on tape

### Database Server Options

The **dbimport** utility supports the following options for a new 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.

## LOAD, UNLOAD, and dbload

The LOAD statement is moderately fast and easy to use, but it can only accept 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**. Figure 2-2 summarizes questions to help you choose among these methods.

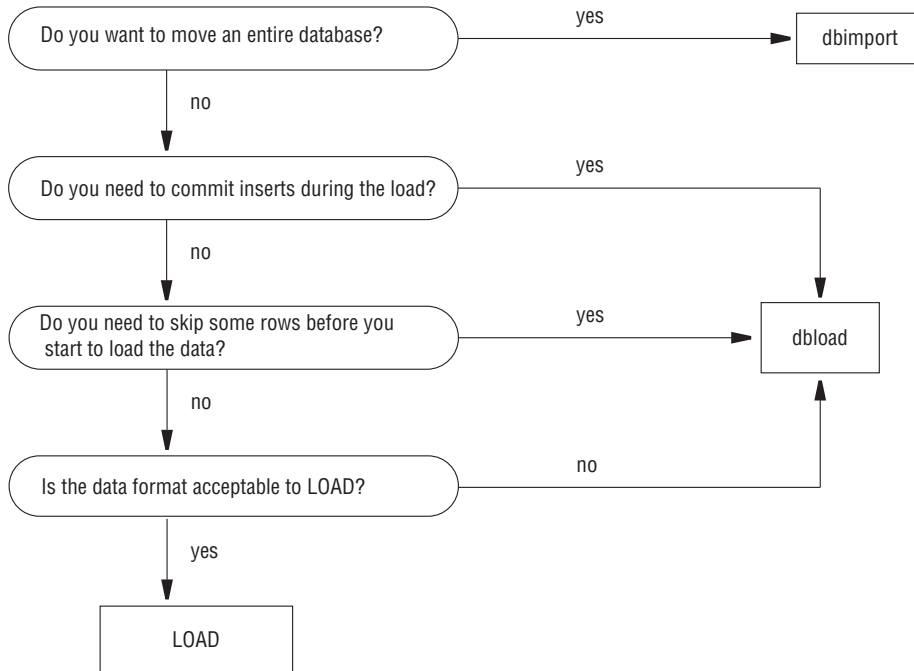


Figure 2-2. Choosing Among *dbimport*, *dbload*, and *LOAD*

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.

## The **dbschema** Utility

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

## Guidelines for Using **dbschema**

### Global Language Support

When the GLS environment variables are set correctly, as the *IBM Informix: GLS User's Guide* describes, **dbschema** can handle foreign characters in Dynamic Server 9.40, 9.30, 9.2x, 7.3x, or 7.24; Dynamic Server, Linux Edition 7.3x; or Workgroup Edition 7.3x or 7.24.

### End of Global Language Support

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

## DB-Access Input from **dbschema** Output

You can use the **dbschema** utility to get the schema of a database and redirect the **dbschema** output to a file. Later, you can feed this file to DB-Access to re-create the database.

## Object Modes and Violation Detection

The **dbschema** output supports object modes and violation detection, as follows:

- The 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 output shows the object mode of objects that are in the disabled state. These objects can be constraints, triggers, or indexes.
- The output shows the object mode of objects that are in the filtering state. These objects can be constraints or unique indexes.

- The output shows 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*.

## External Tables

You can use external tables to unload and load data when no other data migration tool is available.

For detailed information on external tables, see the description of the CREATE EXTERNAL TABLE statement in the *IBM Informix: Guide to SQL Syntax*.

## The High-Performance Loader

The High-Performance Loader (HPL) utility uses parallel processing to perform fast data loading and unloading. The HPL is available with the following database servers:

- Dynamic Server 9.40, 9.30, 9.2x, 7.3x, and 7.24
- Workgroup Edition 7.3x and 7.24

The HPL requires significant preparation time but it is fast. Use the HPL for large migration jobs. The HPL can load data from any ASCII or COBOL file that meets certain format requirements.

You can use the HPL to load from large ASCII or COBOL databases. COBOL is supported up to Version 7.3.

## HPL Tools

In addition to the advantage of speed, 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

## Performance Advantage of the HPL

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. The user, however, must invest significant preparation time before using the HPL, and the HPL program has a



significant start-up time. Therefore, 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.

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

## Nonlogging Raw Tables (IDS 9.x)

You can use nonlogging raw tables in a logging database to speed up the initial loading and validation of data in Dynamic Server 9.40, 9.30, or 9.2x, which creates standard tables that use logging by default. 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.

To create a nonlogging table for loading, you can use the CREATE RAW TABLE statement, or you can use the ALTER TABLE statement to change the table type from STANDARD to RAW. Tables of type RAW do not allow indexes or referential constraints, so the initial loading is faster than with tables of type STANDARD. After the loading of a raw table is complete, you can change it 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*.

## Movement of TEXT and BYTE Data

An Informix database server scans TEXT and BYTE data into an existing table in one of the following ways.

- SQL LOAD statement

- The **dbload** utility
- Informix ESQL/C program
- The HPL
- External table

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

---

## Moving Data Between Computers and Dbspaces

This section discusses moving data between different computers and importing data from non-Informix environments. 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 it lets you 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 into one of the following database servers:

- Dynamic Server 9.40, 9.30, 9.2x, 7.3x, or 7.24
- Dynamic Server, Linux Edition 7.3x
- Workgroup Edition 7.3x or 7.24
- SE

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

IBM Informix Enterprise Gateway with DRDA lets you query databases that conform to the DRDA protocol published by IBM. You can use this Gateway product to query a DRDA database and then insert the results into an Informix database. For example, to import data, execute a SELECT statement to select data from the non-Informix database and then execute an INSERT statement to insert data into the Informix database. For more information,

refer to the *IBM Informix: Enterprise Gateway with DRDA User Manual*.

**UNIX/Linux Only**

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 such as SYBASE SQL Server 10 and ORACLE7 Server. For more information, refer to the *IBM Informix: Enterprise Gateway Manager User Manual*.

**End of UNIX/Linux Only**



---

## **Part 2. Migration to a Later Version of a Database Server**



---

# Chapter 3. Preparing for Migration

Preparing for Migration. . . . .	3-1
Migration Guidelines . . . . .	3-2
Changes in Dynamic Server 9.40, 9.30, and 9.2x. . . . .	3-2
Environment Variables . . . . .	3-3
Configuration Parameters . . . . .	3-4
SQL Reserved Words . . . . .	3-8
System Catalog and sysmaster Changes . . . . .	3-9
Feature Changes. . . . .	3-14
New Features. . . . .	3-18
Storage-Manager Validation and Installation . . . . .	3-27
Migrating Between 32-bit and 64-Bit Database Servers . . . . .	3-28

---

## In This Chapter

This chapter contains the information you need to know before migrating to Dynamic Server 9.40.

You can convert to Dynamic Server 9.40 from the following database servers:

- Dynamic Server 9.30
- Dynamic Server 9.21
- Dynamic Server 9.20
- Dynamic Server 7.31
- Dynamic Server 7.30
- Dynamic Server 7.24 (without Enterprise Replication)

If necessary, you can revert to the database server that you converted from.

**Important:** To convert from another database server, including Universal Server 9.14, you need to convert to an intermediate database server first. For information on which intermediate database servers to use, see “Database Server Migration Paths” on page 1-7.

---

## Preparing for Migration

This section contains:

- Details about the diagnostic information that you need to gather before you upgrade to another version of Dynamic Server
- General guidelines that you must follow when migrating between database servers.

- Information on new features that might affect migration.

## Migration Guidelines

Observe the following guidelines when you migrate to Dynamic Server 9.40:

- Check the release notes for information about the correct operating-system release and any patches that you need for successful installation and operation of the database server.

The release notes are in one of the following directories:

- The **Informix** folder on Windows

To display the release notes, choose **Start > Programs > Informix Dynamic Server 9.40 > Release Notes** from the taskbar.

- **\$INFORMIXDIR/release/en\_us/0333** on UNIX or Linux

- On UNIX or Linux, retain both versions of the IBM Informix product software on disk, if you have enough disk resources. You cannot retain both versions of the IBM Informix product software on disk on Windows.
- Retain the installation media from both versions of the IBM Informix product software.
- Before you convert to the target database server from the source database server, make sure that no open transactions exist in the source database server.

Fast recovery will fail when rolling back open transactions during the conversion. For information about how to close the source database server properly prior to conversion, see “Close All Transactions and Shut Down the Source Database Server” on page 5-6.

- **Important:** Before migration, you must perform a level-0 backup of all storage spaces with the source database server. After you complete the migration, perform another level-0 backup with Dynamic Server 9.40. You need to make a backup with each database server because you can restore a backup only with the same version of the database server on which you did the backup.
- It is recommended that you use a test instance of your database server to test the new version of the database server.
- Verify storage-manager validation for the target database server.

For details, see “Storage-Manager Validation and Installation” on page 3-27.

For additional installation information and guidelines, refer to your *IBM Informix: Installation Guide* and your *IBM Informix: Getting Started Guide*.

## Changes in Dynamic Server 9.40, 9.30, and 9.2x

This section describes changes in Dynamic Server 9.40:

- New and changed environment variables
- New and changed configuration parameters



- New SQL reserved words
- System catalog and **sysmaster** changes
- Feature changes
- New features

### Environment Variables

Dynamic Server 9.40, 9.30, 9.21, and 9.20 introduced new environment variables that might 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: Administrator's Guide*.

Table 3-1 lists the new environment variables in Dynamic Server 9.40, 9.30, 9.21, and 9.20.

*Table 3-1. New Environment Variables*

Version	Environment Variable	Description
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	AFDEBUG	Forces the database system to hang when a failure occurs.
9.21	JAR_TEMP_PATH	Specifies a non-default local file system location for temporary <b>.jar</b> files of the Java virtual machine.
9.21	JAVA_COMPILER	Disables JIT compilation.

Table 3-1. New Environment Variables (continued)

Version	Environment Variable	Description
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.

### Configuration Parameters

Dynamic Server 9.40, 9.30, 9.21, and 9.20 introduced new configuration parameters that might affect your installation. You might also need to adjust the values of existing parameters. For more information on configuration parameters, see the *IBM Informix: Administrator's Reference* and your *IBM Informix: Administrator's Guide*.

Table 3-2 lists the new configuration parameters in Dynamic Server 9.40, 9.30, 9.21, and 9.20. All parameters are located in the ONCONFIG file, unless otherwise noted.

Table 3-2. New Configuration Parameters

Version	New Configuration Parameter	Description
9.40	ASF_SOCTCP_BACKLOG	Specifies the number of connections queuing when using Sockets (SOCTCP).
9.40	CDR_DBSPACE	Defines the default dbspace for the Enterprise Replication <b>syscdr</b> 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.

*Table 3-2. New Configuration Parameters (continued)*

Version	New Configuration Parameter	Description
9.40	ENCRYPT_CIPHER	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 <b>plconfig</b> 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 <b>plconfig</b> file.
9.40	PLOG_OVERFLOW_PATH	Sets the location of the temporary space to extend the physical log during fast recovery.
9.30	CDR_QHDR_DBSPACE	Specifies the dbspace Enterprise Replication uses to store the transaction record headers spooled from the send and receive queues.
9.30	CDR_QDATA_SBSpace	Specifies the sbpace Enterprise Replication uses to store spooled transaction row data.
9.30	DEF_TABLE_LOCKMODE	Sets the lock mode to row for every newly created table.
9.30	DYNAMIC_LOGS	Controls dynamic log allocation. This parameter is not in the default ONCONFIG file.
9.30	SBSPACETEMP	Specifies the name of the default temporary sbpace for storing temporary smart large objects.
9.21	DS_HASHSIZE	Specifies the number of hash buckets in the data-distribution cache for statistics generated by the UPDATE STATISTICS statement.

*Table 3-2. New Configuration Parameters (continued)*

Version	New Configuration Parameter	Description
9.21	DS_POOLSIZE	Specifies the maximum number of entries in each hash bucket for statistics generated by the UPDATE STATISTICS statement.
9.21	STMT_CACHE_HITS	Specifies the number of references to a statement before it is fully inserted in the SQL statement cache.
9.21	STMT_CACHE_NOLIMIT	Controls whether to insert qualified statements into the SQL statement cache after its size is greater than the STMT_CACHE_SIZE value.
9.21	STMT_CACHE_NUMPOOL	Specifies the number of memory pools for the SQL statement cache.
9.20	ALLOW_NEWLINE	Controls whether the database server allows the newline character (\n) in a quoted string.
9.20	DD_HASHMAX	Specifies the maximum number of tables in each hash bucket in the data-dictionary cache.
9.20	DD_HASHSIZE	Specifies the number of hash buckets or lists in the datadictionary cache.
9.20	OPT_GOAL	Specifies optimization goals.
9.20	PC_HASHSIZE	Specifies the number of hash buckets in the caches that the database server uses.
9.20	PC_POOLSIZE	Specifies the maximum number of entries in several memory caches that the database server uses.
9.20	STMT_CACHE	Determines whether the database server uses the SQL statement cache.
9.20	STMT_CACHE_SIZE	Specifies the size of the SQL statement cache.
9.20	SYSSBSPACENAME	Specifies the name of the sbSPACE for statistics collected by the UPDATE STATISTICS statement for certain user-defined data types.

Table 3-3 lists the configuration parameters that were altered or removed in Dynamic Server 9.40, and 9.30. No changes were made to configuration parameters in Dynamic Server 9.21 and 9.20.

*Table 3-3. Altered Configuration Parameters*

Version	Altered Configuration Parameter	Description of Change
9.40	ALARMPROGRAM	Can be set to the <b>alarmprogram.sh</b> 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	The default value is 32 KB.
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.
9.40	LRU_MIN_DIRTY	Can accept a value of type INTEGER or FLOAT.
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	The default value is 32 KB.
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	Removed; 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	DRAUTO	Removed; you must now perform the transition from HDR secondary to standard mode manually.
9.30	JDKVERSION	New default value of 1.3.
9.30	JVPJAVAHOME	New default value of <b>/usr/informix/extend/krakatoa</b> .
9.30	JVPJAVALIB	New default value that is platform dependent.

Table 3-3. Altered Configuration Parameters (continued)

Version	Altered Configuration Parameter	Description of Change
	JVPJAVAVM	New default value that is platform dependent.
	LBU_PRESERVE	Removed; configured an obsolete utility.
	LOGSMAX	Removed.
	LTXHWM	Removed from the default ONCONFIG file, but is still valid.
	LTXEHWM	Removed from the default ONCONFIG file, but is still valid.
	NOAGE	Removed; superseded by the VPCLASS configuration parameter.
	NUMAIOVPS	Removed; superseded by the VPCLASS configuration parameter.
	NUMCPUVPS	Removed; superseded by the VPCLASS configuration parameter.

### SQL Reserved Words

Dynamic Server 9.40, 9.30, 9.21, and 9.20 support 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 list below shows new SQL reserved words. For more information about SQL reserved words, see the *IBM Informix: Guide to SQL Syntax*.

#### Version 9.40

- COLLATION
- CROSS
- FULL
- INSTEAD
- RESTART
- RIGHT

#### Version 9.30

- AVOID\_EXECUTE
- AVOID\_SUBQF
- USE\_SUBQF

**Version 9.21**

- RAW
- STANDARD

**Version 9.20**

- AGGREGATE
- CACHE
- COSTFUNC
- ITEM
- NAME
- REF
- SELCONST

**Version 9.20 and 7.31**

- INNER
- JOIN
- LEFT
- LOCKS
- RETAIN

**Version 9.20 and 7.30**

- ALL\_ROWS
- CASE
- CRCOLS
- DECODE
- FIRST\_ROWS
- MEMORY\_RESIDENT
- NON\_RESIDENT
- NVL
- REPLICATION
- SUBSTR
- SUBSTRING

**System Catalog and sysmaster Changes**

The system catalog tables and **sysmaster** database are different from those for database servers earlier than Dynamic Server 9.20, which includes changes to some column widths, data types, and treatment of null values. Also, some tables have additional columns, and some tables were added or deleted.

**Remote Queries on System Catalog Tables Between 7.x and 9.x:** Certain system catalog tables use data types in Dynamic Server 9.x that are not supported in Dynamic Server 7.x. Remote queries that issue a `SELECT *` statement on these system catalog tables from Dynamic Server 7.x to Dynamic Server 9.x will fail.

For example, the following queries that originate on a Dynamic Server 7.x fail when executed on Dynamic Server 9.x:

```
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:** In Version 7.x, **sysindexes** is a table. In Dynamic Server 9.40, 9.30, and 9.2x, **sysindexes** is a view. The `ALTER TABLE` statement fails for **sysindexes** because this statement is not valid for altering a view.

**Column-Width Changes:** 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:

<b>sysaggregates</b>	<b>sysfragauth</b>	<b>sysroutinelangs</b>
<b>sysams</b>	<b>sysfragments</b>	<b>syssynonyms</b>
<b>sysattrtypes</b>	<b>sysindexes</b>	<b>syssyntable</b>
<b>sysblobs</b>	<b>sysindices</b>	<b>systabamdata</b>
<b>syscasts</b>	<b>sysobjstate</b>	<b>systables</b>
<b>syscolattrs</b>	<b>sysopclasses</b>	<b>systracemsgs</b>
<b>syscolumns</b>	<b>sysopclstr</b>	<b>systriggers</b>
<b>sysconstraints</b>	<b>sysprocedures</b>	<b>sysxtdtypes</b>
<b>sysdomains</b>		

Identifiers changed from `CHAR(18)` to `CHAR(128)` in the following **sysmaster** database tables:



arc_dbSPACE	syscrtadt	syslocks
arc_dbSPACE_set	sysdatabases	sysopendb
arc_phys_dev	sysdblocale	sysprc
arc_rep_table	sysdbspaces	sysproccache
arc_replicate	sysdbspartn	sysptprof
arc_server	sysdbstab	sysssdblock
arc_version	sysdic	sysssqlcurall
arc_vset	sysdiccache	sysssqlcurses
arc_vset_view	sysdistcache	sysssqlstat
flags_text	sysdsc	systabnames
syscfgtab	sysextents	sysstrans
sysconfig	sysextspaces	systxptab

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	sysindices	sysssynonyms
sysams	syslangauth	sysssyntable
syscasts	sysobjstate	systabauth
syscolauth	sysopclasses	systables
sysconstraints	sysopclstr	sysstriggers
sysdomains	sysprocauth	sysusers
sysfragauth	sysprocedures	sysxtdtypeauth
sysindexes	sysroleauth	sysxtdtypes

The following **sysmaster** database tables changed:

sysaudit	sysdiccache	sysrstcb
sysdatabases	sysdistcache	syssscb1st
sysdbspaces	sysdsc	syssessions
sysdbspartn	sysextspaces	systabnames
sysdbstab	sysprc	sysuserthreads
sysdic	sysproccache	

Columns that include pathnames or other values changed from CHAR(128) to CHAR(256) in the following **sysmaster** database tables:

sysadinfo	syscrtadt	sysmchktab
syschktab	sysdrdb	
syschunks	sysdri	

The path for a physical device changed from CHAR(128) to CHAR(260) in the following **sysmaster** database table:

## **arc\_phys\_dev**

Columns widths changed from CHAR(20) to CHAR(128) for longer object names in the following **sysmaster** database tables:

### **sysdrclb**

### **sysdrls**

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 **systabauth** system catalog table is now CHAR(9) instead of CHAR(8). The 9th character indicates the Under privilege.

**Data Type Changes:** The preceding section on column-width changes lists columns that have changed from the CHAR data type to the VARCHAR data type.

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**

### **sysssdblock**

The CHAR data type changed to the STAT data type in the following system catalog table:

### **sysdistrib**

**Changes in Treatment of Null Values:** No nulls allowed changed to nulls allowed for some columns in the following **sysmaster** database tables:

### **arc\_ae\_view**

### **arc\_db\_file\_view**

### **arc\_directory\_view**

### **arc\_file\_copy\_view**

### **arc\_file\_view**

### **arc\_pendreq\_view**

### **arc\_req\_vset\_view**

### **arc\_request\_view**

### **arc\_save\_set\_view**

### **arc\_vol\_lock\_view**

### **arc\_volume\_view**

### **arc\_vset\_user\_view**

### **arc\_vset\_view**

**Columns Added:** One or more columns have been added to the following system catalog tables:

sysams	sysprocedures
sysdistrib	sysroutinelangs

Several columns have been added to the following **sysmaster** database table:

**sysdbstab**

**Tables Added or Deleted:** The following tables have been added to the **sysmaster** database:

logmessage	syscdrctl_txn	syscdrrecv_txn
syscdrack_buf	syscdrprog	syscdrtx
syscdrack_txn	syscdrq	
syscdrctl_buf	syscdrrecv_buf	

The following table has been deleted from the **sysmaster** database:

**arc\_change\_log**

**Changes for Dynamic Server 9.30:** The following tables were been deleted from the **sysmaster** database in Version 9.30:

arc_ae_view	arc_pendreq_view	arc_server
arc_db_file_view	arc_phys_dev	arc_version
arc_dbspace	arc_rep_table	arc_vol_lock_view
arc_dbspace_set	arc_replicate	arc_volume_view
arc_directory_view	arc_req_vset_view	arc_vset
arc_file_view	arc_request_view	arc_vset_user_view
arc_file_copy_view	arc_save_set_view	arc_vset_view

**Changes for Dynamic Server 9.40:** The following new system catalog table was added:

**syssequences**

A new **collation** column has been added to the following system catalog tables:

sysconstraints	sysprocplan	sysstrigbody
sysindices		

## Feature Changes

The following feature changes are in Dynamic Server versions 9.40, 9.30, and 9.2x:

- Detached indexes
- Changes in ON-Bar commands
- Chunk format
- Libraries no longer installed in the `/usr/lib` directory
- VARCHAR column limit
- Stored procedure parameter limit
- Limitation on using routines in distributed transactions
- Case-sensitive name space
- New administration tools
- Management of the SQLHOSTS connectivity information on Windows
- Maximum index key size difference between 9.x and 7.3

**Detached Indexes:** By default, all new indexes that the CREATE INDEX statement creates in Dynamic Server 9.40, 9.30, or 9.2x are detached and stored in separate tablespaces from the data. Indexes created in Version 7.x are attached until you rebuild them.

You cannot revert detached indexes to Version 7.x. To enable reversion to Version 7.x, retain the Version 7.x attached index behavior by setting the environment variable **DEFAULT\_ATTACH** in the application environment. You can attach only B-tree indexes that are nonfragmented and that are on nonfragmented tables (the Version 7.x behavior). All other indexes, including extensibility related indexes such as R-trees and UDT indexes, must be detached.

**Changes in ON-Bar Commands:** Between Dynamic Server 7.x and Dynamic Server 9.x, the syntax of the following ON-Bar commands changed:

Dynamic Server 7.x:	<b>onbar -r -v</b>
Dynamic Server 9.x:	<b>onbar -v</b>
Dynamic Server 7.x:	<b>onbar -l</b>
Dynamic Server 9.x:	<b>onbar -b -l</b>

**Chunk Format:** Dynamic Server 9.40 introduced for the following chunk features that cannot be reverted:

- Chunks greater than 2 GB
- Chunk offsets greater than 2 GB
- Greater than 2047 chunks

To facilitate backward compatibility, these features are not enabled by default. To enable these features, use the **onmode -BC** command.

The **onmode -BC** command has the following options that allow you to stage your migration to these features:

- **onmode -BC 1**

Allows chunks and offsets of greater than 2 GB to be created, but dbspaces and blobspaces that do not contain large chunks remain in the Version 9.3 format. After a chunk larger than 2 GB is added to a dbspace or blobspace, then all chunks added or altered in that dbspace or blobspace are in the new format. To revert, drop any dbspaces or blobspaces using the new chunk features.

- **onmode -BC 2**

Enables the Version 9.4 chunk features for all dbspaces and blobspaces. Any chunk or offset added or modified has the new format. All page writes are in the new format, including those to chunks created prior to Version 9.4. Reversion is no longer possible.

**Libraries No Longer Installed in the /usr/lib Directory:** Dynamic Server 9.40 does not include the following libraries in the **/usr/lib** directory on UNIX:

- Optical storage manager library

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*.

- High-Performance Loader custom-code shared library

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 9.40, the default location of this library is **\$INFORMIXDIR/lib/ipldd09a.so**. For more information, see the *IBM Informix: High-Performance Loader User's Guide*.

**VARCHAR Column Limit:** If you are migrating from Dynamic Server 7.3x, the number of VARCHAR columns per table for Dynamic Server 9.40, 9.30, or 9.2x has been reduced from 231 to 195, within a row size of 32,762 bytes and based on a page size of 2 K. Before you upgrade to Dynamic Server 9.40, ensure that your tables obey this restriction.

**Stored Procedure Parameter Limit:** Dynamic Server 9.40, 9.30, or 9.2x have a limit of 341 parameters for each stored procedure.

**Limitation on Using Routines in Distributed Transactions:** User-defined routines and built-in routines, such as **round()**, cannot be used in distributed

transactions involving Dynamic Server 9.2x or 9.30. For Dynamic Server 9.40, routines in distributed transactions are only supported among Version 9.40 database servers.

**Case-Sensitive Name Space:** If you have case-insensitive schemas, you might need to revise them because Dynamic Server 9.40, 9.30, or 9.2x have a case-sensitive name space. This can affect the resolution of blobspaces and SPL names.

**New Administration Tools:** Dynamic Server 9.40, 9.30, and 9.21 do not support IECC. The IECC functionality for these database servers comes from the following sources:

- Server Studio JE
- IBM Informix Server Administrator (ISA)

*Server Studio JE:* Included in the IBM Informix Dynamic Server bundle is a CD containing Server Studio JE (Java Edition). Server Studio JE is the result of an extensive collaboration between IBM and Advanced Global Systems LTD (AGS).

Before using any Server Studio features, you must first register with AGS. When you use Server Studio to build the first connection to a database server, you will find information on how to register. Once registered, you will receive an email with a certificate to install. You must install this certificate to access Server Studio features.

Once you have registered, you can use the following features for an unlimited period of time:

- Object Explorer
- Schema Editor
- SQL Editor

In addition, during the first 30 days of use, several other modules are enabled on a trial basis. After the 30-day-trial period, the unlicensed add-on modules will not work and will generate expiration messages. To continue using the add-on modules after 30 days, you must obtain a license for those modules directly from Advanced Global Systems LTD (AGS).

Some of the add-on modules that are available inside Server Studio JE are:

- Enhanced SQL Editor
- Database Difference Analyzer
- Dependency Analyzer
- Permission Manager
- XPS 8.3 connectivity

- OnLine 5 connectivity

For a complete list of the currently available add-on modules visit the AGS website at [www.agsltd.com](http://www.agsltd.com).

*IBM Informix Server Administrator:* IBM Informix Server Administrator (ISA) is a browser-based tool that provides Web-based system administration for all Informix 9.40, 9.30, 9.2x, 8.40, 8.3x, 8.21, 7.3x, and 7.24 database servers. ISA provides access to database server command-line functions and presents the output in an easy-to-read format.

With ISA, you can use a browser to perform these common database-server administration tasks:

- Check dbspaces, database catalogs, logs, and other objects.
- Change configuration parameters temporarily or permanently.
- Manage logical and physical logs.
- Examine memory use and add free memory segments.
- Read the message log.
- Change the database server mode between online and offline and intermediate states.
- Back up and restore dbspaces and logical logs.
- Run **onstat** commands to monitor performance.
- Enter SQL statements and examine database schemas.
- Examine and manage user sessions.
- Examine and manage Virtual Processors (VPs).
- Add and remove storage spaces, including spaces that are unique to Dynamic Server database servers.

The database server CD-ROM distributed with your product includes ISA. For information on how to install ISA, see the following file on the CD-ROM.

Operating System	File
UNIX or Linux	/SVR_ADM/README
Windows	\SVR_ADM\readme.txt

#### **Management of the SQLHOSTS Connectivity Information on Windows:**

On Windows, IBM Informix products store the SQLHOSTS information in registry keys. Windows provides the **regedt32** utility, which enables you to manage registry keys, but it is recommended that you do not use **regedt32**.

If you have used **Setnet32** to manage SQLHOSTS information, you can continue to do so; however, **Setnet32** does not enable you to assign a database server to a database server group.

For Enterprise Replication, use ISA.

**Maximum index key size difference between 9.x and 7.3:** Before reverting to Dynamic Server 7.3, you must drop any index with a key size that is greater than 254. Dynamic Server 9.4 supports a maximum index key size of 255 for VARCHAR and NVARCHAR types. However, Dynamic Server 7.3 only supports a maximum index key size of 254.

## **New Features**

This section lists new features in Dynamic Server 9.40, 9.30, 9.21, and 9.20. Read your release notes and documentation notes for late-breaking information on new features.

**New Features in Dynamic Server 9.40:** Dynamic Server 9.40 introduces the following new features:

- Security enhancements
  - Encryption for networked transactions
  - Authentication
- Database server usability enhancements
  - Chunks up to 4 TB
  - Chunk offsets up to 4 TB
  - New chunk limit of 32,766
  - Greater than 2 GB file size
  - The default size of the TAPEBLK and LTAPEBLK configuration parameters is 32 KB
  - Chunks can be added when the root chunk is full by storing metadata in extended chunk reserve pages allocated from non-root chunks in the root dbspace
  - Buffer cleaning can be tuned by setting the LRU configuration parameters to a value of type FLOAT
  - Event alarms can be set with a modifiable shell script, **alarmprogram.sh**
  - Ability to specify up to 32 database server aliases with the DBSERVERALIASES configuration parameter
  - During fast recovery, the physical log space is extended if the physical log overflows
  - Microsoft Transaction Server/XA support
- Performance enhancements
  - PDQ is enabled for hold cursors
  - B-tree cleaner improves transaction processing for logged databases when rows are deleted from a table with indexes
  - Improved priority management for the buffer manager



- Reliability, availability, and supportability features
  - Ability to monitor queries dynamically using the **onmode -Y** command
  - Print the session control block address with the **onstat -g ses** command
  - Display the setting and values of environment variables with the **onstat -g env** command
  - Ability to specify the number of pages to print, whether to print just the page headers, and to print pages from chunks that are online with the **oncheck** utility
  - Display the types and values of host variables in SQL statements, show the stored procedure stack, and show the current SQL statement in a stored procedure using the **onstat -g sql** command
- SQL enhancements
  - Ability to specify an expression or a column in the ORDER BY clause of a SELECT statement that is not listed in the projection of the SELECT statement
  - Change the collation used by the session for comparisons and sorts on NCHAR and NVARCHAR objects
  - Functional indexes can be created on 102 parameters for C UDRs, or 341 parameters for SPL and Java UDRs
  - The LVARCHAR(*m*) data type can be set to sizes larger than 2 KB
  - INSTEAD OF triggers on views
  - Sequence objects
  - Additional in-place alters on clauses of the ALTER TABLE statement
  - DESCRIBE INPUT and DESCRIBE OUTPUT statements to return information about multiple input and output parameters
  - The SET EXPLAIN statement AVOID EXECUTE option displays the query plan without executing the query
  - ANSI join syntax support for cross, right outer, and full outer joins
  - Unions allowed in subqueries of SELECT statements
  - USETABLENAME environment variable to invalidate synonyms in ALTER TABLE and DROP TABLE statements
- GLS enhancements
  - Session-level non-default collation
  - Unicode and Unicode collation support
  - Support for the Chinese locale 18030-2000
- Enterprise Replication enhancements
  - Encrypted transactions implemented with configuration parameters
  - Support for replicating row data types and collection data types
  - Faster queue recovery

- Replication during queue recovery
- Support for large transactions up to 4 TB
- Improved availability by combining High-Availability Data Replication (HDR) and Enterprise Replication
- Dynamic log file
- A **brief** option for the **cdr list replicate** command to display a summary of participants for all replicates
- The **cdr remove** command to remove Enterprise Replication from an HDR server
- The CDR\_QDATA\_SBSPACE configuration parameter now accepts up to 32 sbpace names. Enterprise Replication uses all sbspaces in round-robin order.
- The CDR\_DBSPACE configuration parameter to specify the dbspace of the **syscdr** table
- The CDR\_ENV configuration parameter to set Enterprise Replication environment variables
- The CDR\_LOGDELTA environment variable to determine when the send and receive queues are spooled to disk
- The CDR\_PERFLOG environment variable to enable queue tracing
- The CDR\_ROUTER environment variable to disables intermediate acknowledgements of transactions in hierarchical topologies
- The CDR\_RMSCALEFACT environment variable to set the number of DataSync threads started for each CPU VP
- Extensibility enhancements
  - Ability to create user-defined selectivity functions and to calculate the cost of using an R-tree index so as to allow the optimizer to make accurate decisions about which index to use
  - HDR support for extended data types and UDRs
  - Using an iterator function in the FROM clause of a SELECT statement and returning the result set using a virtual-table interface
  - Naming the return parameters of a UDR
  - Using multiple OUT parameters and statement local variables
- DataBlade API enhancements
  - Ability to return the value of the current database server locale with the **mi\_get\_db\_locale()** function
  - Ability to return the ID of the current transaction with the **mi\_get\_transaction\_id()** function
  - Ability to change the size of an existing memory block with the **mi\_realloc()** function

- Ability to determine whether the current user stack has the specified amount of free space with the **mi\_stack\_limit()** function
- Ability to execute operating system commands in a separate thread with the **mi\_system()** function
- Stream support for files larger than 2 GB
- High-Performance Loader enhancements
  - Use the full capacity of a storage device
  - Ability to set the location of the custom-code shared library file with the **HPL\_DYNAMIC\_LIB\_PATH** configuration parameter
  - Ability to use a different length for data in the input and output arguments of custom-code functions by setting the **HPLAPIVERSION** configuration parameter
- Backup and restore enhancements
  - Renaming chunks to a different path and offset during a cold restore
  - Using the full capacity of a storage device with **ontape**
- Installation enhancements
  - Components previously installed in **/usr/lib** are now installed in **\$INFORMIXDIR/lib**.
  - On UNIX, the installation program prompts the user to avoid overwriting existing client or GLS files that are more recent than those included with the database server.

**New Features in Dynamic Server 9.30:** Dynamic Server 9.30 has the following new features:

- The ability to display the maximum number of connections
- DataBlade API enhancements
  - New **PER\_STMT\_EXEC** and **PER\_STMT\_PREP** memory durations
  - The ability to use **mi\_lo** routines without a connection
  - New **mi\_collection\_card()** function to return the cardinality for a collection (number of items in the collection)
  - Access to files on a client computer one buffer at a time
  - New **mi\_transaction\_state()** function to return the current transaction state (none, implicit, or explicit)
- Enterprise Replication enhancements
  - Improvements to parallel processing
  - External Enterprise Replication conversion
  - SERIAL column primary keys
  - Replicate sets

Pre-Version 9.30 replicate groups are not supported in Version 9.30. Before you migrate to Version 9.30, you must remove any replicate

groups. For instructions, see Chapter 4, “Migrating to Dynamic Server 9.40 with Enterprise Replication,” on page 4-1.

- Replicating only changed columns
- Spooling changes
- In-place processing of ALTER statements to add or drop shadow columns
- Command-line changes to support new features
- Extensibility enhancements
  - New **deepcopy()** function for multirepresentational data types
  - Nearest neighbor queries in R-trees
  - Temporary sbspaces
  - Temporary smart large objects
  - Improved partitioning of user data and metadata in sbspaces
- Java Virtual Machine (JVM) 1.3 support
- Performance enhancements
  - The **onstat -g stm** option
  - The ability to display the query plan without executing the query
  - Optimizer directives for subqueries
  - Dynamic addition of logical logs
  - Performance improvement for smart large objects
- SQL enhancements
  - Configurable default lock modes
  - REVOKE...AS USER statement
  - DELETE statement that does not require the FROM keyword

**New Features in Dynamic Server 9.21:** Dynamic Server 9.21 introduced the following new features:

- Database server features
  - The **onpladm** utility for High-Performance Loader (HPL) jobs
  - SQL statement cache to store SQL statements for re-execution
  - Access to synonyms on remote 7.x database servers through DB–Access
- Extensibility features
  - C++ support for writing user-defined routines (UDRs) with fewer restrictions
  - DataBlade API functions for controlling the virtual processor environment
  - The **mi\_fp\_funcname()** function to get the SQL name of a function
- Java features

- JVM 1.2 support
- Changes in default values of Java configuration parameters
- GLS support for J/Foundation
- The **update\_jars.sql** script
- Run-time environment variables
- JVP virtual-processor classes dropped dynamically
- Support for variable-length opaque types and opaque-type send/receive, import/export, and importbin/exportbin functions
- Support for MaxConnect and new network protocols: **ontliimc** and **onsocimc**

**New Features in Dynamic Server 9.20:** Dynamic Server 9.20 introduced the following new features:

- Extensibility enhancements
  - Dynamic lock allocation
  - SQL enhancements
    - Embedded newline characters in quoted strings
    - Nested dot expressions for row types
    - Triggers on SELECT statements
  - Enhancements to smart large objects
    - Round-robin fragmentation for smart large objects
    - ALTER TABLE for smart large objects
    - Data type conversion: BYTE to BLOB and TEXT to CLOB
    - Change of sbspace characteristics (**onspaces -ch**)
    - Immediate deletion of smart large objects
  - Enhancements to collections
    - Collection constructors that use arbitrary expression elements
    - Collection-derived tables
    - Collection subqueries
  - Enhancements to row types
    - Serial types in row types
    - GRANT/REVOKE UNDER on row types
  - Enhancements to UDRs
    - Ability to write UDRs in the Java language
    - GRANT/REVOKE on UDR external languages
    - ALTER FUNCTION/PROCEDURE/ROUTINE statements
    - User-defined aggregates
  - Extensions to the DataBlade API

Ability to obtain database server environment values with the **mi\_get\_serverenv()** function

Ability to obtain database connection information with the **mi\_get\_connection\_option()** function

Ability to obtain a function descriptor from a routine identifier with the **mi\_funcdesc\_by\_typeid()** function

Ability to obtain a DATETIME or INTERVAL qualifier from a type descriptor with the **mi\_type\_qualifier()** function

Ability to access a collection subquery with the **mi\_collection\_open\_with\_options()** function

Ability to delete a smart large object immediately with the **mi\_lo\_delete\_immediate()** function

Removal of restrictions on callback extensions

- Extensions to the ON-Bar suite

**onmsync** tool for object expiration and synchronization

**onbar** tools: override (with -O option), progress feedback, restartable restore, and external backup and restore

- The **oncheck** and **onlog** utilities for R-tree indexes
- Enhancements to R-tree indexes

A support interface for bulk loading of R-tree indexes

A new SQL UDR that lets users access the bounding box of the root page of an R-tree index

- Performance improvements

- Fuzzy checkpoints
- Parallel statement-local variables (SLVs)
- SQL statement cache
- Compiled expressions
- Improvements for smart large objects

Incremental backup support

Lightweight I/O

Metadata compaction

Byte-range locking

Dirty read

- Improvements for UDRs

Expensive-function optimization

Parallel UDRs

Support for parallel UDR execution in HPL

User-defined statistics routines

- Parallel scan for Virtual-Table Interface (VTI) and Virtual-Index Interface (VII)
- Set read and write for VTI and VII
- Enterprise Replication enhancements
  - Improvements in management of storage queues
  - Faster processing for large transactions
  - Better management to avoid long transactions
  - Improvements in the log reader
  - Updates to the command-line interface
  - Hierarchical routing; direct connections no longer required
  - Support for database servers with intermittent connections
  - Reduced global catalog available
  - New CONNECT and DISCONNECT functions in the Enterprise Replication Manager
- Special features
  - Long identifiers: 128-byte identifiers and 32-byte user names
  - Ability to retain update locks

**Version 9.20 Features from Universal Server 9.14:** Dynamic Server 9.20 includes the following features that were first released in Universal Server 9.1x:

- Lightweight I/O
- Dynamic addition of locks
- Nonroot execution of virtual processors (VPs)
- EXECUTE FUNCTION in FOREACH EXECUTE statement
- Re-entrant triggers
- Limit of SET SESSION AUTHORIZATION scope to the current database
- **INFORMIXCONTIME** environment variable, which needs to be set to a large number to make the DCE CSM work (not necessary for non-CSM environments)
- TP/XA support added
- ON-Bar support for GLS
- **oncheck** enhancement to display access-method data
- **onstat** enhancement to display the access-method cache
- New interface routines added to the Virtual Table Interface and Virtual Index Interface:
  - **mi\_string \*mi\_qual\_funcname(MI\_AM\_QUAL\_DESC \*qd)**
  - **mi\_integer mi\_scan\_nprojs(MI\_AM\_SCAN\_DESC \*sd)**
  - **mi\_smallint \*mi\_scan\_projs(MI\_AM\_SCAN\_DESC \*sd)**

- **MI\_UPDATE\_STAT\_MODE**  
**mi\_tab\_update\_stat\_mode(MI\_AM\_TABLE\_DESC \*td)**
- DataBlade API functions **mi\_current\_command\_name()** , **mi\_file\_errno()** ,  
**and mi\_get\_id()**
- The **superstores\_demo** database
- The **ifxrltree.1.00** DataBlade module

**Version 9.20 Features from Dynamic Server 7.30:** Dynamic Server 9.20 also has features first released in Dynamic Server 7.30:

- ALTER FRAGMENT ATTACH/DETACH enhancements
- In-place ALTER TABLE MODIFY/DROP (for built-in data types)
- External backup and restore and restartable restore
- Performance enhancements, including new optimizer directives, select first *n* rows, SET OPTIMIZATION statement enhancements, memory-resident tables, correlated subquery enhancements, and key-first index scan
- Features for a database server on Windows
  - Multiple residency
  - Raw device support
  - High-Performance Loader (HPL)
  - ON-Bar XBSA certification
  - ON-Bar parallelism
  - Nondomain Administrator install
  - Microsoft cluster support
  - Local-connection support with named pipes
  - Three gigabytes of shared memory
- Application migration
  - UPPER, LOWER, and INITCAP functions for case-insensitive search (for built-in data types)
  - REPLACE, SUBSTR, LPAD, and RPAD functions for string manipulation (for built-in data types)
  - UNION operator in CREATE VIEW statement
  - CASE expression
  - NVL and DECODE functions
  - TO\_CHAR and TO\_DATE date-conversion functions (for built-in data types)
  - **IFX\_UPDDESC** environment variable to describe an UPDATE statement
  - EXECUTE PROCEDURE syntax to update triggering columns
  - New arguments to the **dbinfo()** function to obtain the hostname and version of the database server



- ISM to manage the storage devices and media that contain backups
- Additional information for **onsnmp** Management Information Bases
- Connectivity features
  - Greater network security available through support of third-party security services
  - Client interfaces (ESQL/C, CLI, C++, Java, JDBC, GLS) supported in a single package
  - Compatible with Data Director 3.6 and higher
- Enterprise Replication features
  - Hierarchical routing; direct connections no longer required
  - Additional support and performance enhancements for database servers with intermittent connections
  - Reduced global catalog available
  - Enhanced Global Language Support (GLS) to replicate multiple locales within a single replication environment
  - Command-line utility
  - Scripting view in the Replication Manager graphical interface
  - CONNECT and DISCONNECT functions in the Enterprise Replication application programming interface
- Optical Subsystem shared-library support
- Additional options for the **oncheck** utility
  - The **oncheck -w** option to check and print an index without placing a shared lock on the table
  - The **oncheck -R** option to check the reserved pages, physical-log pages, and logical-log pages

## Storage-Manager Validation and Installation

When you convert or revert an IBM Informix database server, the Storage Manager that you used on the old version might not be validated for the version to which you are migrating. Verify that the storage manager is validated for the target database server version and platform at the following web site: <http://www.ibm.com/software/data/informix/pubs/smv/smv.html>

If the storage manger is not validated, you need to install a certified storage manager before you perform backups with ON-Bar.

Before you convert 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 on UNIX or Linux or the **%INFORMIXDIR%\etc** directory on Windows. If you are using a different directory as **INFORMIXDIR** for the new database server, copy **sm\_versions** to the new **\$INFORMIXDIR/etc** or **%INFORMIXDIR%\etc** directory, 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 conversion. If you are migrating from Version 7.3x before 7.31.xC5 or from Version 7.24, create the **sm\_versions** file by unloading the information from the **sysutils:bar\_version** table.

When you convert to the new database server version, install the storage manager *before* you bring up the database server. That way if you have automatic log backup set up on the database server, ON-Bar can start backing up the logs when the database server comes online.

**Warning:** If you migrate ISM Version 1.0 catalogs to Version 2.0 using the **ism\_catalog** utility, the catalogs become corrupted. When the ISM 2.0 server is restarted after catalog migration, error messages occur in various logs.

Dynamic Server 9.40 uses ISM 2.2. When setting up ISM, you might need setup information about the following features:

- ISMData or ISMLogs name change
- NSRADMIN utility
- Year 2000 compliant status
- ISM installation and certification during migration

**Important:** If you are using NetWare *IPX/SPX*, it should be installed on the same computer as the ISM server.

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

---

## Migrating Between 32-bit and 64-Bit Database Servers

To migrate to or from a 64-bit version of Dynamic Server 9.40 and an earlier database server (64-bit as well as 32-bit), follow the migration process described in Chapter 5, “Converting to Dynamic Server 9.40,” on page 5-1 and Chapter 6, “Reverting from Dynamic Server 9.40,” on page 6-1.

Some additional steps would be required in the following cases:

- If you are migrating from a 32-bit version of Dynamic Server 9.40 to a 64-bit version of Dynamic Server 9.40
- If you are reverting from 64-bit version of Dynamic Server 9.40 to a 9.30, or 9.2x 32-bit version

These steps are required to update certain internal tables. These steps are documented in the platform-specific machine notes that are provided with your database server, as “Additional Documentation” on page xviii of the Introduction indicates.

All UDRs and DataBlade modules that were built in 32-bit mode need to 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 a 32-bit mode, make sure that proper size and alignment of the data structures are used in order to work correctly on a 64-bit computer after recompilation in 64-bit mode. For more information, refer to the machine notes.



---

## Chapter 4. Migrating to Dynamic Server 9.40 with Enterprise Replication

Converting to Dynamic Server 9.40 with Enterprise Replication . . . . .	4-1
Converting Replication of 9.2x User-Defined Data Types. . . . .	4-2
Reverting from Dynamic Server 9.40 with Enterprise Replication . . . . .	4-3

---

### In This Chapter

This chapter describes how to convert to and revert from Dynamic Server 9.40 if you are running Enterprise Replication.

All of the conversion and reversion operations must be performed as user **informix**.

---

### Converting to Dynamic Server 9.40 with Enterprise Replication

You can only use the following procedures to convert to Dynamic Server 9.40 from Dynamic Server 9.40, 9.30, 9.2x, or 7.31.

To convert from an earlier database server, you need to convert to an intermediate database server first. For information on which intermediate database server to use, see “Database Server Migration Paths” on page 1-7. For information on how to convert to the intermediate database server, see the Migration Guide that is included in the IBM Informix documentation for that database server.

#### To prepare for conversion to Dynamic Server 9.40 with Enterprise Replication:

1. If you have replicate groups, which Dynamic Server 9.40 does not support, remove all of them.
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 grouper does not have any pending transactions.
  - Run **onstat -g rqm** to check for queued messages.
4. Shut down Enterprise Replication with the following command:  
`cdr stop`

#### To convert to Dynamic Server 9.40 with Enterprise Replication:

1. Perform the conversion tasks described in “In This Chapter” on page 5-1, including starting Dynamic Server 9.40.
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. Back up the **syscdr** databases with the **dbschema** or the `UNLOAD` statement.
4. Make sure that no replicatable transactions occur before Enterprise Replication starts.
5. Run the conversion script, named **concdr.sh**, in the `$INFORMIXDIR/etc/conv` directory on, or **concdr.bat**, in the `%INFORMIXDIR%\etc\conv` directory on Windows:  

```
% sh concdr.sh from_version 9.4
```

Valid *from\_version* values are: 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 the following file:  
**\$INFORMIXDIR/etc/concdr.out** or **%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
```

**Warning:** After you convert to Dynamic Server 9.40 with Enterprise Replication from Dynamic Server 9.30, 9.2x, or 7.31, do not drop the **syscdr** database. If **syscdr** is dropped, then you cannot revert from Dynamic Server 9.40 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 9.2x User-Defined Data Types

Dynamic Server 9.2x has limited support for the replication of user-defined data types (UDTs). To take advantage of 9.40 UDT replication, the user-defined routines (UDRs) for a UDT must contain **streamwrite()** and **streamread()** functions. After you migrate to Dynamic Server 9.40, implement the **streamwrite()** and **streamread()** functions for any currently replicated UDTs on all database servers within the enterprise.

---

## Reverting from Dynamic Server 9.40 with Enterprise Replication

You can only use the following procedures to revert from Dynamic Server 9.40 to Dynamic Server 9.30, 9.2x, or 7.31.

The procedure reverting to Dynamic Server 9.40 or 9.30 is slightly different than the procedure for reverting to Dynamic Server 9.2x 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).

### To revert to Dynamic Server 9.30 from Dynamic Server 9.40 with Enterprise Replication:

1. Stop applications doing replicatable transactions.
2. Make sure that control and TRG send queues are empty:
  - Run **onstat -g grp** to ensure that grouper does not have any pending transactions.
  - Run **onstat -g rqm** to check for queued messages.
3. Shut down Enterprise Replication with the following command:  
`cdr stop`
4. Back up the **syscdr** databases with **dbschema** or **UNLOAD**.
5. 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 9.4 9.30
```

This script does a reversion test followed by the actual Enterprise Replication reversion.
6. If the reversion test or actual reversion fails, then check the file **\$INFORMIXDIR/etc/revtestcdr.out** or **revcdr.out**, respectively. Attempt reversion after resolving problems reported.
7. Perform database server reversion tasks, as described in “Reverting from Dynamic Server 9.40 to Versions 9.30, 9.2x, 7.3x, or 7.24” on page 6-1.
8. Start Enterprise Replication:  

```
% cdr start
```

### To revert to Dynamic Server 9.2x or 7.31 from Dynamic Server 9.40 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 replicatable transactions.
4. Make sure that control and TRG send queues are empty:

- Run **onstat -g grp** to ensure that 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 9.4 to_version
```

Valid *to\_version* values are:

```
9.21
9.20
7.31
```

This script does a reversion test followed by the actual Enterprise Replication reversion.
  8. If the reversion test or actual reversion fails, then check the file **\$INFORMIXDIR/etc/revtestcdr.out** or **revcdr.out**, respectively. Attempt reversion after resolving problems reported.
  9. Drop the Enterprise Replication sbspaces.
  10. Perform database server reversion tasks, as described in “Reverting from Dynamic Server 9.40 to Versions 9.30, 9.2x, 7.3x, or 7.24” on page 6-1.
  11. Start Enterprise Replication:  

```
% cdr start
```



---

## Chapter 5. Converting to Dynamic Server 9.40

Converting to Version 9.40 . . . . .	5-1
Check and Configure Available Space . . . . .	5-3
Save Copies of the Current Configuration Files . . . . .	5-5
Remove Outstanding In-Place Alters . . . . .	5-5
Close All Transactions and Shut Down the Source Database Server . . . . .	5-6
Check for Any Open Transactions . . . . .	5-7
Verify the Integrity of the Data . . . . .	5-8
Verify the Database Server Mode . . . . .	5-8
Disable High-Availability Data Replication . . . . .	5-8
Make a Final Backup of the Source Database Server . . . . .	5-9
Verify That the Source Database Server Is Offline . . . . .	5-9
On UNIX or Linux, Modify Kernel Parameters . . . . .	5-9
Install Dynamic Server 9.40 . . . . .	5-10
Set Environment Variables . . . . .	5-10
Upgrade the High-Performance Loader onpload Database . . . . .	5-11
Customize Configuration Files . . . . .	5-12
Add Any Communications Support Modules (UNIX/Linux) . . . . .	5-13
Install and Configure Any DataBlade Modules. . . . .	5-13
Initialize Dynamic Server 9.40 . . . . .	5-13
Monitor the Conversion Complete Status . . . . .	5-14
For ON-Bar, Rename the sm_versions.std File . . . . .	5-15
Update Statistics. . . . .	5-15
Verify the Integrity of the Data . . . . .	5-15
Make an Initial Backup of Dynamic Server 9.40 . . . . .	5-16
Tune Dynamic Server 9.40 for Performance . . . . .	5-16
Enable HDR . . . . .	5-16

---

### In This Chapter

This chapter describes the procedures for converting to Dynamic Server 9.40 from any of the following database servers:

- Dynamic Server 9.40
- Dynamic Server 9.30
- Dynamic Server 9.2x
- Dynamic Server 7.31, 7.30, or 7.24

---

### Converting to Version 9.40

When you migrate to Dynamic Server 9.40, you can install and test a database server instance with the same configuration files, environment variables, and **sqlhosts** information that you used for your source database server.

After successful conversion to Dynamic Server 9.40, you might want to modify configuration files and environment variables to take advantage of Dynamic Server 9.40 features. For more information, see the *IBM Informix: Getting Started Guide* and your *IBM Informix: Administrator's Guide*.

**Important:** Read the release notes and the machine notes for any late-breaking information.

To convert to Dynamic Server 9.40, complete the following conversion steps. The subsections below describe these steps in detail.

**To convert to Dynamic Server 9.40:**

1. If necessary, perform Enterprise Replication preconversion tasks. For more information, see Chapter 4, “Migrating to Dynamic Server 9.40 with Enterprise Replication,” on page 4-1.
2. Check and configure available space.
3. Save copies of the current configuration files.
4. Remove outstanding in-place alters.
5. Close all transactions and shut down the source database server.
6. Check for any open transactions.
7. Verify the integrity of the data.
8. Verify the database server mode.
9. If necessary, disable High-Availability Data Replication.
10. Make a final backup of the source database server.
11. Verify that the source database server is offline.

**UNIX/Linux Only**

12. On UNIX or Linux, modify kernel parameters.

**End of UNIX/Linux Only**

13. Install Dynamic Server 9.40.  
Monitor the message log, **online.log**, during the conversion for any error messages.
14. Set environment variables.
15. Upgrade the High-Performance Loader onpload database for 9.40xC3.
16. Customize configuration files.
17. Add any Communications Support Modules.
18. Install and configure any DataBlade modules.
19. Switch to user **informix** and initialize Dynamic Server 9.40.
20. Monitor the conversion complete status.

21. For ON-Bar, rename or edit the **sm\_versions.std** file.
22. Update statistics.
23. Verify the integrity of the data.
24. Make an initial backup of Dynamic Server 9.40.
25. Tune Dynamic Server 9.40 for performance.
26. If necessary, perform Enterprise Replication conversion tasks. For more information, see “Converting to Dynamic Server 9.40 with Enterprise Replication” on page 4-1.
27. If necessary, enable HDR.

Repeat this procedure for each instance of Dynamic Server 9.40 that you plan to run on the computer.

**Warning:** If a serious error occurs during the conversion, it might be necessary to return to the previous version, restore from a backup, and then correct the problem prior to restarting the conversion tasks.

---

## Check and Configure Available Space

During conversion, Dynamic Server drops and then recreates the **sysmaster** database. Depending on which version of Dynamic Server you convert from, the Version 9.40 **sysmaster** database can be significantly larger.

---

### UNIX/Linux Only

Dynamic Server 9.40 requires 3000 free pages of logical-log space (approximately 6000 kilobytes for a 2-kilobyte page size) to build the **sysmaster** database on UNIX or Linux.

---

End of UNIX/Linux Only

---

### Windows Only

Dynamic Server 9.40 requires 1500 to 3000 free pages of logical-log space (approximately 6000 kilobytes for a 4-kilobyte page size) to build the **sysmaster** database on Windows.

---

End of Windows Only

Partition header pages should not be full; key descriptors will occupy slightly more space after conversion to Dynamic Server 9.40.

The root chunk should contain at least ten percent (10%) free space when converting to Dynamic Server 9.40.

In some cases, even if the database server conversion is successful, internal conversion of some databases might fail because of insufficient space for system catalog tables. For more information, see the release notes, as “Additional Documentation” on page xviii of the Introduction indicates.

You need to add any additional free space to the system prior to the conversion. If the dbspaces are nearly full, you need to add space before you start the conversion procedure. When you initialize Dynamic Server 9.40 on the same root dbspace of the earlier database server, Dynamic Server 9.40 automatically converts the **sysmaster** database and then each database individually. For a successful conversion of each database, ensure that 2000 kilobytes 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:

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

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

Use the following 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 partdbsnm(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 dbsnm 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. Configuration files that you might have are listed in Table 5-1. Keep the copies available in case you need to use them.

While you can use an old ONCONFIG configuration files with Version 9.40, IBM recommends that you use the new Version 9.40 ONCONFIG file, or at least examine the file for new parameters.

*Table 5-1. 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 also need to 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**

---

## Remove Outstanding In-Place Alters

The database server cannot convert if there are outstanding in-place alters.

Remove any outstanding in-place alters on tables before converting by running a "dummy" UPDATE statement against each table. A dummy UPDATE statement is an SQL update command that touches rows and generates log traffic, but does not change anything. As example of a "dummy" UPDATE statement is:

```
update table set col = col [where 1=1];
```

To find outstanding partitions having an in-place-alter history, ensure that the OPTCOMPIND configuration parameter was set to 0 before the last restart of the database server. If not, restart the database server with OPTCOMPIND set to 0. Then execute the following SQL statements:

```
database sysmaster;
set isolation to dirty read;
select pg_partnum + pg_pagenum - 1 partn
  from syspaghdr, sysdbspaces a
  where pg_partnum = 1048576 * a.dbsnum + 1
     and pg_next != 0
  into temp pp with no log;
select b.dbsname database, b.tabname table
  from systabnames b, pp where partn = partnum;
```

**Note:** The syntax shown causes the database server to look for partitions having an in-place-alter history, but not necessarily outstanding in-place-alters. The syntax specifies looking for partition pages with `pg_next != 0`, so the syntax returns the same results before and after dummy updates.

---

## Close All Transactions and Shut Down the Source Database Server

Communicate to client users how long you expect the database server to be offline for the migration. Terminate all database server processes and shut down your source database server. This lets users exit and shuts down the database server gracefully. If necessary, you can perform an immediate shutdown of the database server, which does not let users save their work.

Before you convert the source database server to the target database server, make sure that no open transactions exist. Otherwise, fast recovery will fail when rolling back open transactions during the conversion.

### To let users exit and shut down the database server gracefully:

1. Execute the **onmode -sy** command to put the database server in quiescent mode.
2. Wait for all users to exit.
3. Execute the **onmode -l** command to move to the next logical log.
4. Execute the **onmode -c** to force a checkpoint.

5. Make a level-0 backup of the database server.
6. Execute the **onmode -yuk** command to shut down the system.

**To perform an immediate shutdown of the database server:**

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

---

## Check for Any Open Transactions

A shutdown procedure does not guarantee a rollback of all open transactions. To guarantee that the source database server has no open transactions prior to the conversion, you need to put the source database server in quiescent mode. Execute 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 need to execute **oninit -s** before you initialize Dynamic Server 9.40. If the source database server is not left in quiescent mode, 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: Administrator's Guide*.

---

### End of UNIX/Linux Only

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

Only after proper shutdown can you bring the new database server (Dynamic Server 9.40) through the conversion path. Any open transaction during the conversion would cause an execution failure in fast recovery.

---

## Verify the Integrity of the Data

Use the **oncheck** utility to verify the integrity of the data before you make a level-0 backup. If you find any problems with the data, fix them before you make the backup. You can verify the integrity of the reserve pages, extents, system catalog tables, data, and indexes.

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

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

Table 5-2 lists the commands that verify the data integrity.

*Table 5-2. Commands for Verifying the Data Integrity*

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

For information on **oncheck**, see the *IBM Informix: Administrator’s Guide*.

---

## Verify the Database Server Mode

Before you make a backup, execute the following command to verify that your source database server is in quiescent mode:

```
onstat -
```

The first line of the onstat output contains the status of your source database server.

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

The word Quiescent indicates that Dynamic Server is in quiescent mode.

---

## Disable High-Availability Data Replication

If you use High-Availability Data Replication (HDR), you must disable it before conversion to Dynamic Server 9.40.

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



onmode -d standard

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

**Tip:** You do not need to perform the whole conversion process on the secondary database server. For more information, see “Enable HDR” on page 5-16.

---

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

**Important:** Make a final backup for 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. This ensures that backups for the source database server are not confused with backups about to be done for the target database server. Follow the instructions regarding expiration in your storage-manager documentation.

---

## Verify That the Source Database Server Is Offline

The source database server must be offline because the target database server uses the same files. You cannot install the target database server if any of the files that it uses are active.

Check the message log to verify that you obtain the following message for offline mode.

shared memory not initialized...

---

## On UNIX or Linux, Modify Kernel Parameters

You might need to change some of the kernel parameters for your UNIX or Linux operating system before you install Dynamic Server 9.40. To reconfigure the operating system, follow the directions in the machine notes file included on your database server distribution media, and the kernel-configuration instructions for your operating system.

---

## Install Dynamic Server 9.40

On UNIX or Linux, you must be logged in as user **root** to install Dynamic Server 9.40. On Windows, you must be a member of the **Informix-Admin** group. Set the **INFORMIXDIR** environment variable to the directory where you plan to install the database server.

**Warning:** If you install the target database server in the same directory where the source database server resided, the installation script overwrites the older files. If you want to preserve your source database server files, you must install the target database server in a different directory. If you install the target database server in a different directory, you need to change the value of the **INFORMIXDIR** environment variable, or the older version of the database server will start up when you reboot.

Before you overwrite the source database server, you must take the following precautions:

- If you do not have the original media for the source database server, back up the **INFORMIXDIR** directory before you install Dynamic Server 9.40.
- Copy the configuration and **sqlhosts** files from the **etc** directory of **INFORMIXDIR** to another location in the file system.

To install and configure Dynamic Server 9.40, follow the directions in your *IBM Informix: Installation Guide* and in your *IBM Informix: Administrator's Guide*. The installation script installs Dynamic Server 9.40 into the **INFORMIXDIR** directory specified for user **root** on UNIX or Linux or for the **Informix-Admin** group on Windows.

**Important:** Monitor the database server message log, **online.log**, during the conversion for any error messages. If you see an error message, resolve the error condition before you continue the conversion procedure.

When you finish installation as user **root** and are ready to initialize the server, you must switch to user **informix**.

---

## Set Environment Variables

After you install Dynamic Server 9.40, verify that the following environment variables are set to the correct values:

**INFORMIXSERVER**  
**ONCONFIG**

**PATH**  
**INFORMIXSQLHOSTS** (if used)

**Important:** 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.

---

## Upgrade the High-Performance Loader **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 in previous versions. The following table shows the relationship between the database server version and the **onpload** database schema version.

Server Version	onpload Database Version
Pre 7.3	0 or not available
7.31	1
9.2, 9.3, 9.40.xC1, 9.40.xC2	2
9.40.xC3 and later	3

When you upgrade, you must also upgrade the **onpload** database. To upgrade the **onpload** database, perform one of the following tasks:

- Run the **conpload.sh** or **conpload.bat** script, either during server conversion or manually.
- Set the new environment variable **IFX\_ONPLOAD\_AUTO\_UPGRADE** to 1 for the upgrade to happen automatically the first time you invoke 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 invoke an HPL utility before upgrading the **onpload** database, then you receive an error stating that the **onpload** database must be converted.

If you are upgrading from Dynamic Server Version 9.40.xCx, you should either perform a manual conversion using **conpload.sh** or set the **IFX\_ONPLOAD\_AUTO\_UPGRADE** environment variable to 1, because server conversion is not required.

---

## Customize Configuration Files

You can customize your ONCONFIG configuration file and environment variables to take advantage of the new features that Dynamic Server 9.40 introduced. When you initialize Dynamic Server 9.40, use the same configuration that the source database server used. After you observe the performance of Dynamic Server 9.40, you might want to adjust the configuration.

While you can use an old ONCONFIG configuration files with Version 9.4, IBM recommends that you use the new Version 9.4 ONCONFIG file, or at least examine the file for new parameters.

IBM recommends that you 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 conversion. For more details, see “Initialize Dynamic Server 9.40” on page 5-13. After the conversion, change the value of ALARMPROGRAM to **log\_full.sh**.

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 9.40, the default location of this library is **\$INFORMIXDIR/lib/ipldd09a.so**. For more information, see the *IBM Informix: High-Performance Loader User's Guide*.

For information on how to configure Dynamic Server 9.40, see your *IBM Informix: Administrator's Guide*. For information about environment variables, see the *IBM Informix: Guide to SQL Reference*. For information about how to tune the configuration parameters, see the *IBM Informix: Performance Guide*.

**Important:** To facilitate reversion, use the same values for your target database server for ROOTOFFSET, ROOTSIZE, and ROOTPATH that you used for the source 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.

---

## Add Any Communications Support Modules (UNIX/Linux)

You can use a Communications Support Module (CSM) with Dynamic Server 9.40. After you install the CSM components, create entries in the **concsbm.cfg** file and in the options field of the **sqlhosts** file to configure the CSM. For information on how to set up the CSM, see the *IBM Informix: Administrator's Guide*.

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 also have a CSM installed and configured.

---

## Install and Configure Any DataBlade Modules

After you install Dynamic Server 9.40, and before you initialize the database server, install and register any DataBlade modules supplied by IBM or third-party vendors that you want to add to the database server. *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 DataBlade documentation.

---

## Initialize Dynamic Server 9.40

If you just installed Dynamic Server as user **root**, you must switch to user **informix** before initializing the server.

Initialize Dynamic Server 9.40 without disk initialization.

**Warning:** Dynamic Server 9.40 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 9.40 halts 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 conversion. 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 conversion, you need to back it up with **ontape**, the only backup tool you can use at this point. Issue the **ontape -a** command.

Execute the following command to bring Dynamic Server 9.40 online for the first time:

oninit

As Dynamic Server 9.40 comes online for the first time, it modifies certain disk structures. This operation should extend the initialization process by only a minute or two. In the unlikely event that 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 should rebuild the index as instructed.

---

## Monitor the Conversion Complete Status

Check your message log (**online.log**) for status messages that pertain to bringing Dynamic Server 9.40 online. The conversion process ends when the following message is written to **online.log**:

Conversion completed successfully

This message indicates that the conversion 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 each individual database conversion. If a particular database conversion fails, then you should try to connect to the database to find out the exact cause of the failure.

At the end of the conversion of each individual database, the conversion process 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. If the script fails for a database, however, for better performance run the following scripts as user **informix** while connected to the database for converting to Dynamic Server 9.40 from Dynamic Server 7.3x or 7.24:

**\$INFORMIXDIR/etc/dummyups7x.sql**

If you encounter a failure during the conversion, you should revert to the source database server and then restore it from a backup.

The conversion process involves two phases: internal conversion and external conversion. Internal conversion is the process of converting each individual database in the system. External conversion is the process of converting the utilities, such as ON-Bar.

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

**Important:** If the message file notes problems, solve the problems before you continue to the next step.

---

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

After installation of the database server, 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 source database server to the target database server installation.
- If you are using other storage managers, copy your previous `sm_versions` file from the source `$INFORMIXDIR/etc` directory to the target `$INFORMIXDIR/etc` directory.
- If you are upgrading from Version 7.3x before 7.31.xC5 or from Version 7.24, unload the contents of the `sysutils:bar_version` table.

---

## Update Statistics

After a successful conversion, you need to run `UPDATE STATISTICS` on some of the system catalog tables in your databases, as the following lists show.

For a conversion to Dynamic Server 9.40 from a Version 7.3x or Version 7.24 database server, run `UPDATE STATISTICS` on the following system catalog tables in Dynamic Server 9.40:

<code>sysblobs</code>	<code>sysfragments</code>	<code>syssynonyms</code>
<code>syscolauth</code>	<code>sysindices</code>	<code>syssyntable</code>
<code>syscolumns</code>	<code>sysobjstate</code>	<code>systabauth</code>
<code>sysconstraints</code>	<code>sysopclstr</code>	<code>systables</code>
<code>sysdefaults</code>	<code>sysprocauth</code>	<code>systriggers</code>
<code>sysdistrib</code>	<code>sysprocedures</code>	<code>sysusers</code>
<code>sysfragauth</code>	<code>sysroleauth</code>	

---

## Verify the Integrity of the Data

After Dynamic Server finishes converting the system catalog tables, open each database with DB–Access and use `oncheck` to verify that no data was corrupted in the migration process. You can verify the integrity of the reserve pages, extents, system catalog tables, data, indexes, and smart large objects, as Table 5-3 shows.

Table 5-3. Commands for Verifying the Data Integrity

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

---

## Make an Initial Backup of Dynamic Server 9.40

Use a Dynamic Server 9.40 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 source 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 source database server for your target database server.

---

## Tune Dynamic Server 9.40 for Performance

When you finish the level-0 backup, the migration process is complete and users can safely use Dynamic Server 9.40 to access data.

After successful migration to Dynamic Server 9.40, you can tune the database server to obtain maximum performance. If you created sample queries for comparison, you can use them to characterize the performance differences between the source database server and Dynamic Server 9.40. The results of these comparisons might suggest adjustments to configuration parameters or to the layout of databases, tables, and chunks. For details on performance topics, see the *IBM Informix: Performance Guide*.

---

## Enable HDR

You do not need to perform the conversion process on the secondary server. On the secondary server, install the new database server and any user-defined objects or DataBlade modules, and then copy the necessary supporting files from the primary server: for example, the ONCONFIG file. For a list of requirements, see the section on configuring a system for HDR in the *IBM Informix: Administrator's Guide*. The secondary server is converted



automatically when you start HDR. To start HDR, follow the procedure for starting HDR for the first time as described in the *IBM Informix: Administrator's Guide*.



---

## Chapter 6. Reverting from Dynamic Server 9.40

Reverting from Dynamic Server 9.40 to Versions 9.30, 9.2x, 7.3x, or 7.24. . . . .	6-1
Determine Whether Reversion Is Possible. . . . .	6-2
Check and Configure Available Space . . . . .	6-6
Save Copies of the Current Configuration Files (UNIX/Linux). . . . .	6-6
Verify the Integrity of the Data . . . . .	6-6
Back Up Dynamic Server 9.40. . . . .	6-6
Remove BladeManager Extensions . . . . .	6-6
Remove Version 9.40 Features. . . . .	6-7
Disable HDR . . . . .	6-7
Run the Reversion Utility . . . . .	6-7
Modify Configuration Parameters . . . . .	6-8
Reset Environment Variables . . . . .	6-8
Remove Any Communications Support Module Settings (UNIX/Linux) . . . . .	6-8
Reinstall and Start the Target Database Server . . . . .	6-8
Update Statistics . . . . .	6-8
Verify the Integrity of the Data . . . . .	6-9
Back Up the Target Database Server . . . . .	6-9
Return the Target Database Server to Online Mode . . . . .	6-9
Enable HDR . . . . .	6-9

---

### In This Chapter

This chapter describes the steps for reverting from Dynamic Server 9.40 to Dynamic Server 9.30, 9.2x, 7.3x, or 7.24.

---

### Reverting from Dynamic Server 9.40 to Versions 9.30, 9.2x, 7.3x, or 7.24

**Important:** Before you revert from Dynamic Server 9.40:

- Remove all objects from the databases that the old database server does not support, except those that the boot scripts created in the system catalog. Do not remove the objects that the **boot90.sql** and **boot901.sql** scripts created because the reversion utility uses them.
- Read the release notes and the machine notes for any late-breaking information.

If you are reverting to Dynamic Server 7.3, you must drop any index whose key size is greater than 254. Dynamic Server 9.40 supports a maximum index key size of 255 for VARCHAR and NVARCHAR types. However, Dynamic Server 7.3 only supports a maximum index key size of 254.

To revert from Dynamic Server 9.40, complete the following steps. The subsections below describe these steps in detail.

**To revert from Dynamic Server 9.40:**

1. Review the database schema to determine whether reversion is possible.
2. Check and configure available space.
3. Save copies of the current configuration files.
4. Verify the integrity of the data.
5. Back up Dynamic Server 9.40.
6. Remove new features.
7. Remove BladeManager extensions.
8. If necessary, disable HDR.
9. If necessary, perform Enterprise Replication reversion tasks (see “Reverting from Dynamic Server 9.40 with Enterprise Replication” on page 4-3).
10. Run the reversion utility (**onmode -b**).
11. Modify configuration parameters.
12. Reset environment variables.
13. If necessary, remove any Communications Support Module settings.
14. Reinstall and start the target database server.
15. Update statistics.
16. Verify the integrity of the data.
17. Back up the target database server.
18. Return the target database server to online mode.

**Determine Whether Reversion Is Possible**

The easiest reversion scenario is to a Dynamic Server 9.40 database that does not contain any new features. Run the reversion utility and modify the values of the configuration parameters.

You can revert from Dynamic Server 9.40 to a 9.30, 9.2x, 7.3x, or 7.24 database server if you have not added any extensions to the 9.40 database server.

**Warning:** 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.

When you revert to the older version, you see a message similar to the following:

```

This will make all necessary modifications to disk structures
so that the Informix Dynamic Server space will be
compatible with Informix Dynamic Server Version 7.31
Beginning process of reverting system to 7.31 ...
...
Reversion complete.
Install Informix Dynamic Server Version 7.31 before reinitializing.
Do you wish to continue (y/n)? qaonmode -b 7.31: passed

```

In the next-to-last line of the message, *reinitializing* refers to restarting the database server (sometimes referred to as reinitializing shared memory), *not* to reinitializing the existing root dbspace.

Review the database schema to determine whether reversion to the earlier database server is possible. Ask the following questions:

- Does the schema file contain SQL statements that the earlier database server does not support?
- Does the database contain features that the earlier database server does not support, such as long identifiers?
- Have any new SPL routines been created in Dynamic Server 9.40, or have you imported existing routines using **dbimport**?
- Have tables or indexes using expression fragmentation had expressions changed or new fragments added?
- Have any new triggers or procedures been created?
- Have any new check constraints been created?

To review the database schema, execute the **dbschema** utility command. The following example displays complete information about the database **db1**:

```
dbschema -d db1 -ss
```

Database reversion occurs in two phases: a check phase and an actual reversion phase. If a database cannot be reverted, the check phase highlights it and prevents reversion.

The table below lists the restrictions for to a Version 9.30, 9.2x, or 7.x (7.31, 7.30, or 7.24) database server.

*Table 6-1. Reversion Restrictions*

7.x	9.2x	9.30	Restriction
x	x	x	You must drop all triggers and procedures created with the new version before reverting.
x	x	x	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. For more information, see “Chunk Format” on page 3-14.
x	x	x	No files larger than 2 GB should be stored in or in use by the database server.

Table 6-1. Reversion Restrictions (continued)

7.x	9.2x	9.30	Restriction
x	x	x	The TAPESIZE or LTAPESIZE configuration parameters cannot be set to 0.
x	x	x	The ALARMPROGRAM configuration parameter cannot be set to the <b>alarmprogram.sh</b> file.
x	x	x	The LRU_MAX_DIRTY or LRU_MIN_DIRTY configuration parameters must not be set to a value of type FLOAT. To revert, set them to integers.
x	x	x	UDRs should not use multiple IN or OUT parameters. Drop all such UDRs before reversion.
x	x	x	UDRs and stored procedures should not use named return parameters.
x	x	x	Sequence objects should not be in use. Drop all sequences before reversion.
x	x	x	No triggers created with the INSTEAD OF clause should be in use. Drop all such triggers before reversion.
x	x	x	Multiple collations among indexes, stored procedures, triggers, and constraints should not be in use.
x	x	x	Functional indexes cannot contain more than 16 parameters.
x	x	x	High-Data Availability Replication and Enterprise Replication cannot co-exist on the same database server.
x	x	x	The LVARCHAR( <i>n</i> ) data types should not be in use if <i>n</i> is not equal to 2042.
x	x	x	<p>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 need to 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, you need to 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>
x	x	x	No new user-defined or SPL routines should have been created in the converted databases (either implicitly or explicitly). If you plan to use <b>dbexport</b> to export a database containing existing user-defined or SPL routines, you must drop these routines prior to reversion.
x	x	x	No new triggers should have been defined in the converted databases.
x	x	x	<p>No fragment expressions or check constraints created on the 9.40 database server should 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 conversion to Dynamic Server 9.40.</p>
x			Indexes created with Dynamic Server 9.40 and an <b>opclass</b> that supports nearest-neighbor search cannot be reverted to the earlier database server.

Table 6-1. Reversion Restrictions (continued)

7.x	9.2x	9.30	Restriction
x			Reversion fails if, for an index, the value of <b>item_nvarchar</b> is 255 or higher.
x			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.
x			A DataBlade module that uses the PER_STMT_EXEC or PER_STMT_PREP memory duration cannot be used with the earlier database server.
x			You cannot revert a database that was created with the database server from which you are reverting.  Drop the database before you attempt reversion.
x			Select triggers should not be in use.
x			User-defined statistics should not be in use.
x			No long identifiers or long usernames should be in use.  Before reversion, make sure that the R-tree indexes do not use long identifiers as indexed column names, opclass names, or opclass function names.  Also, make sure that the following disk structures do not use long identifiers: <ul style="list-style-type: none"> <li>• Database tablespaces (owner and database name length)</li> <li>• Tablespace tablespaces (owner and tablespace name length)</li> <li>• Dbspaces (owner and dbspace name length) and chunks (path length)</li> </ul>
x			No storage space should have a name more than 18 characters long.
x			No in-place ALTER TABLE statement should be outstanding against any table.  If a user table has an incomplete in-place ALTER operation, then you need to 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, it lists all the tables that need dummy updates. You need to perform a dummy update on each of the tables in the list before you can revert to the older database server.  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 7.3x or 7.24 database server from Dynamic Server 9.40:  <b>\$INFORMIXDIR/etc/dummyupds7x.sql</b>
x			No new routine languages should be defined in the converted databases.
x			No new language authorizations must have been done in the converted databases.
x			No new operator classes, casts, or extended types should be defined in the 9.40 database server.
x			No semidetached indexes should be in the databases.

Table 6-1. Reversion Restrictions (continued)

7.x	9.2x	9.30	Restriction
x			The databases cannot have tables whose primary access method is a user-defined access method.
x			Databases cannot have typed tables.
x			Tables cannot have any user-defined type columns.
x			Tables cannot have named row types with default values.
x			All indexes must be B-tree indexes with a total key length less than or equal to 255.
x			Tables cannot have any functional or VII indexes.
x			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 authorizations, and casts.

## Check and Configure Available Space

For Dynamic Server 9.40, chunk reserve pages can be allocated in non-root chunks. If the 9.40 root chunk is full and chunk reserve pages were allocated in non-root chunks, make sure you have enough space in the root chunk of the target database server. During reversion, all chunk reserve pages are written to the root chunk of the 9.30, 9.2x, 7.3x, or 7.24 database server.

## Save Copies of the Current Configuration Files (UNIX/Linux)

Save copies of the ONCONFIG and **concsbm.cfg** files for when you convert to Dynamic Server 9.40 again. Dynamic Server 9.40 uses the **concsbm.cfg** file to configure CSMs.

## Verify the Integrity of the Data

Execute the following commands to verify the integrity of the data:

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

## Back Up Dynamic Server 9.40

Before you begin the reversion, make a complete level-0 backup of Dynamic Server 9.40. For more information, see the *IBM Informix: Backup and Restore Guide*.

## Remove 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 a 9.40



database server, and you have run BladeManager against a database, you must first run BladeManager and specify the following command to remove the BladeManager extensions:

```
unprep database name
```

## Remove Version 9.40 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 “Determine Whether Reversion Is Possible” on page 6-2.

## Disable HDR

If you use High-Availability Data Replication (HDR), you must disable it before reversion from Dynamic Server 9.40.

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

```
onmode -d standard
```

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

**Tip:** You do not need to perform the whole reversion process on the secondary database server. For more information, see “Enable HDR” on page 6-9.

## Run the Reversion Utility

Dynamic Server 9.40 must be running when you execute the reversion utility. The reversion utility detects and lists any remaining features that are specific to Dynamic Server 9.40. You need to remove these features before reversion can complete.

Execute the reversion utility with one of the following commands:

- `onmode -b 9.3`
- `onmode -b 9.2`
- `onmode -b 7.3`
- `onmode -b 7.2`

After the reversion is complete, Dynamic Server 9.40 is offline. The reversion utility drops the Dynamic Server 9.40 system catalog tables and restores compatibility so that users 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,” on page 12-1.

## Modify Configuration Parameters

Replace the Dynamic Server 9.40 ONCONFIG configuration file with the ONCONFIG file that you used before you converted. 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 “Configuration Parameters” on page 3-4.

## Reset Environment Variables

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

## Remove Any Communications Support Module Settings (UNIX/Linux)

If your Dynamic Server 9.40 instance used CSMs, remove any **csn** option settings that are not supported in the older database server from the **sqlhosts** file entries for the database server. Otherwise, the older database server will return an invalid **sqlhosts** options error. Delete the **concsn.cfg** file if the older database server does not support CSMs.

## Reinstall and Start the Target Database Server

Reinstall and configure the 9.30, 9.2x, 7.3x, or 7.24 database server according to the instructions in your *IBM Informix: Installation Guide* and your *IBM Informix: Administrator's Guide*.

Execute the **oninit -s** command to put the 9.30, 9.2x, 7.3x, or 7.24 database server in quiescent mode.

## Update Statistics

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

For reversion to a 7.3x or 7.24 database server from Dynamic Server 9.40 run UPDATE STATISTICS on the following system catalog tables in the 7.3x or 7.24 database server:

<b>SYSBLOBS</b>	<b>SYSFRAGMENTS</b>	<b>SYSSYNONYMS</b>
<b>SYSCOLAUTH</b>	<b>SYSINDEXES</b>	<b>SYSSYNTABLE</b>
<b>SYSCOLUMNS</b>	<b>SYSOBJSTATE</b>	<b>SYSTABAUTH</b>
<b>SYSCONSTRAINTS</b>	<b>SYSOPCLSTR</b>	<b>SYSTABLES</b>
<b>SYSDEFAULTS</b>	<b>SYSROCAUTH</b>	<b>SYSTRIGGERS</b>
<b>SYSDISTRIB</b>	<b>SYSPROCEDURES</b>	<b>SYSUSERS</b>
<b>SYSFRAGAETH</b>	<b>SYSROLEAUTH</b>	

## Verify the Integrity of the Data

Before you allow users to access the databases, check the integrity of the data. Follow the steps under “Verify the Integrity of the Data” on page 5-8.

## Back Up the Target Database Server

After you complete the reversion, use ON-Bar or **ontape** to make a level-0 backup of the 9.30, 9.2x, 7.3x, or 7.24 database server. 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 Target Database Server to Online Mode

To bring the 9.30, 9.2x, 7.3x, or 7.24 database server online, execute the **onmode -m** command. Then users can access the data.

## Enable HDR

You do not need to perform the reversion process on the secondary server. On the secondary server, install the target database server and any user-defined objects or DataBlade modules, and then copy the necessary supporting files from the primary server: for example, the ONCONFIG file. For a list of requirements, see the section on configuring a system for HDR in the *IBM Informix: Administrator's Guide*. The secondary server is reverted automatically when you start HDR. To start HDR, follow the procedure for starting HDR for the first time in the *IBM Informix: Administrator's Guide*.



---

## **Part 3. Migration of Data Between Database Servers**



---

# Chapter 7. Migrating Between Database Servers and Operating Systems

Choosing a Migration Method . . . . .	7-1
Adjusting Database Tables for File-System Variations . . . . .	7-2
Moving Data to a Database Server on a Different Operating System . . . . .	7-2
Using the Migration Utilities . . . . .	7-3
Adapting Your Programs for UNIX or Windows . . . . .	7-3
Completing Migration . . . . .	7-4
Moving Data Between Dynamic Server and Workgroup Edition on Different Operating Systems . . . . .	7-4

---

## In This Chapter

This chapter describes the steps for moving data between database servers. This chapter covers the following topics:

- Choosing a migration method
- Adjusting database tables for file-system variations
- Moving data to a database server on a different operating system
- Moving data between Dynamic Server and Workgroup Edition on different operating systems

---

## Choosing a Migration Method

UNIX or Linux and Windows store data in different page sizes. When your migration involves 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.

The method 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. You can use the following migration methods:

- **dbexport** and **dbimport**  
To move an entire database, use the **dbexport** and **dbimport** utilities.
- **UNLOAD** and **LOAD**  
To move selected columns or tables, use the **UNLOAD** statement. Use **LOAD** when you do not want to change the data format.
- **UNLOAD**, **dbload**, and **dbschema**  
To move selected columns or tables, use the **UNLOAD** statement. Use **dbload** to change the data format.

- **onunload** and **onload**

To unload data in page-sized chunks, use the **onunload** utility. Use the **onload** utility to move data to an identical database server on a computer of the same type.

- The High-Performance Loader (HPL)

To move selected columns or tables or an entire database, use the HPL.

**Important:** Do not use **onunload** or **onload** with Dynamic Server 9.30 or 9.2x. You can use the HPL for Dynamic Server 9.30 or 9.2x.

For information on migration paths from some older database servers to intermediary servers before migrating to Version 9.40, see “Database Server Migration Paths” on page 1-7.

---

## 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 three-gigabyte table, but your operating system allows only two-gigabyte files, break up your table into separate files before you migrate. For more information, see your *IBM Informix: Administrator's Guide*.

**Important:** Informix database servers support only certified versions of NFS. For information about the NFS products you can use to NFS mount a storage space for an Informix database server, check product compatibility information at [www.ibm.com/software/data/informix/pubs/smv/index.html](http://www.ibm.com/software/data/informix/pubs/smv/index.html).

---

## Moving Data to a Database Server on a Different Operating System

This section describes the steps for moving 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* or your *IBM Informix: Archive and Backup 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.

## Using the Migration Utilities

If you intend to move an entire database on Dynamic Server 9.40, 9.30, 9.2x, 7.3x, or 7.24 between different environments, the **dbexport** and **dbimport** combination is the easiest migration method.

You can also choose one of the following migration utilities:

- If you want to move selected tables or columns, instead of an entire database, use the UNLOAD and LOAD statements 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 **dbexport**, **dbimport**, UNLOAD, LOAD, **dbload**, and **dbschema**, see Part 4, “Data Migration Utilities” of this guide.

## Adapting Your Programs for UNIX or Windows

Certain database server configuration parameters and environment variables are environment dependent, as follows:

- Dynamic Server 9.40, 9.30 and 9.2x support Enterprise Replication.
- Dynamic Server 7.3x supports Enterprise Replication and uses Version 3.0 of IBM Informix Enterprise Command Center (IECC).
- Workgroup Edition 7.3x 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: Administrator's Guide* and the *IBM Informix: Administrator's*

*Reference and the information on environment variables in your IBM Informix: Administrator's Guide and the IBM Informix: Guide to SQL Reference.*

## Completing Migration

The first time the target database server is brought online, the **sysmaster** and **sysutils** databases are built. 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 Part 4, "Data Migration Utilities" of this guide. For information on how to use DB-Access, see the *IBM Informix: DB-Access User's Guide*.

For information about using UNLOAD, **dbschema**, and LOAD to move data between Dynamic Server 7.3x or 7.24 and Workgroup Edition or information about migrating to Workgroup Edition 7.3x from Dynamic Server 7.3x on a Different Operating Systems, see the Version 7.3x Migration Guide, Version 7.3x installation information, and Version 7.3x Release Notes.

---

## Part 4. Data Migration Utilities



---

# Chapter 8. The dbexport and dbimport Utilities

Syntax of the dbexport Command . . . . .	8-2
Syntax of the dbimport Command . . . . .	8-7
Simple Large Objects (IDS 9.x) . . . . .	8-14
Database Locale Changes . . . . .	8-14

---

## In This Chapter

This chapter describes the **dbexport** and **dbimport** utilities and how to use them. You can use **dbexport** and **dbimport** with the following database servers:

- Dynamic Server 9.40, 9.30, 9.2x, 7.3x, or 7.24
- Dynamic Server, Linux Edition 7.3x
- Workgroup Edition 7.3x or 7.24
- SE
- OnLine 5.1x

The **dbexport** utility unloads a database into text files for later import into another database and creates a schema file. 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 edit the schema file to modify the database that **dbimport** creates. The **dbexport** and **dbimport** utilities support Dynamic Server 9.40, 9.30, and 9.2x extended data types.

Dates are stored in four-digit years. By default, **dbexport** exports dates in four-digit year dates unless the environment variable **DBDATE** is set to “mdy2” or to some other value that specifies abbreviated years. We do not recommend this setting for exporting a database 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:** Disable SELECT triggers before exporting a database with **dbexport**. The **dbexport** utility executes SELECT statements during export. The SELECT statement triggers can modify the database content.

**Warning:** 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.

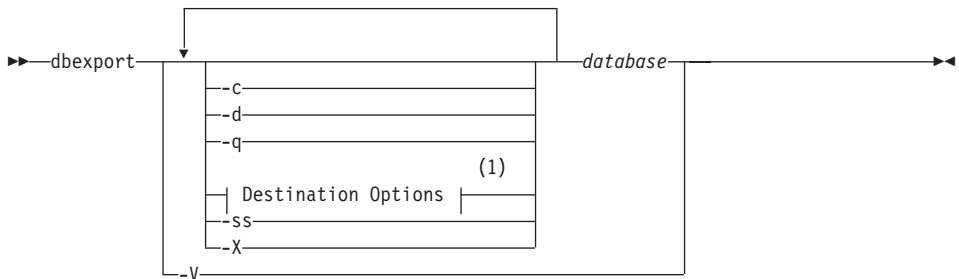
**Note:** If the date context during import is not the same as when these objects were created, you might get explicit errors, or you might not be able to find your data, or a check constraint might not work as expected, and so forth. Many of these problems do not generate errors. The date context for 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*.

### End of Global Language Support

## Syntax of the dbexport Command



### Notes:

- 1 See page 8-4

Element	Purpose	Key Considerations
-c	Makes <b>dbexport</b> complete exporting unless a fatal error occurs	<b>References:</b> For details on this option, see “Errors” on page 8-4.
-d	Makes <b>dbexport</b> export simple-large-object descriptors only, not simple-large-object data	<b>References:</b> For information about simple-large-object descriptors, see the <i>IBM Informix: Optical Subsystem Guide</i> .  <b>Restrictions:</b> Not supported by SE.
-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	<b>References:</b> For details on this option, see “Server-Specific Information” on page 8-4.
-X	Recognizes HEX binary data in character fields	None.
-V	Displays product version information	None.
database	Specifies the name of the database that you want to export	<b>Additional Information:</b> If your locale is set to use multibyte characters, you can use multibyte characters for the database name.  <b>References:</b> 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 8-13.

### End of Global Language Support

You can use delimited identifiers with the **dbexport** utility. The utility detects database objects that are keywords, mixed case, or have special characters and 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 the standard output unless you specify the **-q** option.

During the 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 dbexport

You can press the Interrupt key at any time to cancel **dbexport**. The **dbexport** utility asks for confirmation before it terminates.

## Errors

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

## Server-Specific Information

The **-ss** option generates server-specific information. The **-ss** 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.

## Destination Options

### Destination Options:

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



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

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

The following command exports the database **stores\_demo** to tape with a block size of 16 kilobytes and a tape capacity of 24,000 kilobytes. The schema file is written 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
```

End of UNIX/Linux Only

Windows Only

For Windows, the following command exports the database **stores\_demo** to tape with a block size of 16 kilobytes and a tape capacity of 24,000 kilobytes. 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
```

End of Windows Only

## Contents of the Schema File

The schema 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 Dynamic Server 9.40, 9.30, and 9.2x data types.

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 dbspace 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 executes 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.


**Warning:** Do not delete any comments in the schema file. It is strongly recommended that you do not change any existing comments or add any new comments, or the **dbimport** might abort or produce unpredictable results.

**Note:** The number of rows should 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.

**Tip:** 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.

---

## Syntax of the dbimport Command



### Notes:

- 1 See page 8-9
- 2 See page 8-11

Element	Purpose	Key Considerations
<b>-c</b>	Instructs <b>dbimport</b> to complete importing even though it encounters certain nonfatal errors	<b>References:</b> For more information, see “Errors and Warnings” on page 8-9.
<b>-q</b>	Suppresses the display of error messages, warnings, and generated SQL data-definition statements	None.
<b>-V</b>	Displays product version information	None.
<b>-X</b>	Recognizes HEX binary data in character fields	None.
<b>database</b>	Specifies the name of the database to create	<b>Additional Information:</b> 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.

### End of Global Language Support

## Termination of dbimport

To cancel **dbimport**, press the Interrupt key at any time. The **dbimport** program asks for confirmation before it terminates.

## Errors and Warnings

If you include the **-c** option, **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.

## Input-File Location Options

The input-file location tells **dbimport** where to look for the *database.exp* directory, which contains the files that **dbimport** will import. If you do not specify an input-file location, **dbimport** looks for data files in the directory *database.exp* under the current directory and for the schema file in *database.exp/database.sql*.

### Input-File Location:

-i—directory—			
-t—device—	-b—blocksize—	-s—tapesize—	
-t—pathname—			

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

#### UNIX/Linux Only

The following command imports the **stores\_demo** database from a tape with a block size of 16 kilobytes and a capacity of 24,000 kilobytes. The schema file is read from **/tmp/stores\_demo.imp**.

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

The following command imports the **stores\_demo** database from the **stores\_demo.exp** directory under the **/work/exports** directory. The schema file is assumed to be **/work/exports/stores\_demo.exp/stores\_demo.sql**.

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

## End of UNIX/Linux Only

## Windows Only

The following command imports the **stores\_demo** database from a tape with a block size of 16 kilobytes and a capacity of 24,000 kilobytes. The schema file is read from **C:\temp\stores\_demo.imp**.

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


The following command imports the **stores\_demo** database from the **stores\_demo.exp** directory under the **D:\work\exports** directory. The schema file is assumed to be **D:\work\exports\stores\_demo.exp\stores\_demo.sql**.

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

## End of Windows Only

## Create Options

### Create Options:



Element	Purpose	Key Considerations
<b>-ansi</b>	Creates an ANSI-compliant database in which the ANSI rules for transaction logging are enabled	<b>Additional Information:</b> If you specify the <b>-ansi</b> option, you must also specify the <b>-l logfile</b> option. For more information about ANSI-compliant databases, see the <i>IBM Informix: Guide to SQL Reference</i> .
<b>-d dbspace</b>	Specifies the dbspace where the database is created. The default dbspace location is the <b>rootdbs</b> .	<b>Additional Information:</b> For SE, the database is always in the current directory.
<b>-l</b>	Establishes unbuffered transaction logging for the imported database	<b>References:</b> For more information, see “Database-Logging Mode” on page 8-13.
<b>-l buffered</b>	Establishes buffered transaction logging for the imported database	<b>References:</b> For more information, see “Database-Logging Mode” on page 8-13.
<b>-l logfile</b>	Establishes transaction logging for the imported database and specifies the name of the transaction-log file	<b>Restrictions:</b> For SE, the <i>logfile</i> filename must be an absolute pathname or in the current directory.  <b>References:</b> For more information, see “Database-Logging Mode” on page 8-13.

#### UNIX/Linux Only

The following command imports the **stores\_demo** database from the **/usr/informix/port/stores\_demo.exp** directory. The new database is ANSI compliant, and the transaction-log file is specified as **stores\_demo.log** in **/usr/work**.

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

#### End of UNIX/Linux Only

#### Windows Only

The following command imports the **stores\_demo** database from the **C:\USER\informix\port\stores\_demo.exp** directory. The new database is ANSI compliant, and the transaction-log file is specified as **stores\_demo.log** in **C:\USER\work**.

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

#### End of Windows Only



## Database-Logging Mode

The logging mode is not retained in the schema file. You can specify any of the following options when you use **dbimport** to import a database:

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

For more information, see “Create Options” on page 8-11.

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.

In the following 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 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 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
```

---

## Simple Large Objects (IDS 9.x)

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

---

### End of Windows Only

**Warning:** 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.

---

## Database Locale Changes

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.



# Chapter 9. The dbload Utility

Syntax of the dbload Command . . . . .	9-1
Command File for dbload . . . . .	9-5
Command File to Load Complex Data Types (IDS 9.x) . . . . .	9-14

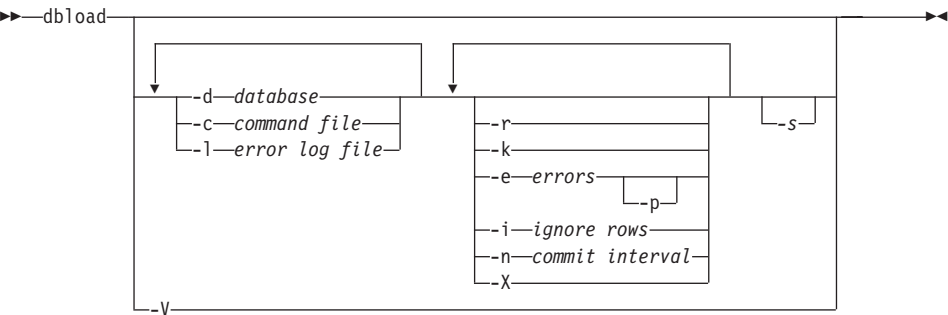
## In This Chapter

This chapter describes the **dbload** utility and how to use it. You can use **dbload** with the following database servers:

- Dynamic Server 9.40, 9.30, 9.2x, 7.3x or 7.24
- Dynamic Server, Linux Edition 7.3x
- Workgroup Edition 7.3x or 7.24
- SE
- OnLine 5.1x

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 Dynamic Server 9.40, 9.30 and 9.2x.

## Syntax of the dbload Command



Element	Purpose	Key Considerations
-c <i>command file</i>	Specifies the filename or pathname of a <b>dbload</b> command file	<b>References:</b> For information about building the command file, see “Command File for dbload” on page 9-5.

Element	Purpose	Key Considerations
<b>-d database</b>	Specifies the name of the database to receive the data	<b>Additional Information:</b> 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> .
<b>-e errors</b>	Specifies the number of bad rows that <b>dbload</b> reads before terminating . The default value for <i>errors</i> is 10.	<b>References:</b> For more information, see “Bad-Row Limit” on page 9-3.
<b>-i ignore rows</b>	Specifies the number of rows to ignore in the input file	<b>References:</b> For more information, see “Rows to Ignore” on page 9-3.
<b>-k</b>	Instructs <b>dbload</b> to lock the tables listed in the command file in exclusive mode during the load operation	<p><b>References:</b> For more information, see “Table Locking” on page 9-3.</p> <p><b>Restrictions:</b> You cannot use the <b>-k</b> option with the <b>-r</b> option because the <b>-r</b> option specifies that no tables are locked during the load operation.</p>
<b>-l error log file</b>	Specifies the filename or pathname of an error log file	<p><b>Restrictions:</b> If you specify an existing file, its contents are overwritten. If you specify a file that does not exist, <b>dbload</b> creates the file.</p> <p><b>Additional Information:</b> The error log file stores diagnostic information and any input file rows that <b>dbload</b> cannot insert into the database.</p>
<b>-n commit interval</b>	<p>Specifies the commit interval in number of rows</p> <p>The default interval is 100 rows.</p>	<p><b>Additional Information:</b> If your database supports transactions, <b>dbload</b> commits a transaction after the specified number of new rows is read and inserted. A message appears after each commit.</p> <p><b>References:</b> For information about transactions, see the <i>IBM Informix: Guide to SQL Tutorial</i>.</p>
<b>-p</b>	Prompts for instructions if the number of bad rows exceeds the limit	<b>References:</b> For more information, see “Bad-Row Limit” on page 9-3.
<b>-r</b>	Prevents <b>dbload</b> from locking the tables during a load, thus enabling other users to update data in the table during the load	<p><b>Additional Information:</b> For more information, see “Table Locking” on page 9-3.</p> <p><b>Restrictions:</b> You cannot use the <b>-r</b> option with the <b>-k</b> option because the <b>-r</b> option specifies that the tables are not locked during the load operation while the <b>-k</b> option specifies that the tables are locked in exclusive mode.</p>
<b>-s</b>	Checks the syntax of the statements in the command file without inserting data	<b>Additional Information:</b> The standard output displays the command file with any errors marked where they are found.

Element	Purpose	Key Considerations
-V	Displays product version information	
-X	Recognizes HEX binary data in character fields	

**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.

## Table Locking

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.

To override this default lock mode, specify the **-k** option. The **-k** option instructs **dbload** to lock the tables in exclusive mode rather than shared mode during the load operation.

## Rows to Ignore

The **-i** option instructs **dbload** to read and ignore the specified number of new-line characters in the input file before it begins to process. 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. It is also useful if header information in the input file precedes the data records.

## Bad-Row Limit

The **-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.

## Guidelines for Using dbload

This section includes the following guidelines for using the **dbload** utility:

- Termination of dbload
- Network names
- Simple large objects
- Indexes
- Delimited identifiers
- SE Example

### Termination of dbload

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).

### 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 or coserver.

### 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.

### Delimited Identifiers

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.



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.

### Example

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

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

---

## Command File for dbload

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 need to 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.

## FILE and INSERT Statements: 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**. To see the **.unl** input data files, refer to the directory **\$INFORMIXDIR/demo/prod\_name** (UNIX or Linux) or **%INFORMIXDIR%\demo\prod\_name** (Windows).

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

## Syntax for the Delimiter Form

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	<b>Restrictions:</b> If the delimiter specified by <b>c</b> 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 <b>c</b> is specified as a square bracket ([) , 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.

**Warning:** 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 testtbl
      (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 execute **dbload** with this command file must have the Insert privilege on the named table.

#### How to Write a dbload Command File in Delimiter Form

The first **FILE** and **INSERT** statement set in the delimiter example on page 9-6 is repeated in the following example:

```

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

```

The **FILE** statement describes the **stock.unl** data rows as composed of six fields, each separated by a vertical bar (|) as the delimiter. 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|HR0|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 in the **col4** position (assuming that the column permits nulls). 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

## FILE and INSERT Statements: Character-Position Form

The examples in this section 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	CA	94086408-789-8075		
Denver	CO	80219303-936-7731		
Blue Island	NY	60406312-944-5691		
Brighton	MA	02135617-232-4159		
Tempe	AZ	85253xxx-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 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	<b>Restrictions:</b> 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 <b>dbload</b> should substitute a null	<b>Restrictions:</b> 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 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 SQLS

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 executes **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

The first FILE and INSERT statement set in the character-position example on page 9-10 is repeated in the following example:

```

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 FILE statement defines six data fields from the **cust\_loc\_data** table data rows. 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 Figure 9-2.




Figure 9-2. A Sample Data File

The FILE statement defines the following data fields, which are derived from the data rows in Figure 9-2.

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 Figure 9-2 and the FILE statement on page 9-12 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 (IDS 9.x)

This section describes how to write **dbload** command files that load columns that contain complex data types into tables. The examples cover how to use **dbload** with named row types, unnamed row types, sets, and lists.

### Using dbload 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. In fact, you can follow these steps for any user-defined data type.

This example uses a table **person** that contains one column with a named row type. 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|5820 Easy Ave #100|Fremont|CA|94502|01/13/1983|
```

1. 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
```

2. Use DB–Access to re-create the **person** table in the new database. For detailed steps, see “DB–Access Input from dbschema Output” on page 10-13.

3. 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');
```

4. Execute 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 dbload with Unnamed Row Types

You can use **dbload** with unnamed row types. 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('Candy','S','Cane')|ROW('7 Willy Wonka
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('Craig','X','Smith')",
"row('1200 Cheese Ave.','Rainy City','OR','97200')");
```

## Using dbload with Collection Data Types

You can use **dbload** with collection data types such as SET, LIST, and MULTISET.

### SET Data Type Example

In a SET, each element is unique, and no nulls are allowed. The number of elements in a SET can vary. 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

A list is an ordered collection of elements that allows duplicate 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.

## Using dbload with Other Data Types

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

**Warning:** 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.

**Note:** 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 should be able to load it.

---

## Chapter 10. The dbschema Utility

Syntax of the dbschema Command . . . . .	10-2
Database Schema Creation . . . . .	10-3
Creating Schemas for Databases Across a UNIX or Linux Network . . . . .	10-4
Changing the Owner of an Object . . . . .	10-4
Server-Specific Information . . . . .	10-5
User-Defined and Complex Data Types (IDS 9.x) . . . . .	10-5
Sequence Creation . . . . .	10-6
Synonym Creation . . . . .	10-7
Privileges . . . . .	10-7
Granting Privileges . . . . .	10-8
Displaying Privilege Information for a Role. . . . .	10-8
Table, View, or Procedure Creation. . . . .	10-8
Table Information . . . . .	10-9
Role Creation . . . . .	10-10
Distribution Information for Tables . . . . .	10-10
Example Output . . . . .	10-11
Distribution Description . . . . .	10-11
Distribution Information. . . . .	10-12
Overflow Information . . . . .	10-13
DB-Access Input from dbschema Output . . . . .	10-13
Inserting a Table into a Database Example. . . . .	10-13
Re-Creating the Schema of a Database . . . . .	10-13

---

### In This Chapter

This chapter describes the **dbschema** utility and how to use it. You can use **dbschema** with the following database servers:

- Dynamic Server 9.40, 9.30, 9.2x, 7.3x or 7.24
- Dynamic Server, Linux Edition 7.3x
- Workgroup Edition 7.3x or 7.24
- SE
- OnLine 5.1x


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, synonym, sequence, 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

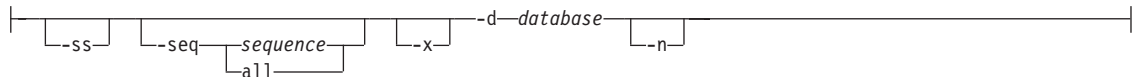
## Syntax of the dbschema Command



### Table Options:



### Database Options:



### Notes:

- 1 See page 10-5
- 2 See page 10-7
- 3 See page 10-7
- 4 See page 10-8
- 5 See *IBM Informix Guide to SQL: Syntax*
- 6 See page 10-10



Element	Purpose	Additional Information
<b>all</b>	Directs <b>dbschema</b> to include all the tables or sequence objects in the database, or all the user-defined data types in the display of distributions	None.
<b>-d database</b>	Specifies the database or co-server to which the schema applies. The <i>database</i> can be on a remote database server.	<b>References:</b> 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> .
<b>filename</b>	Specifies the filename to contain the <b>dbschema</b> output	If you omit <i>filename</i> , <b>dbschema</b> sends output to the screen. If you specify a <i>filename</i> , <b>dbschema</b> creates a file named <i>filename</i> to contain the <b>dbschema</b> output.
<b>-hd</b>	Displays the distribution as data values	<b>References:</b> For more information, refer to "Distribution Information for Tables" on page 10-10.
<b>-ss</b>	Generates server-specific information	<b>Restrictions:</b> This option is ignored if no table schema is generated.  <b>References:</b> For more information, refer to "Server-Specific Information" on page 10-5.
<b>-seq sequence</b>	Generates the DDL statement to define the specified <i>sequence</i> object	<b>References:</b> For more information, refer to "Sequence Creation" on page 10-6.
<b>-u</b>	Prints the definitions of user-defined data types	<b>References:</b> For more information, refer to "User-Defined and Complex Data Types (IDS 9.x)" on page 10-5.
<b>-ui</b>	Prints the definitions of user-defined data types, including type inheritance	<b>References:</b> For more information, refer to "User-Defined and Complex Data Types (IDS 9.x)" on page 10-5.
<b>-V</b>	Displays product version information	None.
<b>-x</b>	Expands dbslice names into dbspace name lists in <b>-ss</b> 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.

#### Dynamic Server Versions 9.20 to 9.40

- Display user-defined and row data types with or without type inheritance.
- Display the CREATE SEQUENCE statement defining the specified *sequence* object, or defining all sequence objects in the database.

#### End of Dynamic Server Versions 9.20 to 9.40

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 *pathname* to a file in which compile-time warnings are stored. This pathname 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*.

## Server-Specific Information

The **-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 **-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 (IDS 9.x)

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 is executed 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** executed 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



Element	Purpose	Key Considerations
<b>-seq sequence</b>	Displays the CREATE SEQUENCE statement defining <i>sequence</i>	None.
<b>-seq all</b>	Displays all CREATE SEQUENCE statements for the database	None.

Executing **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

### Synonyms:

|--s--| ownername |-----|  
          |all|

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 is executed 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

### Privileges:

|--p--| user |-----|  
          |all|

Element	Purpose	Key Considerations
-p <i>user</i>	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	<b>Restriction:</b> You cannot specify a specific list of users with the <b>-p</b> 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

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

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.

The following **dbschema** command and output show the privileges that were granted for the **calen** role:


```
sharky% dbschema -p calen -d stores_demo
```

```
DBSCHEMA Schema Utility      INFORMIX-SQL Version 7.22
Copyright (C) IBM Corporation, 1984-1996
Software Serial Number RDS#N0000000
```

```
grant alter on table1 to 'calen'
```

---

## Table, View, or Procedure Creation



Element	Purpose	Key Considerations
<b>-f all</b>	Limits the SQL statement output to those statements that are needed to replicate all functions and procedures	None.
<b>-f function</b>	Limits the SQL statement output to only those statements that are needed to replicate the specified function	None.
<b>-f procedure</b>	Limits the SQL statement output to only those statements that are needed to replicate the specified procedure	None.
<b>-ff all</b>	Limits the SQL statement output to those statements that are needed to replicate all functions	None.
<b>-fp all</b>	Limits the SQL statement output to those statements that are needed to replicate all procedures	None.
<b>-t table</b>	Limits the SQL statement output to only those statements that are needed to replicate the specified table	None.
<b>-t view</b>	Limits the SQL statement output to only those statements that are needed to replicate the specified view	None.
<b>-t all</b>	Includes in the SQL statement output all statements that are needed to replicate all tables and views	None.
<b>-ti table</b>	Includes in the SQL statement output all statements that are needed to replicate all table levels	None.
<b>-ti all</b>	Includes in the SQL statement output all statements that are needed to replicate all tables and views  Functionally equivalent to <b>-t all</b> .	None.

For more information about the CREATE PROCEDURE and CREATE FUNCTION statements, see the *IBM Informix: Guide to SQL Syntax*.

## Table Information

When you use the **-ss** option, you can retrieve information about fragmented tables, the lock mode, and extent sizes.

The following **dbschema** output shows the expressions specified for fragmented table.

```
DBSCHEMA Schema Utility          INFORMIX-SQL Version X.XX.UC1
Copyright (C) IBM Corporation, 1984-1997
{ 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".tl from "public";

```

---

## Role Creation

### Roles:



Element	Purpose	Key Considerations
<b>-r role</b>	Displays the CREATE ROLE and GRANT statements that are needed to replicate and grant the specified role	<b>Restriction:</b> You cannot specify a list of users or roles with the <b>-r</b> option. You can specify either one role or all roles. SE does not support the <b>-r</b> option.
<b>-r all</b>	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          INFORMIX-SQL Version 7.22
Copyright (C) IBM Corporation, 1984-1996
Software Serial Number RDS#N0000000
create role calen;

```

```

grant calen to cathl with grant option;
grant calen to judith ;
grant calen to sallyc ;

```

---

## Distribution Information for Tables

To display the distribution information that is stored for a table in a database, use the **-hd** option with the name of the table. If you specify the **ALL** keyword for the table name, the distributions for all the tables in the database are displayed.

Distribution information is stored only if you have run the **UPDATE STATISTICS...MEDIUM** or **HIGH** statement for one or more columns of a table. For information about the **UPDATE STATISTICS** statement, refer to 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 Output

```
Distribution for cathl.invoices.invoice_num
```

```
Constructed on 03/10/2001
```

```
High Mode, 10.000000 Resolution
```

```
--- DISTRIBUTION ---
```

```
(
1: ( 16,      7,      5)
2: ( 16,      6,      11)
3: ( 16,      8,      17)
4: ( 16,      8,      25)
5: ( 16,      7,      38)
6: ( 16,      8,      52)
7: ( 16,     12,      73)
8: ( 16,     12,      95)
9: ( 16,     11,     139)
10: ( 10,      5,     182)
    ( 10,      5,     200)
```

```
--- OVERFLOW ---
```

```
1: (  5,      56)
2: (  6,      63)
}
```

## Distribution Description

The first part of the **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 cathl.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 date on which the distributions are constructed is listed. In this example, the date is 03/10/1995, which is the date when the UPDATE STATISTICS statement that generated the distributions was executed. You can use this date to tell how outdated your distributions are. Although the system records the date, it does not record the time.

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 is executed with 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 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:

	(		5)
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.

---

## DB-Access Input from dbschema Output

You can use the **dbschema** utility to get the schema of a database and redirect the **dbschema** output to a file. Later, you can use this file as input to DB-Access to re-create the database.

### Inserting a Table into a Database Example

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

Syntax of the UNLOAD Statement . . . . .	11-1
Syntax of the LOAD Statement . . . . .	11-2

---


## In This Chapter

- This chapter shows the syntax of the SQL UNLOAD and LOAD statements. You can use UNLOAD and LOAD with the following database servers:
- Dynamic Server 9.40, 9.3x, 9.2x, 7.3x or 7.24
  - Dynamic Server, Linux Edition 7.3x
  - Workgroup Edition 7.3x or 7.24
  - SE
  - OnLine 5.1x

---

## Syntax of the UNLOAD Statement

You can use the UNLOAD statement in DB–Access to unload selected rows from a table into a text file.



### Notes:

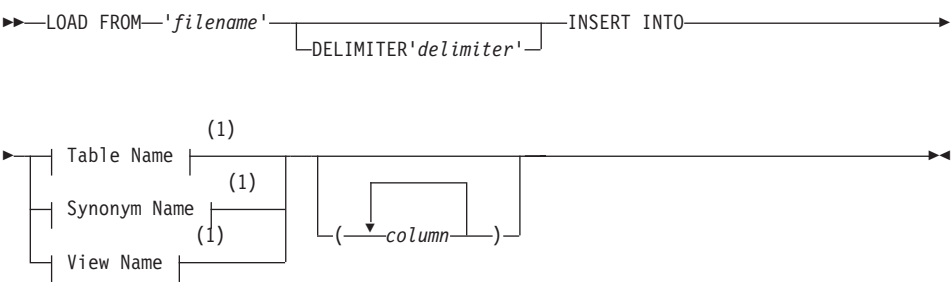
- 1 See the *IBM Informix: Guide to SQL Syntax*.

Element	Purpose	Key Considerations
<i>delimiter</i>	Character to use as delimiter	<b>Restrictions:</b> 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

You can use the LOAD statement in DB–Access to append rows to an existing table of a database.



## Notes:

- 1 See SQLS

Element	Purpose	Key Considerations
<i>column</i>	The name of a column to receive data from <i>filename</i>	<b>Restrictions:</b> Must be a column in the specified table or view.
<i>delimiter</i>	Character to use as delimiter	<b>Restrictions:</b> 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*.

---

# Chapter 12. The onmode Utility

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

Preparation for Reversion. . . . . 12-2

Syntax of the onmode -b Command . . . . . 12-2

---

## In This Chapter

- This chapter describes the **-b** option of the **onmode** utility and how to use it. You can use **onmode -b** with the following database servers:
- Dynamic Server 9.40, 9.30, 9.2x, 7.3x or 7.24
  - Dynamic Server, Linux Edition 7.3x
  - Workgroup Edition 7.3x or 7.24
  - SE
  - OnLine 5.1x

After you convert to a newer database server, you can use the **-b** option of the **onmode** utility for reversion to the older database server from which you converted. The **onmode** utility modifies the data in an Informix database so that the earlier version of the database server can access it. For information about other **onmode** options, see your *IBM Informix: Administrator's Guide*.

---

## 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 -b** command restores the databases to a format that is compatible with the earlier version. You must revert the databases before users can access the data with the earlier database server version. The utility does not revert changes made to the layout of the data that do not affect compatibility.

UNIX/Linux Only

You must be user **root** or user **informix** to execute **onmode**.

End of UNIX/Linux Only

Windows Only

You must be a member of the **Informix-Admin** group to execute **onmode**.

End of Windows Only


---

# 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



## Notes:

- 1 For all other onmode options, see your *IBM Informix: Administrator's Reference*.



Element	Purpose	Key Considerations
<b>-b 6.0</b>	Changes the database to the Version 6.0 format	<b>Additional Information:</b> See the 9.2/8.3 version of the <i>IBM Informix: Migration Guide</i> .
<b>-b 6.0A</b>	Changes the database to the Version 6.0 ALS format	<b>Additional Information:</b> See the 9.2/8.3 version of the <i>IBM Informix: Migration Guide</i> .
<b>-b 7.1</b>	Changes the database to the Version 7.10.UC1 format, which is compatible with all 7.10.UCx formats	<b>Additional Information:</b> See the 9.2/8.3 version of the <i>IBM Informix: Migration Guide</i> .
<b>-b 7.1UD1</b>	Changes the database to the Version 7.1UD1 format, which is compatible with 7.11, 7.12, 7.13, and 7.14 formats	<b>Additional Information:</b> See the 9.2/8.3 version of the <i>IBM Informix: Migration Guide</i> .
<b>-b 7.2</b>	Changes the database to the Version 7.2x format	<b>Additional Information:</b> See “Reverting from Dynamic Server 9.40 to Versions 9.30, 9.2x, 7.3x, or 7.24” on page 6-1.
<b>-b 7.3</b>	Changes the database to the Version 7.3x format	<b>Additional Information:</b> See “Reverting from Dynamic Server 9.40 to Versions 9.30, 9.2x, 7.3x, or 7.24” on page 6-1.
<b>-b 9.1</b>	Changes the database to the Version 9.14 format	<b>Additional Information:</b> See “Reverting from Dynamic Server 9.40 to Versions 9.30, 9.2x, 7.3x, or 7.24” on page 6-1.
<b>-b 9.2</b>	Changes the database to the Version 9.2x format	<b>Additional Information:</b> See “Reverting from Dynamic Server 9.40 to Versions 9.30, 9.2x, 7.3x, or 7.24” on page 6-1.
<b>-b 9.3</b>	Changes the database to the Version 9.30 format	<b>Additional Information:</b> See “Reverting from Dynamic Server 9.40 to Versions 9.30, 9.2x, 7.3x, or 7.24” on page 6-1.

**Note:** You cannot use redirection to a file when you execute **onmode -b 7.2** because this command does not function with redirection.

**Note:** To list the available options for your database server, type **onmode -b -**.



---

## Chapter 13. The onunload and onload Utilities

How onunload and onload Work . . . . .	13-1
Syntax of the onunload Command. . . . .	13-2
Destination Parameters . . . . .	13-3
Constraints That Affect onunload . . . . .	13-3
Database or Table Unloading . . . . .	13-4
Unloading a Database . . . . .	13-4
Unloading a Table . . . . .	13-4
Logging Mode . . . . .	13-5
Locking During Unload Operation. . . . .	13-5
Syntax of the onload Command . . . . .	13-5
Source Parameters . . . . .	13-6
Create Options . . . . .	13-7
Constraints That Affect onload . . . . .	13-8
Constraints That Affect onload and onunload . . . . .	13-9
Restrictions That Affect onload and onunload . . . . .	13-10
Logging During Loading . . . . .	13-10
Movement of Simple Large Objects to a BlobSpace . . . . .	13-11
Ownership and Privileges . . . . .	13-11
Exclusive Locking During Load Operation. . . . .	13-11
Steps for Using onunload and onload . . . . .	13-11

---

### In This Chapter

This chapter describes the **onunload** and **onload** utilities and how to use them. These utilities unload and load databases and tables. You can use **onunload** and **onload** with the following database servers:

- Dynamic Server 9.40, 9.30, 9.2x, 7.3x or 7.24
- Dynamic Server, Linux Edition 7.3x and Workgroup Edition 7.3x or 7.24

**Important:** You can use **onunload** and **onload** with Dynamic Server 9.40, 9.30, or 9.2x only if the databases contain only legacy data types and no extended data types.

---

### How onunload and onload Work


The **onunload** utility unloads data from a database. The **onunload** utility writes a database or table into a file on tape or disk. 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.

**Note:** You cannot use **onunload** and **onload** to move data between different versions of Dynamic Server.

The **onload** utility loads data that was created with the **onunload** command into the database server. 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 blobspace to another blobspace.

## Syntax of the **onunload** Command



### Notes:

1 See page 13-3

Element	Purpose	Key Considerations
<i>database</i>	Specifies the name of a database	<b>Additional Information:</b> The database name cannot be qualified by a database server name ( <i>database@dbservername</i> ). <b>References:</b> 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	<b>Additional Information:</b> The owner name must not include invalid characters. <b>References:</b> For pathname syntax, see your operating-system documentation.
<i>table</i>	Specifies the name of the table	<b>Restriction:</b> The table must exist. <b>References:</b> 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: Administrator's Guide*.)

## Destination Parameters

### Destination Parameters:



### Notes:

- 1 Only one occurrence of each option allowed. More than one option may occur in a single invocation.

Element	Purpose	Key Considerations
<b>-b <i>blocksize</i></b>	Specifies in kilobytes the block size of the tape device	<b>Restrictions:</b> The <i>blocksize</i> must be an integer.  <b>Additional Information:</b> This option overrides the default value in TAPEBLK or LTAPEBLK.
<b>-l</b>	Directs <b>onunload</b> to read the values for tape device, block size, and tape size from LTAPEDEV, LTAPEBLK, and LTAPESIZE, respectively	None.
<b>-s <i>tapesize</i></b>	Specifies in kilobytes the amount of data that can be stored on the tape	<b>Restrictions:</b> The <i>tapesize</i> must be an integer. If you do not specify 0, then the maximum <i>tapesize</i> is 2,097,151 kilobytes  <b>Additional Information:</b> This option overrides the default value in TAPESIZE or LTAPESIZE. To write to the end of the tape, specify a tape size of 0.
<b>-t <i>source</i></b>	Specifies the pathname of the file on disk or of the tape device where the input tape is mounted	<b>Additional Information:</b> This option overrides the tape device specified by TAPEDEV or LTAPEDEV. It must be a valid pathname.

## Constraints That Affect onunload

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 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 9.40, 9.30, or 9.2x 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.

## 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**.)

### Unloading 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.

### Unloading a Table

If you unload a table, **onunload** unloads the table data and information from the following system catalog tables:

**systables**  
**syscolumns**  
**sysindexes**  
**sysblobs**

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.

## 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 blobspace to another blobspace.

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: Administrator's Guide*.)

If you do not specify creation options, **onload** stores the database or table in the root dbspace.


## 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.

---

## Syntax of the onload Command



### Notes:

- 1 See page 13-6
- 2 See page 13-7

Element	Purpose	Key Considerations
<b>-d dbspace</b>	Loads a database or table into the specified dbspace	<b>Restriction:</b> The tape being loaded must contain the specified database or table.
<b>database</b>	Specifies the name of the database	<b>Restriction:</b> The database name cannot include a database server name, such as <i>database@dbservername</i> .  <b>References:</b> Syntax must conform to the Identifier segment; see the <i>IBM Informix: Guide to SQL Syntax</i> .
<b>owner.</b>	Specifies the owner of the table	<b>Restriction:</b> The owner name must not include invalid characters.  <b>References:</b> For pathname syntax, refer to your operating-system documentation.
<b>table</b>	Specifies the name of the table	<b>Restriction:</b> The table must not exist.  <b>References:</b> Syntax must conform to the Table Name segment; see the <i>IBM Informix: Guide to SQL Syntax</i> .

## Source Parameters

### Source Parameters:



### Notes:

- 1 Only one occurrence of each option allowed. More than one option may occur in a single invocation.





Element	Purpose	Key Considerations
<b>-c</b> <i>oldcnstrnt</i> <i>newcnstrnt</i>	Directs <b>onload</b> to rename the specified constraint.	None.
<b>-i</b> <i>oldindex</i> <i>newindex</i>	Directs <b>onload</b> to rename the table index when it stores the index on disk.	<b>Additional Information:</b> Use the <b>-i</b> option to rename indexes during the load to avoid conflict with existing index names.  <b>References:</b> Syntax must conform to the Identifier segment; see the <i>IBM Informix: Guide to SQL Syntax</i> .
<b>-fd</b> <i>olddbbsp</i> <i>newdbbsp</i>	Allows you to move a data fragment from one dbspace to another.	<b>Restriction:</b> The new dbspace must exist and not already contain another data fragment for the table.  <b>Additional Information:</b> This option is used with parallel data query (PDQ) and table fragmentation.
<b>-fi</b> <i>indexname</i> <i>olddbbs</i> <i>newdbbsp</i>	Allows you to move index fragments from one dbspace to another.	<b>Restriction:</b> The new dbspace must exist and not already contain another index fragment for the table.  <b>Additional Information:</b> This option is used with PDQ and table fragmentation.
<i>database</i>	Specifies the name of the database	<b>Restriction:</b> The database name cannot include a database server name, such as <i>database@dbservername</i> .  <b>References:</b> 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	<b>Restriction:</b> The owner name must not include illegal characters.  <b>References:</b> For pathname syntax, refer to your operating-system documentation.
<i>table</i>	Specifies the name of the table	<b>Restriction:</b> The table must not exist.  <b>References:</b> 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 the following constraints:

- The **onload** utility can only create 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. While 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 executes **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.

---

## Constraints That Affect onload and onunload

---

### Global Language Support

---



---

### Native Language Support

---

You can use **onunload** and **onload** to move data between databases if the NLS and GLS locales are identical. For example, 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** and **onunload** report errors. However, if both the NLS and GLS tables were created with the same French locale, **onload** and **onunload** would work.

---

### End of Native Language Support

---



---

### End of Global Language Support

---

The tape that **onload** reads contains binary data that is stored in disk-page-sized units. For this reason, the computer where the original database resides (where you use **onunload**) and the computer where the target database will reside (where you use **onload**) must share the following characteristics:

- The same page size
- The same representation of numeric data
- The same byte alignment for structures and unions

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 last instead of first, or different float-type representations), the contents of the data page could be misinterpreted.

## Restrictions That Affect **onload** and **onunload**

The **onload** and **onunload** utilities have the following restrictions:

- The original database and the target database must be from the same version of the database server.

### Global Language Support

- You cannot use **onload** and **onunload** to move data between non-GLS and GLS locales.

### End of Global Language Support

- Do not use **onload** and **onunload** to move data between two Dynamic Server 9.40, 9.30, or 9.2x databases if they contain extended data types. Use the HPL instead to move 9.40, 9.30, or 9.2x data. However, you can use **onload** and **onunload** with this data if the databases contain only legacy data types.

**Note:** You cannot use the **onload** and **onunload** utilities to move data from one version to another. You also cannot use these utilities to move data between different types of database servers.

## Logging During Loading

The **onload** utility performs all its loading within a transaction. This feature allows the changes to be rolled back if an error occurs.

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, 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*.)

## 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.

---

## Steps for Using **onunload** and **onload**

This section describes the procedure for using **onunload** and **onload** to move a database. You can use these commands to move either a complete database or a table from one computer to another. The syntax and description of the **onunload** utility starts on page 13-2. The syntax and description of the **onload** utility starts on page 13-5.

### 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 two kilobytes on certain UNIX systems and four kilobytes 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 two kilobytes on certain UNIX systems and four kilobytes on Windows NT.)
2. Decide where to store the unloaded data. (See step 2 of the previous section.)
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



---

## Notices

IBM may not offer the products, services, or features discussed in this document in all 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 give 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 information 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 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. 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 IBM's application programming interfaces.

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

AIX; DB2; DB2 Universal Database; Distributed Relational Database Architecture; NUMA-Q; OS/2, OS/390, and OS/400; IBM Informix®; C-ISAM®; Foundation.2000™; IBM Informix® 4GL; IBM Informix® DataBlade® Module; Client SDK™; Cloudscape™; Cloudsync™; IBM Informix® Connect; IBM Informix® Driver for JDBC; Dynamic Connect™; IBM Informix® Dynamic Scalable Architecture™ (DSA); IBM Informix® Dynamic Server™; IBM Informix® Enterprise Gateway Manager (Enterprise Gateway Manager); IBM Informix® Extended Parallel Server™; i.Financial Services™; J/Foundation™; MaxConnect™; Object Translator™; Red Brick™; IBM Informix® SE; IBM Informix® SQL; InformiXML™; RedBack®; SystemBuilder™; U2™; UniData®; UniVerse®; wintegrate® are trademarks or registered trademarks of International Business Machines Corporation.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

Windows, Windows NT, and Excel are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Other company, product, and service names used in this publication may be trademarks or service marks of others.

---

# Index

## A

Abbreviated years 8-1  
AFDEBUG environment variable 3-3  
AFF\_NPROCS configuration parameter 3-7  
AFF\_SPROC configuration parameter 3-7  
ALARMPROGRAM configuration parameter 3-7, 5-12, 5-13  
ALL keyword, with dbschema utility 10-3  
ALLOW\_NEWLINE configuration parameter 3-6  
ANSI compliance  
    level xx  
Applications, client  
    *See also* Client applications.  
    Communications Support Module 5-13  
ASCII file, produced by unloading 2-4  
ASF\_SOCTCP\_BACKLOG configuration  
    parameter 3-4  
Authentication policy 5-13, 6-8

## B

Backing up  
    Dynamic Server 9.40 5-16, 6-6  
    logical logs 5-13  
    ON-Bar 5-9, 5-16, 6-9  
    ontape utility 5-9, 5-16, 6-9  
    source database 5-9  
    source database server 5-9  
Binary files, loading 2-5, 2-7, 13-9  
BladeManager  
    installing and registering DataBlade modules 5-13  
    removing extensions during reversion 6-6  
Blobs, moving with onload 13-13  
Blobspace  
    backing up 3-2  
    moving, with onunload and onload 13-13  
Boldface type xiii  
boot90.sql and boot901.sql 6-1  
Bringing down  
    Dynamic Server 5-6, 5-9  
Bringing up  
    Dynamic Server 9.40 5-13  
    target database server after reversion 6-8

## C

Case-sensitive name space 3-16  
CDR\_DBSPACE configuration parameter 3-4  
CDR\_ENV configuration parameter 3-4  
CDR\_LOGBUFFERS configuration parameter 3-7  
CDR\_LOGDELTA configuration parameter 3-7

CDR\_LOGDELTA environment variable 3-3  
CDR\_MAX\_DYNAMIC\_LOGS configuration  
    parameter 3-4  
CDR\_NIFRETRY configuration parameter 3-7  
CDR\_NUMCONNECT configuration parameter 3-7  
CDR\_PERFLOG environment variable 3-3  
CDR\_QDATA\_SBFLAGS configuration parameter 3-7  
CDR\_QDATA\_SBSPACE configuration parameter 3-5, 3-7  
CDR\_QHDR\_DBSPACE configuration parameter 3-5  
CDR\_RMSCALEFACT environment variable 3-3  
CDR\_ROUTER environment variable 3-3  
Character-position form of FILE and INSERT  
    statements 9-10  
Checking available space  
    Dynamic Server 9.40 5-3  
Checking database integrity 5-16  
Choosing a migration method 1-1  
Chunk format 3-14  
Chunk reserve pages  
    and reversion from Dynamic Server 6-6  
Client applications  
    Communications Support Module 5-13  
Client locale.  
    *See* Locale.  
Code-set conversion, HPL 2-12  
Code, sample, conventions for xvii  
Command file, dbload 9-5  
Command-line conventions  
    how to read xvi  
    sample diagram xvi  
Commands  
    *See* SQL statement.  
Communications Support Module  
    client applications 5-13  
    configuring 5-13  
    removing 6-8  
    saving configuration 6-6  
Compliance  
    with industry standards xx  
conscm.cfg file 5-13  
    creating 5-13  
    deleting 6-8  
    saving 6-6  
Configuration file  
    customizing 5-12  
    replacing after reversion from Dynamic Server  
        9.40 6-8  
    saving a copy of current version 6-6

## Configuration file *(continued)*

- saving during conversion 5-5
- Configuration parameter 3-5
  - ALARMPROGRAM 5-12, 5-13
  - ALLOW\_NEWLINE 3-6
- altered in Dynamic Server 9.30
  - JDKVERSION 3-7
  - JVPJAVAHOME 3-7
  - JVPJAVALIB 3-7
  - JVPJAVAVM 3-8
- altered in Dynamic Server 9.40
  - ALARMPROGRAM 3-7
  - CDR\_QDATA\_SBSPACE 3-7
  - DBSERVERALIASES 3-7
  - LRU\_MAX\_DIRTY 3-7
  - LRU\_MIN\_DIRTY 3-7
  - LTAPEBLK 3-7
  - LTAPESIZE 3-7
  - OPTICAL\_LIB\_PATH 3-7
  - TAPEBLK 3-7
  - TAPESIZE 3-7
- ASF\_SOCTCP\_BACKLOG 3-4
- CDR\_DBSPACE 3-4
- CDR\_ENV 3-4
- CDR\_MAX\_DYNAMIC\_LOGS 3-4
- CDR\_QDATA\_SBSPACE 3-5
- CDR\_QHDR\_DBSPACE 3-5
- DD\_HASHMAX 3-6
- DD\_HASHSIZE 3-6
- DEF\_TABLE\_LOCKMODE 3-5
- DS\_HASHSIZE 3-5
- DS\_POOLSIZE 3-6
- DYNAMIC\_LOGS 3-5
- ENCRYPT\_CDR 3-4
- ENCRYPT\_CIPHER 3-5
- ENCRYPT\_MAC 3-5
- ENCRYPT\_MACFILE 3-5
- ENCRYPT\_SWITCH 3-5
- HPL\_DYNAMIC\_LIB\_PATH 3-5, 5-12
- LTAPEBLK 3-18
- new in Dynamic Server 9.20 3-6
- new in Dynamic Server 9.21 3-5
- new in Dynamic Server 9.30 3-5
- new in Dynamic Server 9.40 3-4
- OPT\_GOAL 3-6
- OPTICAL\_LIB\_PATH 5-12
- PC\_HASHSIZE 3-6
- PC\_POOLSIZE 3-6
- PLOG\_OVERFLOW\_PATH 3-5
- removed in Dynamic Server 9.30
  - AFF\_NPROCS 3-7
  - AFF\_SPROC 3-7
  - CDR\_LOGBUFFERS 3-7
  - CDR\_LOGDELTA 3-7
  - CDR\_NIFRETRY 3-7

## Configuration parameter *(continued)*

- removed in Dynamic Server 9.30 *(continued)*
  - CDR\_NUMCONNECT 3-7
  - DRAUTO 3-7
  - LBU\_PRESERVE 3-8
  - LOGSMAX 3-8
  - LTXEHWM 3-8
  - LTXHWM 3-8
  - NOAGE 3-8
  - NUMAIOVPS 3-8
  - NUMCPUVPS 3-8
- removed in Dynamic Server 9.40
  - CDR\_QDATA\_SBFLAGS 3-7
- ROOTOFFSET 5-12
- ROOTPATH 5-12
- ROOTSIZE 5-12
- SBSPACETEMP 3-5
- STMT\_CACHE 3-6
- STMT\_CACHE\_HITS 3-6
- STMT\_CACHE\_NOLIMIT 3-6
- STMT\_CACHE\_NUMPOOL 3-6
- STMT\_CACHE\_SIZE 3-6
- SYSSBSPACENAME 3-6
- TAPEBLK 3-18

## Contact information xx

## Conventions

- documentation xii
- syntax diagrams xiv
- syntax notation xiv

## Converting

- from Dynamic Server 7.31 to Dynamic Server 9.40
  - with Enterprise Replication 4-1, 4-2
- from Dynamic Server 9.2x to Dynamic Server 9.40 5-1
- from Dynamic Server 9.2x to Dynamic Server 9.40
  - with Enterprise Replication 4-2
- from Dynamic Server 9.30 to Dynamic Server 9.40 5-1
- from Dynamic Server 9.30 to Dynamic Server 9.40
  - with Enterprise Replication 4-2
- monitoring status with online.log 5-14
- to Dynamic Server 9.40 5-1, 5-17
- to Dynamic Server 9.40 with Enterprise Replication 4-1

## CREATE PROCEDURE statement, use of dbschema 10-4

## CREATE SEQUENCE statement 10-4

## CREATE SYNONYM statement, use of dbschema 10-3

## CREATE TABLE statement

- use of dbschema 10-4

## CREATE VIEW statement, use of dbschema 10-4

## Creating a dbload command file 9-5

## Customizing configuration file 5-12



## D

Data integrity 5-15

Data type

user-defined

converting from Dynamic Server 9.2x 4-2

Database

ownership, set by onload 13-11

verifying integrity of 5-16

Database locale.

See Locale.

Database server

platforms and versions 1-1

Datablade API

enhancements 3-21

DataBlade API

enhancements 3-20

DataBlade modules

installing and registering 5-13

dbexport

SELECT triggers, disabling 8-1

dbexport utility 8-1

-c option 8-3, 8-4

-d option 8-3

-q option 8-3

-ss option 8-3, 8-4

-V option 8-3

-X option 8-3

description of 2-9, 8-1

destination options 8-4

schema output 8-6

syntax 8-2

dbimport utility 8-1

-c option 8-8, 8-9

-l option 8-13

-q option 8-8

-V option 8-8

-X option 8-8

changing the locale 8-14

create options 8-11

database logging mode 8-13

description of 2-9, 8-1

importing from another computer 2-14

input file location options 8-9

input-file location options 8-9

Interrupt key 8-9

syntax 8-7

using with GLS 8-8

using with NLS 8-14

dbload utility

-c command file option 9-1

-d database option 9-2

-e errors option 9-2

-e option 9-3

-i ignore rows option 9-2

-i option 9-3

dbload utility (*continued*)

-k option 9-2

-l error log file option 9-2

-n commit interval option 9-2

-p option 9-2

-r option 9-2, 9-3

-s option 9-2

-V option 9-3

-X option 9-3

compared to LOAD 2-10

creating a command file 9-5

description of 9-1

FILE statement 9-5

ignoring rows 9-3

importing from another computer 2-14

INSERT statement

compared to SQL INSERT statement 9-12

using 9-5, 9-7

Interrupt key 9-4

moving data to Workgroup Edition 7-4

number errors to allow 9-3

speed, increasing 9-4

syntax 9-1

table locking 9-3

using 7-3

writing a command file

in character-position form 9-12

in delimiter form 9-8

dbschema utility

-ss option 10-5

create schema for a database 6-3, 10-3

description of 2-11, 10-1

distribution information 10-10

moving data to Workgroup Edition 7-4

obtaining

privilege schema 10-7

sequence schema 10-6

synonym schema 10-7

owner conventions 10-4

specifying a table, view, or procedure 10-8

specifying the role schema 10-10

using 7-3

DBSERVERALIASES configuration parameter 3-7

Dbospace

backing up 3-2

moving tables to another dbospace 13-13

DD\_HASHMAX configuration parameter 3-6

DD\_HASHSIZE configuration parameter 3-6

DEF\_TABLE\_LOCKMODE configuration

parameter 3-5

Default locale xii

Delimiter form of FILE and INSERT statements 9-6

Dependencies, software x

Detached indexes 3-14

- Diagnostics table, displayed by dbschema utility 2-11, 10-13
  - Directory, installation 5-10
  - Disabling
    - High-Availability Data Replication 5-8
  - Display schema for a database 10-3
  - Distributed transactions, Limitation on using routines 3-16
  - Distribution information for tables 10-10
  - Documentation notes xix
  - Documentation, types of
    - documentation notes xix
    - machine notes xix
    - online manuals xviii
    - release notes xix
  - DRAUTO configuration parameter 3-7
  - DRDA database 2-14
  - DS\_HASHSIZE configuration parameter 3-5
  - DS\_POOLSIZE configuration parameter 3-6
  - Dummy UPDATE statement 5-6
  - Dynamic Server
    - platforms and versions 1-1
    - reverting 6-9
    - reverting from current version 6-2
  - Dynamic Server 7.3x
    - moving data between computers 2-8
    - using dbexport and dbimport 2-8
  - Dynamic Server 9.40
    - authentication policies 5-13, 6-8
    - backing up 5-16, 6-6
    - bringing down 5-9
    - bringing up 5-13
    - DataBlade modules, installing 5-13
    - initializing 5-13
    - installing 5-10
    - performance tuning 5-16
    - release notes 3-2
    - saving configuration file 6-6
    - space requirements 5-3
    - starting 5-13
  - Dynamic Server 9.x
    - importing data 2-14
    - moving data between computers 2-8, 2-14
    - onload and onunload 2-7
    - using dbexport and dbimport 2-8
    - using dbimport 2-14
    - using dbload 2-14
    - using onload and onunload 2-7
  - DYNAMIC\_LOGS configuration parameter 3-5
- ## E
- en\_us.8859-1 locale xii
    - See U.S. English locale.
  - ENCRYPT\_CDR configuration parameter 3-4
  - ENCRYPT\_CIPHER configuration parameter 3-5
  - ENCRYPT\_MAC configuration parameter 3-5
  - ENCRYPT\_MACFILE configuration parameter 3-5
  - ENCRYPT\_SWITCH configuration parameter 3-5
  - Enterprise Gateway Manager 2-15
  - Enterprise Gateway with DRDA 2-14
  - Enterprise Replication
    - converting to Dynamic Server 9.40 4-1
    - enhancements 3-19, 3-21
  - Environment variables
    - AFDEBUG 3-3
    - CDR\_LOGDELTA 3-3
    - CDR\_PERFLOG 3-3
    - CDR\_RMSCALEFACT 3-3
    - CDR\_ROUTER 3-3
    - CLIENT\_LOCALE 8-2
    - DB\_LOCALE 8-2, 8-14
    - DBCENTURY 8-1
    - DBDATE 8-1
    - DBTEMP 8-14
    - IFX\_DEF\_TABLE\_LOCKMODE 3-3
    - IFX\_LONGID 3-4
    - IFX\_UPDDESC 3-4
    - INFORMIXSERVER 5-10
    - INFORMIXSQLHOSTS 4-1, 5-10
    - JAR\_TEMP\_PATH 3-3
    - JAVA\_COMPILER 3-3
    - JVM\_MAX\_HEAP\_SIZE 3-4
    - new in Dynamic Server 9.20 3-4
    - new in Dynamic Server 9.21 3-3
    - new in Dynamic Server 9.30 3-3
    - new in Dynamic Server 9.40 3-3
    - ONCONFIG 5-10
    - PATH 5-10
    - reverting Dynamic Server 9.40 6-8
    - SERVER\_LOCALE 8-2
    - STMT\_CACHE 3-4
    - TEMP 8-14
    - TMP 8-14
    - USETABLENAME 3-3
  - Example, calculating free space 5-4
  - Extents, checking 5-16
  - Extracting schema information 7-1
- ## F
- Fast recovery, initiating 5-7
  - Features of
    - Dynamic Server 9.40, new 3-21
  - FILE statement
    - character-position form 9-10
    - delimiter form 9-6
    - syntax for
      - character-position form 9-11
      - delimiter form 9-6
    - with dbload 9-5

## Files

- oncfg 5-5
- onconfig 6-6
- ONCONFIG 5-5, 5-12
- sqlhosts 5-5, 5-11
- ttermcap
  - saving 5-5
- termcap 5-5
- finderr utility xx
- Formula
  - free space 5-4
- Free space required 5-4

## G

- Gateway, importing from another computer 2-14
- Global Language Support (GLS) xi
  - and dbimport utility 8-8
  - using onload and onunload 2-7, 13-9
- GLS
  - new features 3-19
- GLS locales 8-2
- GLS.
  - See Global Language Support.
- GRANT statement
  - role privileges 10-8
  - use of dbschema 10-4

## H

- Help xviii
- High-Availability Data Replication
  - disabling 5-8
  - disabling during reversion 6-7
  - enabling 5-16
- High-Performance Loader (HPL)
  - custom-code shared library 3-15
  - ipload utility 2-12
  - loading ASCII or COBOL data 2-12
  - onpladm utility 2-12
  - using 2-12
- HPL\_DYNAMIC\_LIB\_PATH configuration
  - parameter 3-5, 5-12
- HPLAPIVERSION 3-5
- HPLAPIVERSION configuration parameter 3-5

## I

- IBM Informix Server Administrator 3-17
- IBM Informix Storage Manager 3-27
  - migrating catalogs 3-28
- IECC.
  - See IBM Informix Enterprise Command Center.
- IFX\_DEF\_TABLE\_LOCKMODE environment
  - variable 3-3
- IFX\_LONGID environment variable 3-4
- IFX\_UPDESC environment variable 3-4
- Importing non-Informix data 2-14

## In-Place alters

- finding 5-6
- removing 5-6

## Index

- and dbload utility 9-3
- checking 5-16
- rebuilding 5-14

## Indexes, detached 3-14

- Industry standards, compliance with xx
- INFORMIXDIR directory 5-10
- INFORMIXDIR/bin directory xii
- INFORMIXSERVER environment variable 5-10
- INFORMIXSQLHOSTS environment variable 5-10
- Initializing Dynamic Server 9.40 5-13

## INSERT statement

- character-position form 9-10
- delimiter form 9-6
- syntax for
  - character-position form 9-11
  - delimiter form 9-6
- with dbload 9-5

## Installation directory 5-10

## Installation enhancements 3-21

## Installing

- Dynamic Server 9.40 5-10
- storage manager 3-27

## Interrupt key

- with dbimport 8-9
- with dbload 9-4

## ipload utility 2-12

## ISM.

- See IBM Informix Storage Manager.

## ISO 8859-1 code set xii

## J

- JAR\_TEMP\_PATH environment variable 3-3
- JAVA\_COMPILER environment variable 3-3
- JDKVERSION configuration parameter 3-7
- JVM\_MAX\_HEAP\_SIZE environment variable 3-4
- JVPJAVAHOME configuration parameter 3-7
- JVPJAVALIB configuration parameter 3-7
- JVPJAVAVM configuration parameter 3-8

## K

## Kernal parameters on UNIX or Linux 5-9

## Keywords

- in syntax diagrams xvii

## L

## LBU\_PRESERVE configuration parameter 3-8

## Level-0 backup 5-9, 6-9

- after moving data 13-12
- Dynamic Server 9.40 5-16

## Linux

- migrating on 1-4

- LOAD SQL statement
  - syntax 11-2
- Loading
  - ASCII files 2-4, 2-14
  - binary data 2-5, 2-7, 13-9
  - COBOL data 2-4, 2-12
- Locale xii
- Locking, set by onload 13-11
- Logging, and
  - onload and onunload 13-12
- Logical logs
  - backing up 5-13
  - out of space 5-13
  - restoring 5-16
  - space required for Dynamic Server 9.40 5-3
- LOGSMAX configuration parameter 3-8
- LRU\_MAX\_DIRTY configuration parameter 3-7
- LRU\_MIN\_DIRTY configuration parameter 3-7
- LTAPEBLK configuration parameter 3-7
- LTAPEBLK parameter 3-18
- LTAPEDEV parameter, with onunload/onload 13-12
- LTAPESIZE configuration parameter 3-7
- LTXEHWM configuration parameter 3-8
- LTXHWM configuration parameter 3-8

## M

- Machine notes xix
- Message file for error messages xx
- Migrating
  - between 32-bit and 64-bit database servers 3-28
- Migrating a database
  - using dbexport and dbimport 7-3
  - using UNLOAD, dbschema, and LOAD 7-3
- Migrating a database to
  - overview 2-1
  - same platforms 2-5
  - Workgroup Edition 7-4
- Migration
  - before you begin 3-1
  - on Linux 1-4
  - on same operating system 1-4
  - on UNIX 1-4
  - on Windows 2000 1-5
  - on Windows NT 1-5
  - on Windows XP 1-5
  - onload utility 13-2, 13-5
  - onunload utility 13-1
  - paths 1-1, 1-7
- Migration procedure
  - between UNIX and Windows NT 7-2
  - general guidelines for 3-2
- Mode
  - checking 5-8
- Monitoring conversion status 5-14
- Moving data
  - blobspaces 13-13
  - choosing a migration method 1-1
  - constraints 2-4
  - oncheck utility 13-12
  - steps for using onunload and onload 13-11

## N

- Name space
  - case-sensitive 3-16
- Native Language Support (NLS)
  - populating with dbimport 8-14
- New features
  - 7.30 features in Dynamic Server 9.20 3-26
  - Java Virtual Machine 1.3 support 3-22
  - MaxConnect support in Version 9.21 3-23
  - of Dynamic Server 9.20 3-23, 3-25
  - of Dynamic Server 9.21 3-22
  - of Dynamic Server 9.30 3-21
  - of Dynamic Server 9.40 3-18, 3-21
  - Universal Server 9.1x features in Dynamic Server 9.20 3-25
- NFS.
  - See Network File System.
- NLS.
  - See Native Language Support.
- NOAGE configuration parameter 3-8
- Non-Informix data, importing 2-14
- NUMAIOVPS configuration parameter 3-8
- NUMCPUVPS configuration parameter 3-8

## O

- ON-Bar 5-5, 5-15
- ON-Bar utility
  - backing up
    - Dynamic Server 9.40 5-16
  - backing up Dynamic Server 5-9
  - backing up Dynamic Server 9.40 5-16, 6-6
- ON-Bar Utility
  - command changes between Version 7.x and 9.x 3-14
- oncfg file in \$INFORMIXDIR/etc 5-5
- oncfg file in %INFORMIXDIR%/etc 5-5
- oncheck utility
  - cc database\_name option 5-8, 5-16
  - cD database\_name option 5-8, 5-16
  - ce option 5-8, 5-16
  - cI database\_name option 5-8, 5-16
  - cr option 5-8, 5-16
  - cs sbspace\_name option 5-16
  - cS sbspace\_name option 5-16
  - before moving data 13-12
  - rebuilding table indexes 5-14
  - verifying database integrity 5-8, 5-15
- ONCONFIG environment variable 5-10

- onconfig file
  - saving 6-6
- ONCONFIG file 5-5, 5-12
- ONCONFIG file.
  - See Configuration file.
- ONCONFIG parameters.
  - See Configuration parameter.
- oninit utility
  - s option 5-7
  - starting Dynamic Server 9.40 5-14
- Online help xviii
- Online manuals xviii
- online.log 5-14
- onload and onunload utilities 2-7, 13-13
- onload utility
  - constraints on use 13-9
  - create options 13-5, 13-7
  - logging status 13-10
  - moving a database 13-11
  - moving a table 13-12, 13-13
  - moving locales 2-7, 13-9
  - moving to another dbspace 13-13
  - ownership and privileges 13-11
  - specifying tape parameters 13-6
  - steps for using 13-11
  - syntax 13-5
  - using between computers 2-5, 2-7
- onmode utility
  - b option 6-7, 12-2
  - d standard option 5-8
  - ky option 5-6
  - m option 6-9
  - sy option 5-6
  - yuk option 5-7
  - reverting from Dynamic Server 9.40 6-7
  - shutting down 5-6
  - shutting down Dynamic Server 5-6
- onmode-b command 12-1
- onpladm utility 2-12
- onstat
  - See onstat utility.
- onstat utility 5-8
- ontape utility
  - a option 5-13
  - backing up Dynamic Server 9.40 5-16, 6-6
  - backing up source database server 5-9
- onunload utility
  - constraints on use 13-3, 13-9
  - locking 13-5
  - logging mode 13-5
  - moving a database 13-11
  - moving a table 13-12, 13-13
  - moving locales 2-7, 13-9
  - moving to another dbspace 13-13
  - ownership and privileges 13-4
- onunload utility (*continued*)
  - steps for using 13-11
  - syntax 13-2
  - system catalog tables 13-4
  - use without tape parameters 13-2
  - using 2-5, 2-7
  - what is included with a
    - database 13-4
    - table 13-4
- Operating system
  - locale.
    - See Global language support.
  - reconfiguring 5-9
- OPT\_GOAL 3-6
- OPT\_GOAL configuration parameter 3-6
- Optical storage manager library 3-15
- OPTICAL\_LIB\_PATH configuration parameter 3-7, 5-12
- Owner, in dbschema 10-4

## P

- PATH environment variable 5-10
- PC\_HASHSIZE configuration parameter 3-6
- PC\_POOLSIZE configuration parameter 3-6
- Performance comparison 5-16
- Performance tuning 5-16
- Planning
  - data migration overview 2-1
- Platforms, moving data between
  - compatible computers 2-5, 2-8, 2-14
- PLOG\_OVERFLOW\_PATH configuration
  - parameter 3-5
- Privileges 10-8
  - access privileges for tables 13-4
  - required for onload 13-11
  - required for onunload 13-4
- Program group
  - Documentation notes xix
  - Release notes xix

## Q

- Quiescent mode 5-8

## R

- Reconfiguring the operating system 5-9
- Registering DataBlade modules 5-13
- Release notes xix
  - Dynamic Server 9.40 3-2
  - location of 3-2
- Reserve pages, checking 5-16
- Reversion utility 6-7
- Reverting
  - chunk reserve pages 6-6
  - determining if it is possible 6-2
  - disabling High-Availability Data Replication 6-7
  - from Dynamic Server 9.40 6-1, 6-2, 6-9

## Reverting (*continued*)

- from Dynamic Server 9.40 with Enterprise Replication to Dynamic Server 7.31 4-3, 4-4
  - from Dynamic Server 9.40 with Enterprise Replication to Dynamic Server 9.2x 4-3, 4-4
  - from Dynamic Server 9.40 with Enterprise Replication to Dynamic Server 9.30 4-3
  - removing 9.40 features 6-7
  - restrictions for reverting from Dynamic Server 9.40 6-6
  - restrictions for reverting to prior versions 6-3
  - using the onmode -b command 12-1
- ROOTOFFSET configuration parameter 5-12
- ROOTPATH configuration parameter 5-12
- ROOTSIZE configuration parameter 5-12
- Running the reversion utility 6-7

## S

- Sample-code conventions xvii
- SBSAPACETEMP configuration parameter 3-5
- SELECT triggers, disabling with dbexport 8-1
- SERIAL data type, treatment by dbschema 10-4
- Server Studio JE 3-16
- Simple large objects, moving with onload 13-13
- sm\_versions.std file
- renaming to sm\_versions 5-15
- Software dependencies x
- Space

- checking what is available 5-3
- configuring 5-3
- for sysmaster database 5-3

## SQL

enhancements 3-19

## SQL code xvii

## SQL reserved words 3-8

## SQL statement

- CREATE SEQUENCE 10-4
- FILE 9-5
- INSERT 9-5
- UPDATE STATISTICS
- data distributions 10-10

## SQLHOSTS Connectivity Information 3-17

## sqlhosts file, UNIX

- changing name or path 5-11
- csn option 5-13, 6-8
- save a copy when migrating 5-5

## Starting

- Dynamic Server 9.40 5-13
- target database server after reversion 6-8

## STMT\_CACHE configuration parameter 3-6

## STMT\_CACHE environment variable 3-4

## STMT\_CACHE\_HITS configuration parameter 3-6

## STMT\_CACHE\_NOLIMIT configuration parameter 3-6

## STMT\_CACHE\_NUMPOOL configuration

- parameter 3-6

## STMT\_CACHE\_SIZE configuration parameter 3-6

## Storage manager 3-27

## Stored Procedure Parameter Limit 3-15

## stores\_demo database xii

## superstores\_demo database xii

## Syntax diagrams

- conventions for xiv
- keywords in xvii
- variables in xvii

## sysindexes

- changes to 3-10

## sysmaster database 5-3

- and logical logs 5-13

- changes to 3-9, 3-13

- space required to build 5-3

- when created 7-4

## SYSSBSPACENAME configuration parameter 3-6

## System catalogs

- boot scripts 6-1
- changes to 3-9, 3-10
- checking tables 5-16

## System requirements

- database x
- software x

## sysutils database 7-4

## T

## TAPEBLK configuration parameter 3-7

## TAPEBLK parameter 3-18

## TAPEDEV parameter, with onunload and onload 13-12

## TAPESIZE configuration parameter 3-7

## tctermcap file

- saving 5-5

## termcap file in \$INFORMIXDIR/etc 5-5

## Tools, migration 2-4

## Transactions

- checking for open ones 5-7

## Tuning

- Dynamic Server 9.40 5-16

## U

## ulimit 8-4

## UNIX

- migrating on 1-4

## UNLOAD SQL statement

- moving data 7-4

- syntax 2-9, 11-1

- using 7-3

## UPDATE statement

- dummy 5-6

## UPDATE STATISTICS statement

- data distributions 10-10

- UPDATE STATISTICS statement (*continued*)
  - using after a successful conversion 5-15
  - using after reversion 6-8
- Upgrading, ISM 3-27
- Upgrading,
  - See Converting.
- User-defined data type, converting from Dynamic Server 9.2x 4-2
- Users, types of x
- USETABLENAME environment variable 3-3
- Utilities
  - dbexport 2-9, 7-1, 8-1
  - dbexport syntax 8-2
  - dbimport 2-9, 7-1, 8-1
  - dbimport syntax 8-7
  - dbload 7-1, 9-1
  - dbschema 2-11, 7-1, 10-1
  - onload 7-2
  - onload and onunload 2-7, 13-13
  - onpladm 2-12
  - onpload 2-12
  - onunload 7-2
- Utilities for data migration
  - dbexport 7-1, 8-1
  - dbimport 7-1, 8-1
  - dbload 7-1, 9-1
  - dbschema 7-1, 10-1
  - onload 7-2, 13-6
  - onunload 7-2, 13-2

## V

- Validated storage manager 3-27
- VARCHAR Column Limit 3-15
- Variables, in syntax diagrams xvii
- Verifying data integrity 5-15
- Verifying storage-manager validation 3-27
- Versions
  - database servers 1-1
- View, display description with dbschema 10-3

## W

- Windows 2000
  - migrating 1-5
- Windows NT
  - migrating 1-5
- Windows XP
  - migrating on 1-5
- Workgroup Edition 7.24
  - moving data between computers 2-8
  - using onload and onunload 2-8
- Workgroup Edition 7.3x
  - moving data between computers 2-8
  - using onload and onunload 2-8
- Writing a dbload command file
  - in character-position form 9-12

- Writing a dbload command file (*continued*)
  - in delimiter form 9-8

## X

- X/Open compliance level xx









Printed in USA

G251-2440-00

