

**HCL Informix 15.0.0**

**Informix Migration Guide**



# Contents

<b>Chapter 1. Migrating Informix® database systems.....</b>	<b>3</b>
Overview of migration.....	3
Overview of moving data.....	3
Migration to and reversion from 15.0.0.....	8
Preparing for migration to Version 15.0.0.....	8
and Upgrading.....	22
High-availability cluster migration.....	25
Migrating to 15.0.0.....	45
Reverting from Version 15.0.0.....	59
Migration of data between database servers.....	69
Migrating database servers to a new operating system.....	69
Data migration utilities.....	71
External tables.....	71
The dbexport and dbimport utilities.....	72
The dbload utility.....	88
The dbschema utility.....	103
The LOAD and UNLOAD statements.....	120
The onunload and onload utilities.....	122
The onmode utility reversion option.....	134
The onrestorept utility.....	135
<b>Index.....</b>	<b>137</b>

# Chapter 1. Migrating Informix® database systems

The *HCL Informix® Migration Guide* describes how to move data manually between databases, servers, and computers.

These topics are intended for database server administrators or database administrators who are responsible for upgrading the database server or migrating data. These topics assume that you have the following background:

- A working knowledge of your computer, your operating system, and the utilities that your operating system provides
- Some experience with database server administration, operating-system administration, or network administration

## Overview of migration

### Overview of moving data

If you are installing the new version of the database server on another computer or operating system (non-in-place migration), you can use one of several tools and utilities to move data from your current database server.

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

Before you move data, consider these issues:

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

For information about how to move data between database servers on different operating systems, also see [Migrating database servers to a new operating system on page 69](#).

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

### Prerequisites before moving data

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

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

### Data-migration tools

Informix® provides tools, utilities, and SQL statements that you can use to move data from one HCL Informix® database to another or from one operating system to another.

You might want to use a data-migration tool when you have different page sizes or code pages. For example, UNIX™ or Linux and Windows store data in different page sizes.

When your migration involves migrating between different operating systems, you must export data and its schema information from one database server and import the exported data into the other database server.

Normally, if you are migrating on the same operating system, you do not need to load and unload data.

You can use the following tools to move data:

- The **dbexport** and **dbimport** utilities
- The **dbload** utility
- The **onunload** and **onload** utilities
- UNLOAD and LOAD statements
- Nonlogging raw tables

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

- The **dbimport** and **dbload** utilities
- External tables that you create with the CREATE EXTERNAL TABLE statement

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



**Note:** **dbimport** cannot be used to move datablade data between Informix versions.

**Table 1. Comparison of tools for moving data**

Tool	Description	Advantages	Disadvantages
<b>dbexport</b> and <b>dbimport</b> utility	Imports or exports a database to a text file that is stored on disk or tape	Can modify the database schema and change the data format Can move data between operating systems Optional logging Can import data from non-Informix® sources	Faster performance than the <b>dbload</b> utility, but slower performance than the <b>onload</b> utility Moves the entire database
<b>dbload</b> utility	Transfers data from one or more text files into	Can modify database schema	Slower performance than the <b>dbexport</b> , <b>dbimport</b> , and <b>onload</b> utilities

Table 1. Comparison of tools for moving data (continued)

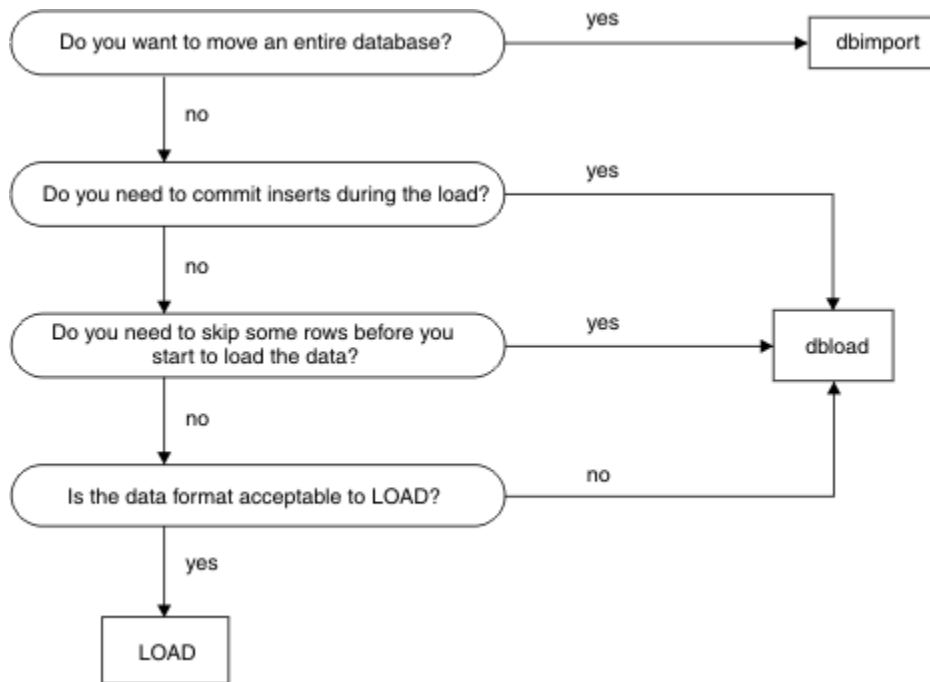
Tool	Description	Advantages	Disadvantages
	one or more existing tables	Can move data between operating systems  Optional logging  Moderately easy to use  Can import data from non-Informix® sources	
<b>onunload</b> and <b>onload</b> utilities	Unloads data from a database into a file on tape or disk; loads data, which was created with the <b>onunload</b> command, into the database server	Fast performance Optional logging	Only moves data between database servers of the same version on the same operating system  Cannot modify the database schema  Logging must be turned off  Difficult to use
UNLOAD and LOAD statements	Unloads and loads specified rows	Can modify database schema  Can move data between operating systems  Easy to use  Optional logging	Only accepts specified data formats
Nonlogging raw tables	Loads certain kinds of large tables	Can load very large data warehousing tables quickly	Does not support primary constraints, unique constraints, and rollback  Requires SQL  Not recommended for use within a transaction
External tables	Enables you to read and write from a source that is external to	Performs express (high-speed) and deluxe (data-checking) transfers	Requires SQL

**Table 1. Comparison of tools for moving data (continued)**

Tool	Description	Advantages	Disadvantages
	the database server, providing an SQL interface to data in text files managed by the operating system or to data from a FIFO device.		

If you are choosing a tool for loading data, the questions shown in [Figure 1: Choosing among dbimport, dbload, and LOAD on page 6](#) will help you make a decision.

**Figure 1. Choosing among dbimport, dbload, and LOAD**



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

## When TEXT and BYTE data is scanned, not compressed

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

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

## Moving data between computers and dbspaces

You can move data between different computers, and you can import data from environments other than the Informix® database server. Except when you use external tables, you must unload your data to ASCII files before you move the data to another computer.

If you are moving data into the 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 an application that is not based on Informix®, you might need to use the UNLOAD statement because you can specify the delimiter that is used in the data files.

## Importing data from a non-Informix® source

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

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

In addition to **dbimport** and **dbload**, HPL provides ways to access information from non-Informix® sources.

## Moving data by using distributed SQL

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



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

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

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

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

For example:

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

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

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

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

## Migration to and reversion from 15.0.0

### Preparing for migration to Version 15.0.0

Before you install the new version of Informix®, you must prepare the database server environment for migration by performing specified pre-migration tasks. If you are also migrating from 32-bit to 64-Bit database servers, you must perform additional tasks.

### Preparing for migration

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

#### Before you begin

- Check the [IBM® Support Portal](#) for the latest patches that you must install before you migrate or upgrade Informix® software.
- If you use Enterprise Replication, you must first prepare your replication environment for migration. For more information, see [and Upgrading on page 22](#).

#### About this task

Review and complete all tasks that apply:

1. [Reviewing changes in Informix product functionality on page 9](#).
2. [Checking and configuring available space on page 9](#).
3. [Configuring for recovery of restore point data in case an upgrade fails on page 11](#).



4. Renaming user-defined routines (UDRs) that have the following names: CHARINDEX() , LEFT(), RIGHT(), INSTR(), DEGREES(), RADIANS(), REVERSE(), SUBSTRING\_INDEX(), LEN(), and SPACE().

These names are reserved for built-in SQL string manipulation functions.

5. Adjusting settings:

- a. If you use UNICODE, ensure that the GL\_USEGLU environment variable on the source server is set to the same value as the GL\_USEGLU environment variable on the target server.
- b. If the source version of the database server contains the IFX\_EXTEND\_ROLE configuration parameter, which controls authorization to register modules or external UDRs, disable the parameter by setting it to 0 (off).

6. [Saving copies of the current configuration files on page 12.](#)
7. [Preparing 12.10 BSON columns with DATE fields for upgrade on page 13.](#)
8. [Resolving outstanding in-place alter operations on page 14.](#)
9. [Closing all transactions and shutting down the source database server on page 16.](#)
10. [Initiating fast recovery to verify that no open transactions exist on page 16.](#)
11. [Verifying the integrity of the data on page 17.](#)
12. [Verifying that the database server is in quiescent mode on page 18.](#)
13. [Making a final backup of the source database server on page 18.](#)



**Important:** Complete the previous step in case you have to revert to the source database server.

14. [Verifying that the source database server is offline on page 18.](#)

## What to do next

If you use high-availability clusters, you must complete additional preparations. See [High-availability cluster migration on page 25.](#)

## Reviewing changes in Informix® product functionality

Changes to Informix® product functionality might affect your plans for migrating to the latest version of the product.

### About this task

Changes in functionality in Informix® 15.0.0 can potentially impact your applications, scripts, maintenance processes, and other aspects that are related to your database server environment.

Changes to functionality that was introduced before Informix® 15.0.0 can also affect your plans.

## Checking and configuring available space

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

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

When you migrate to Version 15.0.0, you need the following space for building **sysmaster**, **sysutils**, and **sysadmin** databases:

- 21892 KB of logical-log spaces (or 10946 pages) for 2 K page platforms
- 26468 KB of logical-log spaces (or 6617 pages) for 4 K page platforms

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

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

If your migration fails because there is insufficient space in the partition header page, you must unload your data before you try to migrate again. Then you must manually load the data into the new version.

The root chunk should contain at least ten percent free space when converting to the new version of the server.

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

Add any additional free space to the system prior to the migration. If the dbspaces are nearly full, add space before you start the migration procedure. When you start the new version of Informix® on the same root dbspace of the earlier database server, Informix® automatically converts the **sysmaster** database and then each database individually.

For a successful conversion of each database, ensure that 2000 KB of free space per database is available in each dbspace where a database resides.

To ensure enough free space is available:

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

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

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

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

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

You can run SQL statements to determine the free space that each dbspace requires and the free space available. These statements return the free-space calculation in page-size units. The **free\_space\_req** column value is the free-space requirement, and the **free\_space\_avail** column value is the free space available.

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

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

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

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



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

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

## Configuring for recovery of restore point data in case an upgrade fails

By default, the `CONVERSION_GUARD` configuration parameter is enabled and a temporary directory is specified in the `RESTORE_POINT_DIR` configuration parameter. These configuration parameters specify information that can use if an upgrade fails. You can change the default values of these configuration parameters before beginning an upgrade.

### Before you begin

**Prerequisites:** The directory specified in the `RESTORE_POINT_DIR` configuration parameter must be empty before the upgrade begins, but not when recovering from a failed update.



### Important:

After a failed upgrade, do not empty the `RESTORE_POINT_DIR` directory before you attempt to run the `onrestorept` utility. The server must be offline after a failed upgrade.

### About this task

You can change the value of the `CONVERSION_GUARD` configuration parameter or the directory for restore point files before beginning an upgrade. The default value for the `CONVERSION_GUARD` configuration parameter in the `ONCONFIG` file is `(2)`, and the default directory where the server will store the restore point data is `$INFORMIXDIR/tmp`. You must change this information before beginning an upgrade. You cannot change it during an upgrade.

To change information:

1. If necessary for your environment, change the value of the `CONVERSION_GUARD` configuration parameter.

When the `CONVERSION_GUARD` configuration parameter is set to `2` (the default value), the server will continue the upgrade even if an error related to capturing restore point data occurs, for example, because the server has insufficient space to store the restore point data.

However, if the `CONVERSION_GUARD` configuration parameter is set to `2` and the upgrade to the new version of the server fails, you can use the **onrestorept** utility to restore your data.

However, if you set the `CONVERSION_GUARD` configuration parameter to `2`, conversion guard operations fail (for example, because the server has insufficient space to store restore point data), and the upgrade to the new version fails, you cannot use the **onrestorept** utility to restore your data.

2. In the `RESTORE_POINT_DIR` configuration parameter, specify the complete path name for a directory that will store restore point files.

The server will store restore point files in a subdirectory of the specified directory, with the server number as the subdirectory name.

## What to do next

If the `CONVERSION_GUARD` configuration parameter is set to `1` and an upgrade fails, you can run the **onrestorept** utility to restore the instance back to its original state just before the start of the upgrade.

If the `CONVERSION_GUARD` configuration parameter is set to `1` and conversion guard operations fail (for example, because the server has insufficient space to store restore point data), the upgrade to the new version will also fail.

If any restore point files from a previous upgrade exist, you must remove them before you begin an upgrade.

Even if you enable the `CONVERSION_GUARD` configuration parameter, you should still make level 0 backup of your files in case you need to revert after a successful upgrade or in case a catastrophic error occurs and you cannot revert.

## Saving copies of the current configuration files

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

Although you can use an old `ONCONFIG` configuration file with Informix® Version 15.0.0, you should use the new Version 15.0.0 `ONCONFIG` file, or at least examine the file for new parameters.

Configuration files that you might have are listed in [Table 2: Configuration files to save from the source database server on page 12](#).

**Table 2. Configuration files to save from the source database server**

UNIX™ or Linux™	Windows™
<code>\$INFORMIXDIR/etc/\$ONCONFIG</code>	<code>%INFORMIXDIR%\etc\%ONCONFIG%</code>
<code>\$INFORMIXDIR/etc/onconfig.std</code>	<code>%INFORMIXDIR%\etc\onconfig.std</code>
<code>\$INFORMIXDIR/etc/oncfg*</code>	<code>%INFORMIXDIR%\etc\oncfg*</code>
<code>\$INFORMIXDIR/etc/sm_versions</code>	<code>%INFORMIXDIR%\etc\sm_versions</code>
<code>\$INFORMIXDIR/aaodir/adtcfg</code>	<code>%INFORMIXDIR%\aaodir\adtcfg.*</code>
<code>\$INFORMIXDIR/dbssodir/adtmasks</code>	<code>%INFORMIXDIR%\dbssodir\adtmasks.*</code>

**Table 2. Configuration files to save from the source database server (continued)**

UNIX™ or Linux™	Windows™
\$INFORMIXDIR/etc/sqlhosts or \$INFORMIXSQLHOSTS	%INFORMIXDIR%\etc\sqlhosts or \$INFORMIXSQLHOSTS
\$INFORMIXDIR/etc/tctermcap	
\$INFORMIXDIR/etc/termcap	

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

**UNIX™ or :**

**\$INFORMIXDIR/etc/ixbar.srvnum**

**Windows™:**

**%INFORMIXDIR%\etc\ixbar.srvnum**

The HCL Informix® Primary Storage Manager does not use the `sm_versions` file. If you plan to use the , you do not need the `sm_versions` file. However, if you use the IBM Spectrum Protect or a third-party storage manager, you do need the `sm_versions` file.

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

## Preparing 12.10 BSON columns with DATE fields for upgrade

Before you upgrade from Informix® 12.10 you must unload binary JSON (BSON) columns with DATE fields into JSON format so that you can load them into Informix® 15.0.0

### About this task

Perform the following steps on the 12.10 server.

1. Create an external table with a similar name as the original table and with JSON (instead of BSON) format for the date.

For example, assume that the original table named **datetab** has a BSON column named **i** that has DATE fields in it. Use the following statement to create an empty, external table named **ext\_datetab** that has a JSON column with DATE fields. The DATAFILES clause specifies the location and name of the delimited data file, which in this example is `disk:/tmp/dat.unl`.

```
create external table ext_datetab (j int, i json) using
  (datafiles ("disk:/tmp/dat.unl"),
   format "delimited");
```

2. Unload the data from the original table into the external table.

For example:

```
insert into ext_datetab select j, i::json from datetab;
```

## What to do next

Complete other pre-migration steps. After you upgrade to the new server, you must load the JSON columns with DATE fields from the external table into a new table in BSON format.

## Resolve outstanding in-place alter operations

Resolving outstanding in-place alter operations is a requirement before *converting* to Version 15.0.0. If it becomes necessary to *revert* to a previous version you must also resolve any new outstanding in-place alter operations first.

An in-place alter operation is *outstanding* when data pages still exist with the prior definition, a state that can be detected with the oncheck -pT command. An in-place alter operation is *new* if the ALTER statement is executed in the higher database version.

If the reversion process detects new outstanding in-place alter operations, reversion fails and the message log will contain a list of all tables whose in-place alter operations must be resolved before the reversion will succeed.

If you are reverting from version 12.10.xC4 or later, you can remove in-place alter operations by running the admin( ) or task( ) SQL administration command with the table update\_ipa or fragment update\_ipa argument. You can include the parallel option to run the operation in parallel. For example, the following statement removes in-place alter operations in parallel from a table that is named **auto**:

```
EXECUTE FUNCTION task('table update_ipa parallel','auto');
```

If you are reverting from an earlier version of 12.10, you can resolve outstanding in-place alter operations by running sample UPDATE statements. Sample UPDATE statements force any outstanding in-place alter operations to complete by updating the rows in the affected tables. To generate a sample UPDATE statement, create an UPDATE statement in which a column in the table is set to its own value. This forces the row to be updated to the latest schema without changing column values. Because the database server always alters rows to the latest schema, a single pass through the table that updates all rows completes all outstanding in-place alter operations.

The sample UPDATE statement differs from a standard UPDATE statement because it does not change the data. A standard UPDATE statement usually changes the value of the affected row.

For example, to create a sample update, specify:

```
UPDATE tab1 SET col1=col1 WHERE 1=1 ;
```

You must ensure that the column selected is a numeric data type (for example, INTEGER or SMALLINT) and not a character data type.

If a table is large, a single update of the whole table can cause a long transaction. To avoid a long transaction, update the table in pieces, by ranges of some column, with this statement:

```
... WHERE {id_column} BETWEEN {low_value} AND {step_value}
```

For example, specify:

```
UPDATE tab1 SET col1=col1 WHERE col1 BETWEEN 1 AND 100;  
UPDATE tab1 SET col1=col1 WHERE col1 BETWEEN 101 AND 200;
```

Ensure that the UPDATE statements include the entire data set.

If the table is replicated with Enterprise Replication, the database server replicates all updated rows unnecessarily. To avoid replication, update the table as follows:

```
BEGIN WORK WITHOUT REPLICATION;
...
COMMIT WORK;
```

When all the pending in-place alter operations are resolved, run the oncheck -pT command again for each table. In the output of the command, check information in the `Version` section. The number of data pages should match with current version. Also, all other table versions should have `count=0` for the number of data pages that the version is accessing.

For example, if you run the oncheck -pT testdb:tab1 command after outstanding in-place alter operations are resolved, you might see information similar to the information in this segment of sample output:

TBLspace Report for testdb:root.tab1

```
Physical Address      1:860
Creation date         06/23/2011 14:23:08
TBLspace Flags        800801      Page Locking
                        TBLspace use 4 bit bit-maps

Maximum row size      29
Number of special columns 0
Number of keys        0
Number of extents     1
Current serial value   1
Current SERIAL8 value  1
Current BIGSERIAL value 1
Current REFID value    1
Pagesize (k)          2
First extent size     8
Next extent size      8
Number of pages allocated 8
Number of pages used   4
Number of data pages   3
  << Number of data pages used is 3 >>
Number of rows        6
Partition partnum     1048981
Partition lockid      1048981
```

Extents

Logical Page	Physical Page	Size	Physical Pages
0	1:1895	8	8

TBLspace Usage Report for testdb:root.tab1

Type	Pages	Empty	Semi-Full	Full	Very-Full
Free	4				
Bit-Map	1				
Index	0				
Data (Home)	3				
-----					
Total Pages	8				

Unused Space Summary

```

Unused data slots                                177

Home Data Page Version Summary

Version                                Count

3 (oldest)                                0
  << Other version should show data page count=0>>
4                                          0
  << Other version should show data page count=0>>
5 (current)                               3
  << Current should always match the number of data pages>>

```

## Closing all transactions and shutting down the source database server

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

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

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

To let users exit and shut down the database server gracefully

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

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

```

onmode -l
onmode -c
onmode -ky

```

## Initiating fast recovery to verify that no open transactions exist

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

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

```

oninit -s

```



### UNIX/Linux Only

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

You must run the **oninit -s** command before you initialize the new version of Informix®. If any transactions remain when you try to initialize the new database server, the following error message appears when you try to initialize the new database server, and the server goes offline:

```
An open transaction was detected when the database server changed log versions.
Start the previous version of the database server in quiescent mode and then shut
down the server gracefully, before migrating to this version of the server.
```

For more information about fast recovery, see your *HCL® Informix® Administrator's Guide*.

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

Only after proper shutdown can you bring the new database server (Informix® Version 15.0.0) through the migration path. Any transaction that is open during the migration causes an execution failure in fast recovery.

## Verifying the integrity of the data

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

To obtain the database names, use the following statements with :

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

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

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

[Table 3: Commands for verifying the data integrity on page 17](#) lists the **oncheck** commands that verify the data integrity.

**Table 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 <i>database_name</i></b>
Check data	<b>oncheck -cD <i>database_name</i></b>
Check indexes	<b>oncheck -cI <i>database_name</i></b>

## Verifying that the database server is in quiescent mode

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

Run the `onstat -` command to verify that the database server is in quiescent mode.

The first line of the `onstat` output shows the status of your source database server. If the server is in quiescent mode, the status line includes this information:

```
Quiescent -- Up
```

## Making a final backup of the source database server

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

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



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

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

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

## Verifying that the source database server is offline

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

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

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

## Pre-migration checklist of diagnostic information

Before you migrate to a newer version of Informix®, gather diagnostic information, especially if you have large, complex applications. This information will be useful to verify database server behavior after migration. This information will also be useful if you need help from HCL Software Support.

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

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

**Table 4. Checklist of information to get before migrating**

Information to Get Before Migrating	Done
Get the SQL query plans for all regularly used queries, especially complex queries, by using SET EXPLAIN ON.	
Run the <b>dbschema -d -hd</b> command for all critical tables.  The output contains distribution information.	
Get <b>oncheck -pr</b> output that dumps all of the root reserved pages.	
Make a copy of the ONCONFIG configuration file.  A copy of the ONCONFIG file is essential if you need to revert to an earlier version of the database server. In addition, a copy of this file is useful because <b>oncheck -pr</b> does not dump all of the configuration parameters.	
Prepare a list of all the environment variables that are set using the <b>env</b> command.	
During times of peak usage: <ul style="list-style-type: none"> <li>• Obtain an <b>online.log</b> snippet, with some checkpoint durations in it</li> <li>• Run <b>onstat -aF, -g all</b>, and <b>-g stk all</b>.</li> </ul>	
During times of peak usage, run the following <b>onstat</b> commands repeatedly with the <b>-r repeat</b> option for a period of about three to five minutes: <ul style="list-style-type: none"> <li>• <b>onstat -u</b>, to see the total number of <b>sqlxecs</b> used</li> <li>• <b>onstat -p</b>, for read and write cache rates, to detect deadlocks and the number of sequential scans</li> <li>• <b>onstat -g nta</b>, a consolidated output of <b>-g ntui, ntt, ntm</b> and <b>ntd</b></li> <li>• <b>onstat -g nsc, -g nsd</b>, and <b>-g nss</b> for the status of shared memory connections</li> <li>• <b>onstat -P, -g tpf</b>, and <b>-g ppf</b></li> <li>• <b>vmstat, iostat</b> and <b>sar</b>, for cpu utilization</li> <li>• <b>timex</b> of all queries that you regularly run</li> </ul>	

## Migrating from 32-bit to 64-bit database servers

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

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

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

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

## Migrating from 32-bit to 64-bit with collection types that use the SMALLINT data type

If you are moving your database from a 32-bit computer to a 64-bit computer and your database contains collection types that use the SMALLINT data type, you must take extra steps to prevent memory corruption. Collection types are the ROW, LIST, SET, and MULTiset data types. This restriction applies if you are upgrading from an older version of Informix® on 32-bit to the current version of Informix® on 64-bit, or if you are moving from 32-bit to 64-bit on the current version of Informix®.

To migrate a database with SMALLINT collection types from 32-bit to 64-bit, use one of the following methods:

- Export and import the data.
  1. Export the data from the 32-bit computer.
  2. Import the data onto the 64-bit computer.
- Drop and recreate specific collection types and database objects.
  1. Drop the collection types that use the SMALLINT data type and all other database objects that reference them (such as tables, columns, SPL routines, triggers, indexes, and so on).
  2. Recreate the collections types and all the other necessary database objects.

## CSDK Library Updates for 15.0.0

In version 15.0.0, some of the CSDK libraries have undergone name changes.

The pattern of changes mainly involves the suffixes being updated from "s09" to "s15" and the addition of "15a" in some file names.

If you are referencing these libraries in your build scripts, ensure that you use the new names. The following table shows the new names:

Library	Old Filename (4.50.x)	New Filename (15.0.x)
lib/c++	ic++s02a.so	ic++s15a.so
	ic++t02a.so	ic++t15a.so
	libifc++.so	libifc++15a.so
	libthc++.so	libthc++15a.so
	iclis09b.so	iclis15a.so
	iclit09b.so	iclit15a.so
	idmrs09a.so	idmrs15a.so
	libifcli.so	libifcli15a.so
	libifdmr.so	libifdmr15a.so
	libixcli.so	libixcli15a.so

	libtcli.so	libtcli15a.so	
	libthcli.so	libthcli15a.so	
lib/client/csm	igsss11a.so	igsss15a.so	
	libixgss.so	libixgss15a.so	
lib/dmi	idmis09a.so	idmis15a.so	
	idmit09a.so	idmit15a.so	
	libifdmi.so	libifdmi15a.so	
	libthdmi.so	libthdmi15a.so	
lib/esql	icobsqsls09b.so	icobsqsls15a.so	
	ifgisqsls09b.so	ifgisqsls15a.so	
	igens09a.so	igens15a.so	
	igent09a.so	igent15a.so	
	iosls09a.so	iosls15a.so	
	ioslt09a.so	ioslt15a.so	
	isqsls09b.so	isqsls15a.so	
	isqlt09b.so	isqlt15a.so	
	isxas09a.so	isxas15a.so	
	isxat09a.so	isxat15a.so	
	libifcobsqsl.so	libifcobsqsl15a.so	
	libiffgisqsl.so	libiffgisqsl15a.so	
	libifgen.so	libifgen15a.so	
	libifos.so	libifos15a.so	
	libifsql.so	libifsql15a.so	
	libifxa.so	libifxa15a.so	
	libixcobsqsl.so	libixcobsqsl15a.so	
	libixfgisqsl.so	libixfgisqsl15a.so	
	libixgen.so	libixgen15a.so	
	libixos.so	libixos15a.so	
	libixsql.so	libixsql15a.so	
	libixxa.so	libixxa15a.so	
	libtgen.so	libtgen15a.so	
	libthgen.so	libthgen15a.so	
	libthos.so	libthos15a.so	
	libthsql.so	libthsql15a.so	
	libthxa.so	libthxa15a.so	
	libtos.so	libtos15a.so	
	libtsql.so	libtsql15a.so	
	libtxa.so	libtxa15a.so	
lib	iasfs09b.so	iasfs15a.so	
	iasft09b.so	iasft15a.so	
	icsss09a.so	icsss15a.so	
	libifasf.so	libifasf15a.so	
	libifcss.so	libifcss15a.so	
	libixasf.so	libixasf15a.so	
	libtasf.so	libtasf15a.so	
	libthasf.so	libthasf15a.so	
-----			



**Important:** The libraries (.NetFramework/.NetCore) included with CSDK 15.0.x will only be compatible with CSDK 15.0.0 or later versions. There will be no backward compatibility.

## and Upgrading

The server nodes in your Enterprise Replication (ER) domain can run different versions and fixpack levels. On the other hand, new features might only be available between servers running the new version, and administrative 'cdr' commands should be run from the newest version environment.

If the server being upgraded is running Enterprise Replication in combination with High-availability cluster replication (HDR, RSS, SDS), the cluster acts as one ER server and ER will be active only on the respective cluster primary. Should this primary role change during the cluster upgrade, so would the active ER server. This should be considered particularly [when planning an online cluster upgrade on page 28](#) to minimize downtime, likely with the intention to also have ER replicating throughout the process.

There is no immediate requirement to upgrade the various servers in your Enterprise Replication domain in any coordinated manner. Should the newer version provide any new or improved functionality, this might be available only on or between servers running this version; as far as the novelty is restricted to the "cdr" command line utility, it might be sufficient to use the new version's utility in order to enable the functionality across more servers.



**Note:** When upgrading only a subset of servers to version 15.x, especially if LARGE tables (see also [SMALL and LARGE tables](#)) are meant to participate in replication, only the version 15.x "cdr" utility will be able to work with these tables, for instance when defining, checking or syncing replicates.

When upgrading a server running ER, ER must explicitly be stopped before starting the upgrade process, and restarted manually only once all conversion steps are completed successfully.

These topics describe the additional tasks that you must perform when upgrading to and reverting from Version 15.0.0 if you are running .

## Preparing to upgrade with Enterprise Replication

If you use Enterprise Replication, you must do replication-related tasks to prepare for migration.

### Before you begin

You must do all migration operations as user **informix**, unless otherwise noted.

Only a DBSA can run the cdr check queue command. With a non-root installation, the user who installs the server is the equivalent of the DBSA, unless the user delegates DBSA privileges to a different user.

To prepare for migration with Enterprise Replication:

1. Stop applications that are performing replicable transactions.
2. Make sure that the replication queues are empty.

If you are migrating from Informix® 12.10.xC1 or later releases, run the following commands to check for queued messages and transactions:

- a. `cdr check queue -q cntrlq targetserver`
- b. `cdr check queue -q sendq targetserver`
- c. `cdr check queue -q recvq targetserver`

If you are migrating from an earlier version of Informix®, run the following commands:

- a. Run `onstat -g grp` to ensure that the Enterprise Replication grouper does not have any pending transactions. The grouper evaluates the log records, rebuilds the individual log records in to the original transaction, packages the transaction, and queues the transaction for transmission.
  - b. Run `onstat -g rqm` to check for queued messages.
3. Shut down by running the `cdr stop` command. ER will now remain in this stopped state until explicitly restarted once the upgrade is completed.

## Results

Now you can complete the steps in [Preparing for migration on page 8](#) and, if necessary, in [Migrating from 32-bit to 64-bit database servers on page 19](#).

## Preparing Replication of user-defined data types

User-defined data types can also be replicated using the required support libraries.

If you are replicating your own user-defined data types, make sure the new version environment has the required support function libraries installed. For more information, see [Replication of opaque user-defined data types](#).

## Migrating with Enterprise Replication

If you use Enterprise Replication, you must complete replication-related tasks when you migrate to a new version of .

### About this task

Prerequisites:

- Complete the steps in [Preparing for migration on page 8](#).
- Complete the steps in [Preparing to upgrade with Enterprise Replication on page 22](#)
- Perform all migration operations as user **informix**.
- All servers in the Enterprise Replication domain must be available, in order for the replication queues to be able to drain.

To migrate with Enterprise Replication:



**Note:**



If you are migrating server from a version lower than 14.10.xC6 to 14.10.xC6 or a later version, after upgrading server to 14.10FC6 or a later version, for the first time when you start ER, make sure to start ER using `cdr cleanstart` command to force required schema changes to syscdr database.

Starting with 14.10.xC6, send queue spooling of smaller transactions (< ~30kB) will no longer occur to sblobs within sbspaces, but instead go to a new LVARCHAR column within the spool table. As a consequence of this, when spooling occurs, the space consumption might no longer occur mainly in the spool sbspaces (CDR\_QDATA\_SBSPACE), but rather in the spool table's dbspace (CDR\_QHDR\_DBSPACE, defaulting to rootdbs). Having a dedicated, possibly extendable dbspace configured for this purpose is therefor more important now and will take effect with `cdr cleanstart`.

1. Perform the tasks that are described in [Migrating to the new version of Informix on page 45](#), including starting the new version of the server.
2. Make sure that no replicable transactions occur before starts.

Especially with the upgrade to version 15.x.x.x from a previous version, the log snooper cannot start snooping at a log position from the previous version's life, which is one reason why replication should be started for the first time after conversion using `cdr cleanstart`, directing replication to start snooping at the current log position and meaning that any new transactions prior to this position would not be replicated.

The database server will also demand using `cdr cleanstart` rather than simple `cdr start` for the first startup of replication when upgrading from a version prior to 14.10.xC6.

3. After successful conversion, if you are upgrading server to 14.10xC6 or later version, start by running the `cdr cleanstart` command, otherwise run `cdr start` command. All required internal ER conversion now will be performed automatically by the database server.
4. If you upgraded the servers in a grid from version 11.70 to version 12.10 and you want to copy external files to the servers in the grid, you must enable the ability to copy external files. To enable the copying of external files, run this command:

```
cdr modify grid grid_name --enablegridcopy server_groupname
```

## Results



**Important:** After you convert to the new version of with , do not drop the **syscdr** database. If you intend to upgrade without replication and re-define it after the upgrade, do the required undefining, including deletion of the replication server definition, prior to the upgrade still in the old version.

## Reverting with Enterprise Replication

If you use Enterprise Replication, you must complete replication-related tasks when you revert from the new version of .

### Before you begin



Prerequisites:

- Perform all reversion operations as user **informix**.
- Enterprise Replication must be running, or you must delete the replication server before you revert.
- All servers in the Enterprise Replication domain must be available.
- Replication queues must be nearly empty.

### About this task

If you want to revert to a version earlier than when Enterprise Replication was defined on this server, you must remove Enterprise Replication from this server before reverting. After you remove Enterprise Replication, you can revert your server using the instructions at [Reverting from Version 15.0.0 on page 59](#).

To revert from Version 15.0.0 with Enterprise Replication:

1. Release or remove any grid statements that are deferred from propagation by running the `ifx_grid_release()` or `ifx_grid_remove()` function.
2. Stop applications that are doing replicable transactions.
3. Remove Enterprise Replication features from the current release that cannot be reverted. For more information, see [#unique\\_39](#).
4. Run the `onstat -g cat repls` command to check whether if Enterprise Replication is in alter mode. If so, run the `cdr alter -off` command.
5. Delete shadow replicates.
6. Run `cdr stop` to shut down Enterprise Replication.
7. Perform database server reversion tasks, as described in [Reverting from Version 15.0.0 on page 65](#).
8. Run `onmode -l` and `onmode -c` to prevent the database server from failing when you start Enterprise Replication.
9. After reversion completed and old version started successfully, start by running the `cdr start` command. When reverting to a version prior to version 15.x.x.x, use `cdr cleanstart` instead to avoid the log snooper reading the newer version's log files.

## High-availability cluster migration

You must coordinate the migration and reversion of all servers that are involved in high-availability clusters. You can use a rolling upgrade in some cases to update servers while the cluster is online, with minimal interruption to client applications. Otherwise, you must schedule cluster downtime to upgrade the servers. A cluster must be offline to reverse an upgrade or revert to an earlier release.

A rolling upgrade can minimize downtime, but the tradeoff is that you must spend time to prepare for a rolling upgrade. The effort to prepare your servers for a rolling upgrade depends on how closely your current configuration matches the procedure requirements. In some cases, you might find it easier to plan for downtime during a period of low activity instead of setting up your environment for a rolling upgrade. The downtime that is required depends on your system, but smaller clusters usually require less downtime.

Use the following table to help you determine which upgrade procedure to use. Also, review the requirements and limitations that are documented in each procedure.

**Table 5. Comparison of upgrade procedures for high-availability clusters**

Upgrade to	Procedure	Use	Procedure overview
Next consecutive fix pack or PID of the same major version	<a href="#">Rolling upgrade of an online cluster to the next fix pack or PID (UNIX, Linux) on page 28</a>	As of Informix® 12.10.xC5, you can apply the next consecutive fix pack or interim fix (PID) of the same major version to all the servers while the cluster remains online.	You must stop and restart each server in the cluster, including the primary. Specifically, you upgrade each secondary server, and then you stop the primary server and promote an upgraded secondary server to primary. Next, you upgrade the original primary server, bring it online as a secondary server, and promote it back to primary, if necessary. Interruption to client applications is minimized because transactions are redirected to active servers by Connection Manager or through the client applications. The new database server capabilities can be used in the cluster after all servers are upgraded.
Any later fix pack or PID of the same major version	<a href="#">Upgrading an offline cluster to a fix pack or PID on page 32</a>	<p>You can use this procedure within a major version to apply a fix pack or PID in offline clusters in the following situations:</p> <ul style="list-style-type: none"> <li>• Your cluster does not meet the requirements for a rolling upgrade.</li> <li>• The limitations that are imposed by a rolling upgrade are impractical for your environment.</li> <li>• You prefer or need to bring the servers in the cluster offline before you upgrade them.</li> </ul>	You stop all servers in the cluster during a period when downtime is acceptable. While the cluster is offline, you upgrade all the servers. You then start the primary followed by the secondary servers. You don't have to rebuild the cluster or clone the servers.
New major version	<a href="#">Migrating an</a>	You can use this procedure to migrate the servers in your offline cluster to a new major	You stop the servers in a specific order during a period when planned downtime is acceptable. Then, if you are migrating

Table 5. Comparison of upgrade procedures for high-availability clusters

(continued)

Upgr ade to	Proce dure	Use	Procedure overview
jor vers ion  A fix pack or PID that requi res conv ers ion	<a href="#">offline clus ter to a new major vers ion on p age 33</a>	version. For example, you can migrate from Informix® 12.10 to Informix® 15.0.0.  Also, you can use this procedure to apply a fix pack or PID that requires standard conversion procedures. Usually conversion is not required for a fix pack or PID. The machine notes indicate whether conversion is required.	from Informix 12.10 or later version, You stop all servers in the cluster during a period when downtime is acceptable. While the cluster if offline, you upgrade all the servers. You then start the primary followed by the secondary servers. You don't have to rebuild the cluster or clone the servers.
Any supp or ted rele ase	<a href="#">Roll ing upgr ade of an online clus ter with Enterp rise Replic ation on p age 36</a>	This alternative rolling upgrade procedure requires a working knowledge of Enterprise Replication. You can use this procedure to upgrade online clusters to any supported product release. You can use this procedure to upgrade to a new major version or to apply any fix pack or PID (not just the next consecutive one) within a release. This procedure is more flexible than the other rolling upgrade procedure; however, it is more complex to set up.	You temporarily convert the primary and secondary servers to stand-alone Enterprise Replication servers. The upgrade occurs without incurring any downtime because Enterprise Replication supports replication between different versions of the server software.

## Preparing to migrate, upgrade, or revert clusters

If you use high-availability clusters, you must coordinate the migration of all of the servers that are involved in a high-availability cluster, and you must perform additional steps when preparing to migrate.

## About this task

Prerequisites:

- Perform all migration operations as user **informix**.
- Download the product package from [at](#) , or download [recommended fixes for server products](#).
- If you are migrating to a new version of the server, complete all steps in [Preparing for migration on page 8](#).

To prepare for migration or reversion of a cluster:

1. Prepare each server in the cluster:

- a. Install the target server software in a different location from where the source database server software is installed. Do not install the target server software over the source server.



**Tip:** Install the new version of HCL® Informix® Client Software Development Kit (Client SDK), which contains the Connection Manager that you must use with the new version of the server. If you want to use only the Connection Manager from the new version of , install in a new location and use the Connection Manager from that location.

- b. Copy the `onconfig` and `sqlhosts` configuration files to the target installation directory (for example, **\$INFORMIXDIR/etc**).
- c. Copy any keystore (.p12), stash (.sth), and trusted host files to the new **\$INFORMIXDIR/etc** directory.
- d. Install on the target server any user-defined objects or modules that are used on the source server.

2. Back up your primary server. You can perform this step in the following ways:

**Choose from:**

- Use ON-Bar or ontape to make a level-0 backup on the primary source server.
- If you have a high-availability cluster with a High-availability Data Replication (HDR) secondary server, you can use the HDR secondary server as a standby server for any contingencies that occur while you upgrade the primary server. However, if it is necessary to use an HDR secondary server as a standby server for contingencies, do not perform updates on the standby server while migration or reversion is in progress, because the updates cannot be replicated and will be lost. Additionally, nonlogged objects on the primary server will not exist on the secondary server.



**Attention:** Do not use RS secondary servers as backup servers, because transactions could be lost.

## Rolling upgrade of an online cluster to the next fix pack or PID (UNIX, Linux)

Note that it is not possible to perform a rolling upgrade to version 15.0.0. Consult the following information only if you are upgrading an Informix server to version 14.10 or earlier.

## About this procedure

As of Informix® 12.10.xC5, you can apply the next consecutive fix pack or interim fix (PID) to your current version of the database server. During the rolling upgrade, the cluster remains online even though the servers in the cluster run different levels of the software. The new capabilities can be used in the cluster after all the servers in the cluster are upgraded.

This rolling upgrade procedure works only on UNIX and Linux platforms. Use this procedure to apply the next *consecutive* fix pack or PID in the current major version. For example, you can do a rolling upgrade from 12.10.xC4 to 12.10.xC5. Otherwise, you must use a different procedure that is documented at [High-availability cluster migration on page 25](#).

Do *not* use this procedure in the following situations:

- To apply a fix pack or PID that requires conversion.
- To upgrade to 12.10.xC5 from 12.10.xC3, 12.10.xC2, or 12.10.xC1.
- To upgrade an earlier major version of the server to a later major version of the database server, for example, to upgrade from version 11.70 to 12.10.
- To upgrade from a patch release or special build, unless advised to do so by IBM® Software Support.

You must stop and restart each server in the cluster, including the primary, as part of this procedure. Specifically, you upgrade each secondary server, and then you stop the primary server and promote an upgraded secondary server to primary. Next, you upgrade the original primary server, bring it online as a secondary server, and promote it back to primary, if necessary.

You must have the following permission on all the servers in the cluster: user **root** or user **informix**.

This procedure consists of these steps:

- [Prepare for a rolling upgrade on page 29](#)
- [Upgrade the servers on page 30](#)
- [Return the cluster to its original configuration on page 31](#)

## Prepare for a rolling upgrade


Use the following steps to plan for and prepare the servers for rolling upgrade:


1. Complete the steps in [Preparing to migrate, upgrade, or revert clusters on page 27](#), which include installing the new software on all the servers in the cluster, copying the appropriate configuration files, and backing up the primary server.
2. If you are migrating from a server version prior to 14.10.xC8, or from a Client SDK version prior to 4.50.xC8, you must upgrade all Connection Managers before upgrading any server in your cluster. Upgrade them one at a time to avoid a service interruption that would affect your applications. All Connection Managers should be running with version 4.50.xC8 or later before continuing to the next step.
3. Configure client redirection to minimize interruption of service.

Set up redirection and connectivity for clients by using the method that works best for your environment.

If Connection Manager controls the connection redirection in the cluster: Ensure that every service level agreement (SLA) definition in the Connection Manager configuration file can redirect to at least one server other than the one you are about to update. For example, assume that you have an SLA with only one secondary. Before you upgrade the secondary server in that SLA, update the SLA to include the cluster primary (PRI).

4. Ensure that the primary server has an appropriate amount of disk space for the logical log records that are created during the entire upgrade process. The space that is required depends on your environment.

 **Attention:** If a log wrap occurs during the rolling upgrade procedure, you must apply the fix pack or PID while the cluster is offline.


 **Tip:** Examine the online log to get an estimate of your data activity during normal operations. You might want to ensure that you have enough space for data activity for a day. Also, you might find it convenient to plan the rolling upgrade for a period of low traffic.

5. Prepare the secondary server that will become the primary when you upgrade the original primary server. You must use an SD secondary or a fully synchronous HDR secondary server that has transactional consistency with the original primary server.
  - a. If the cluster contains an SD secondary server, you don't need to do any additional preparation to that server.
  - b. If the cluster contains an HDR secondary server, make sure that it runs in fully synchronous (SYNC) mode.
  - c. If the cluster contains only RS secondary servers in addition to the primary server, you must change one of the RS secondary servers to an HDR secondary server in SYNC mode.

## Upgrade the servers

Follow the steps in this section to upgrade the cluster servers in the following order:

1. Remote standalone (RS) secondary server
2. HDR secondary server
3. Shared disk (SD) secondary server
4. Primary server

 **Important:** Upgrade the primary server only after *all* the secondary servers are upgraded and tested. After you upgrade the primary server, if you want to revert to your original environment you must take the cluster offline.

1. Run the `onmode -c` command to force a checkpoint for each server.
2. If a wire listener is running on the server that you want to upgrade, stop that wire listener.
3. Stop the server that you want to upgrade. If you can wait for all connections to exit gracefully before stopping the server, use the `onmode -kuy` command. Otherwise, use the `onmode -ky` command to stop the server.

- **When you stop a secondary server:** If [redirection is configured on page 29](#) for the cluster, the client application automatically connects to another active server in the cluster.
- **When you stop the primary server:** If failover is configured for the cluster, a secondary server is promoted automatically to primary. Otherwise, you can run the `onmode -d make primary` command to promote a prepared secondary server to primary.



**Tip:** If the primary is offline before the failover, you must use the `onmode -d make primary force` command.

4. Set your environment to use the fix pack or PID that you installed on the server.
  - a. Set the **INFORMIXDIR** environment variable to the full path name for the target installation.
  - b. Update all environment variables that depend on the **INFORMIXDIR** environment variable. At a minimum, update these environment variables: **PATH**, **DBLANG**, **INFORMIXSQLHOSTS**, and any platform-specific library path environment variables, such as **LD\_LIBRARY\_PATH**.

5. Start the upgraded server.

**To start an upgraded secondary server:** Use the `oninit` command.

**To start an upgraded original primary server:** Start the original primary server as a secondary server. For convenience, start it as the server type that was promoted to primary during the rolling upgrade. For example, if you promoted an HDR server to primary for the rolling upgrade, start the original primary as an HDR secondary server.

- To start the upgraded server as an SD secondary server, run the `oninit -SDS` command.
- To start the upgraded server as an HDR secondary server, run the `oninit -PHY` command, and then run the following command: `onmode -d secondary primary_server secondary_server`

After the server starts, it runs the new version of the software and automatically reconnects to the cluster.

6. To verify that the upgraded secondary server is active in the cluster, run the `onstat -g cluster` command on both the server you upgraded and on the primary server.
7. If you stopped the wire listener to upgrade this server, restart the wire listener.

After you upgrade all the servers and restart them, the original primary server is running as a secondary server.

## Return the cluster to its original configuration

Use the following steps if you want the cluster to operate as it did before you prepared it for a rolling upgrade.

1. Manually promote the secondary server that was the original primary to be the primary server again.
  - a. Run the `onmode -c` command to force a checkpoint.
  - b. Run the `onmode -d make primary` command to promote the secondary server to primary.
2. Undo any changes that you made when you prepared the servers for a rolling upgrade. Some of these optional steps might not apply to you.

- Adjust the amount of disk space that is allocated for logical log records.
- Convert the HDR secondary server back to an RS secondary server.
- Change the HDR secondary server back to ASYNC mode from SYNC mode.
- Change the Connection Manager SLA definitions.

## Upgrading an offline cluster to a fix pack or PID

You can bring a high-availability cluster offline to apply any PID or fix pack of the same version of the database server. For example, you must bring the clusters offline to upgrade the servers to the current fix pack if earlier fix packs were not applied to the server. This procedure is an alternative to the rolling upgrade procedure.

### Before you begin

Prerequisites:

- Verify that you are upgrading to a PID or fix pack in which standard conversion procedures are not necessary. If the PID or fix pack requires you to complete standard conversion procedures, or if you want to upgrade to a new major release such as version 15.0.0, go to [Migrating an offline cluster to a new major version on page 33](#) instead of following the procedures in this topic.
- Complete the steps in [Preparing to migrate, upgrade, or revert clusters on page 27](#).
- Perform all migration operations as user **informix**.

To upgrade offline clusters to a new PID or fix pack:

1. Stop the Connection Manager by issuing the `oncmsm -k connection_manager_name` command.
2. If you are using a High-availability Data Replication (HDR) secondary server as a backup server in case of contingencies:
  - a. Quiesce the primary server by issuing an `onmode -sy` command to prevent user connections to the server.
  - b. Force a checkpoint by issuing an `onmode -c` command on the primary server.
3. Stop secondary servers in the cluster in the following order:
  - a. If you have remote standalone (RS) servers, stop them by issuing the `onmode -ky` command.
  - b. If you have shared disk (SD) servers, stop them by issuing the `onmode -ky` command.
  - c. If you have an HDR secondary server, stop it issuing the `onmode -ky` command.
4. Stop the primary server by issuing the `onmode -ky` command.
5. On each server, set the `INFORMIXDIR` environment variable to the full path name for the target installation.
6. Ensure that all of the necessary configuration files are available in the target installation.
7. Start the servers in the cluster and perform additional tasks in the following order:
  - a. Start the primary server by running an `oninit` command.
  - b. Wait for primary server to be in online (multi-user) mode.
  - c. Start the Connection Manager by running an `oncmsm` command.
  - d. Start the HDR secondary server by running an `oninit` command.



- e. Start SD servers by running an oninit command.
- f. Start RS servers by running an oninit command.

## Migrating an offline cluster to a new major version

When migrating from Informix 11.70 or a later version to 15.0.0, you stop all servers in the cluster during a period when downtime is acceptable. While the cluster is offline, you upgrade all the servers. You then start the primary followed by the secondary servers. You need not rebuild the cluster or clone the servers.

### Before you begin

Prerequisites:

- Complete the steps in [Preparing for migration on page 8](#).
- Complete the steps in [Preparing to migrate, upgrade, or revert clusters on page 27](#).
- Perform all migration operations as user **informix**.

### About this task

Be sure to stop and start the servers in the cluster in the order shown in the following procedure.

To migrate to a new version with high-availability clusters:

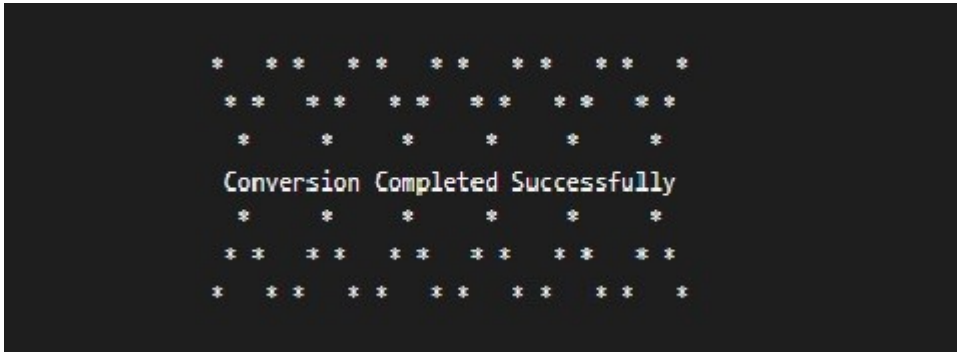
1. Stop the Connection Manager by issuing the `oncmsh -k connection_manager_name` command.
2. If you are using a High-availability Data Replication (HDR) secondary server as a backup server in case of contingencies:
  - a. Quiesce the primary server by issuing an `onmode -sy` command to prevent user connections to the server.
  - b. Force a checkpoint by issuing an `onmode -c` command on the primary server.



**Important:** Do not proceed to the next step until that checkpoint has been processed on all secondary servers.

3. Stop the secondary servers in the cluster in the following order:
  - a. If you have remote standalone (RS) secondary servers, stop them by issuing the `onmode -ky` command.
  - b. If you have shared disk (SD) servers, stop them by issuing the `onmode -ky` command.
  - c. If you have an HDR secondary server, stop it by issuing the `onmode -ky` command.
4. Stop the primary server by issuing the `onmode -ky` command.
5. On each server, set the `INFORMIXDIR` environment variable to the full path name for the target installation.
6. Ensure that all of the necessary configuration files are available in the target installation.
7. Optional: Enable quick reversion to a consistent restore point if the migration fails. Do this by setting the `CONVERSION_GUARD` and `RESTORE_POINT_DIR` configuration parameters. (For more information, see [Configuring for recovery of restore point data in case an upgrade fails on page 11](#).)
8. Start the primary server by issuing an `oninit` command.

9. Ensure that conversion of the primary was completed successfully and that the server is in multi-user mode. You should see the following in the message log:



10. For HDR and RSS servers,
- Enable quick reversion to a consistent restore point if the migration fails. Do this by setting the `CONVERSION_GUARD` and `RESTORE_POINT_DIR` configuration parameters. (For more information, see [Configuring for recovery of restore point data in case an upgrade fails on page 11.](#))
  - Start secondary server by issuing an `oninit` command.
  - Check `onstat -g cluster` output at primary server, and once HDR/RSS secondary server connected to primary server, force a checkpoint at primary server using '`onmode -c`' command and make sure checkpoint is replayed at HDR/RSS secondary server.
  - Note that `onmode -b` is not an available reversion tool for HDR/RSS secondaries, and that rebuilding these secondaries in the lower version is required after reverting the primary, should that become necessary.
11. Start the Connection Manager by issuing an `oncmsm` command.
12. Start SD secondary servers by issuing `oninit` commands.
13. Perform any additional standard migration tasks described in [Migrating to the new version of Informix on page 45](#) and in [Completing required post-migration tasks on page 54](#).
14. If you have Connection Manager configuration files that you created with a version of Informix® Client Software Development Kit (Client SDK) that is prior to version 3.70.xC3, you must convert the files. You must convert them because the older files are incompatible with the current version of the Connection Manager. For more information, see [Converting older formats of the Connection Manager configuration file to the current format on page](#) .

## Results

The in-place migration of all servers in the cluster is now complete. They are now running in what is called [compatibility mode on page 50](#). Legacy applications may be tested in this mode and should work as they did in the previous version. Reversion of the primary server, if it becomes necessary, is allowed while in compatibility mode. New applications may also be tested in compatibility mode. Be aware, however, that many 15.0.0 features are not enabled.

In order to enable all version 15.0.0 features one must move the primary server [out of compatibility mode and into native mode on page 35](#).



**Note:**



Reversion is not supported for HDR and RSS secondary servers. To revert a cluster you must shut down all secondary servers, revert the primary server, and then rebuild the cluster by recreating all secondary servers. You can use the **ifxclone** utility to recreate secondary servers.

## Moving Your Migrated Version 15.0.0 Instance from Compatibility Mode to Native Mode

To enable all version 15.0.0 features one must move the instance into *native mode*.

In-place migration to Version 15.0.0 is a two step process. The first step converts critical disk structures and places your instance in [compatibility mode on page 50](#). While in this mode your legacy applications may be tested, and they should behave as they did in the prior version. Reversion to that prior version, should it become necessary, is allowed while in compatibility mode. New applications may also be tested. Be aware, however, that while in compatibility mode many Version 15.0.0 features such as larger tables, chunks, and page sizes are not supported. To enable all Version 15.0.0 features one must move the instance into *native mode*, after which it is no longer possible to revert the instance in-place. Nor is it possible to return to compatibility mode without physically restoring an archive.

Moving your instance from compatibility mode to native mode may be done one of several ways:

1. Take a whole system level-0 backup using onbar, and include the -IV1 argument on the command line, as in:

```
onbar -b -w -L 0 -IV1
```

2. Take a level-0 backup using ontape, and include the -IV1 argument on the command line, as in:

```
ontape -s -L 0 -IV1
```

3. Take a "fake" archive using onbar, and include the -IV1 argument on the command line, as in:

```
onbar -b -F -IV1
```

4. Take a "fake" archive using ontape, and include the -IV1 argument on the command line, as in:

```
onmode -wf TAPEDEV=/dev/null
```

```
ontape -s -L 0 -IV1
```

5. Through SQL, take a "fake" archive using the sysadmin "archive fake" command, and include the optional IV1 keyword, as in:

```
DATABASE sysadmin;
```

```
EXECUTE FUNCTION task("archive fake IV1");
```

### Results

Taking a real archive during this step is recommended. Physically restoring this archive, should it become necessary, would return the instance to compatibility mode provided no logs are then rolled forward.

The act of moving a primary instance from compatibility mode to native mode generates a log record that is then propagated to all HDR, RSS, and SDS secondaries, moving them into native mode as well.

Once in native mode all Version 15.0.0 features are enabled.

## Rolling upgrade of an online cluster with Enterprise Replication

You can perform a rolling upgrade in a high-availability cluster by temporarily converting the primary and secondary servers to stand-alone Enterprise Replication servers. The upgrade occurs without incurring any downtime because Enterprise Replication supports replication between different versions of the server software. You can use this approach to upgrade to a new major version, or to apply fix packs or interim fixes (PIDs).

### Before you begin

During this procedure, you convert the primary server and the secondary servers to standalone Enterprise Replication servers. You then upgrade the software on the secondary server, stop Enterprise Replication, and then clone the server using the `ifxclone` command.

See the following link for additional information about upgrading clusters: <http://www.ibm.com/developerworks/data/library/techarticle/dm-1012rollingupgrade/index.html>.

The following prerequisites apply when upgrading software on a cluster:

- Non-logged databases are not supported.
- Raw or unlogged tables are not supported.
- Typed tables are not supported unless the typed table contains a primary key.
- UDTs that do not support Enterprise Replication are not supported.
- The `CDR_QDATA_SBSPACE` configuration parameter must be set on both the primary and secondary servers.
- The `sqlhosts` file must define a server group.
- The primary and secondary servers must belong to different groups.
- For versions of Informix® software earlier than 11.50.xC7, converting a primary and secondary server pair to Enterprise Replication is not supported if a table does not have a primary key.
- For versions of Informix® software earlier than 11.10, the `sec2er` command is not supported.

If you are upgrading Informix® software version 11.50.xC7 or later, the `sec2er` command adds a primary key to any table that does not already have one defined. For large tables, adding the primary key can take a long time, during which you will not see any server activity. In addition, the `sec2er` command requires exclusive access to the table while adding the primary key and user transactions will be blocked from accessing the table. You might want to manually create primary keys on any large table before running the `sec2er` command. If you have tables that were created with the `DELIMIDENT` environment variable set, and the tables do not have primary keys, then you must manually create the primary keys for those tables before running the `sec2er` command.

### About this task

There are different steps involved in the upgrade process depending on whether you are using the Connection Manager:

- For high-availability clusters that use the Connection Manager to redirect user connections:

There are two options you can choose based on your requirements. You can:

- Add a new Connection Manager instance to manage user connections while the cluster is upgraded. This involves configuring a new 3.70 Connection Manager instance that supports Enterprise Replication and has corresponding changes to the `sqlhosts` file or other connection mechanisms for user applications. If users already have a Connection Manager group support infrastructure to manage their user connections, they can easily add the new Connection Manager for Enterprise Replication to their existing Connection Manager group to ensure that no user connection downtime occurs during the upgrade process.
  - Use your existing cluster Connection Manager or Connection Manager groups throughout the upgrade process, without making any changes to the Connection Manager configuration, applications, or application connection mechanisms. This option has a 10-second down time for user connections, but if that is acceptable, you can avoid the overhead of adding a new Connection Manager instance and the configuration changes that go with it.
- For clusters not using the Connection Manager to redirect user connections:

Users must take steps to move user connections to the appropriate servers during the upgrade process.

### Performing a server upgrade when the Connection Manager is not in use

In this example, the terms *server1* and *server2* refer to server names rather than machine names.

Some additional steps are required to upgrade Informix® software:

1. On the primary server (**server1**), perform a check to see whether the servers can be split into Enterprise Replication servers by running the following command:

```
cdr check sec2er -c server1 --print server2
```

The command examines the primary and secondary servers and determines if it is possible to convert them to Enterprise Replication. The command displays warnings and errors that explain conditions that may prevent the servers from converting to Enterprise Replication. The `--print` option prints the commands that will be run when the `cdr start` command runs. You should fix any warnings or errors and then run the command again before performing the next step.

2. Run the following command from an Informix® 11.70 or later server:

```
cdr start sec2er -c server1 server2
```

The `sec2er` command converts the primary and secondary servers into standalone servers and configures and starts Enterprise Replication. Enterprise Replication keeps the data on the servers synchronized; however, any table created after the `sec2er` command is run will not be replicated.

3. On the former secondary server (*server2*), upgrade the Informix® software. The steps to upgrade the server are as follows:
  - a. Stop replication by running the following command:

```
cdr stop
```

- b. Back up the logical logs:

```
ontape -a
```

- c. Stop the server that contains the older version of the Informix® software:

```
onmode -kuy
```

- d. Log on to the server with the newly installed Informix® software.

- e. Start the server and let the conversion complete successfully :

```
oninit
```

- f. Run the concdr.sh script to convert the syscdr database from the old software version to the new version:

```
concdr.sh old_version new_version
```

- g. Start replication on the former secondary server (**server 2**) after it has been upgraded:

```
cdr start
```

Because Enterprise Replication supports replication between dissimilar versions of the server software, the upgraded secondary server (**server2**) replicates data with the former primary server (**server1**), so that data updates are replicated on both servers.

4. Move client application connections from the former primary server (**server1**) to the upgraded server (**server2**).
5. On the primary server (**server1**) use the onmode -k command to take the database server to offline mode.

```
onmode -k
```

6. On the former secondary server (**server2**) run the following command to stop Enterprise Replication:

```
cdr stop
```

7. You can now clone the upgraded server to set up the other secondary servers in your cluster. Clone the newly upgraded server (**server2**) by running the ifxclone utility on **server1**. Use the -d (disposition) parameter to create a standalone, RSS, or HDR secondary server. In the following examples, assume that the TCP/IP address for **server1** is 192.168.0.1 on port 123, and the address for **server2** is 192.168.0.2 on port 456.

- To create a standalone server:

```
ifxclone -T -S server2 -I 192.168.0.2 -P 456 -t server1  
-i 192.168.0.1 -p 123
```

- To create an RS secondary server, specify the disposition by using the **-d RSS** option:

```
ifxclone -T -S server2 -I 192.168.0.2 -P 456 -t server1  
-i 192.168.0.1 -p 123 -d RSS
```

- To create an HDR secondary server, specify the disposition by using the **-d HDR** option:

```
ifxclone -T -S server2 -I 192.168.0.2 -P 456 -t server1  
-i 192.168.0.1 -p 123 -d HDR
```

At this point, the cluster is running on the upgraded server. Clients can move applications from server2 if necessary.

### Performing a server upgrade when the Connection Manager is in use

Refer to the following steps when clients are using the Connection Manager without a Connection Manager group defined in the existing setup.

For this example, assume that the following Connection Manager configuration file is defined:

```
NAME cm1
LOG 1
LOGFILE /tmp/cm1.log

CLUSTER cluster1
{
    INFORMIXSERVER ids1,ids2
    SLA oltp DBSERVERS=primary
    SLA webapp DBSERVERS=HDR
    SLA report DBSERVERS=(primary,HDR)
    FOC ORDER=ENABLED PRIORITY=1
}
```

1. On the primary server (**server1**), perform a check to see whether the servers can be split into Enterprise Replication servers by running the following command:

```
cdr check sec2er -c server1 --print server2
```

When the above command is run, the primary and secondary servers are examined to determine whether it is possible to convert them to Enterprise Replication. The command displays warnings and errors that explain conditions that might prevent the servers from converting to Enterprise Replication. The `--print` option prints the commands that will be run when the `cdr start sec2er` command runs. You should fix any warnings or errors and then run the command again before performing the next step.

2. Reload the Connection Manager so that it directs all client connections to the primary server. Here is the revised Connection Manager configuration file:

```
NAME cm1
LOG 1
LOGFILE /tmp/cm1.log

CLUSTER cluster1
{
    INFORMIXSERVER ids1,ids2
    SLA oltp DBSERVERS=primary
    SLA webapp DBSERVERS=primary,HDR
    SLA report DBSERVERS=primary,HDR
    FOC ORDER=ENABLED PRIORITY=1
}
```

3. Run the following command from an Informix® 11.70 or later server:

```
cdr start sec2er -c server1 server2
```

The `sec2er` command converts the primary and secondary servers into standalone servers and configures and starts Enterprise Replication. Enterprise Replication keeps the data on the servers synchronized; however, any table created after the `sec2er` command is run will not be replicated.

4. On the former secondary server (**server2**), upgrade the Informix® software.

Because Enterprise Replication supports replication between dissimilar versions of the server software, the upgraded secondary server (**server2**) replicates data with the former primary server (**server1**), so that data updates are replicated to both servers.

5. Move client application connections from the former primary server (**server1**) to the upgraded server (**server2**).
  - a. Create a new Connection Manager configuration file for Enterprise Replication. The following shows a sample Enterprise Replication Connection Manager configuration file. The SLA names are same as for cm1:

```
NAME cm2
LOG 1
LOGFILE /tmp/cm2.log
MACRO list=g_server2,g_server1

REPLSET replset_1
{
  INFORMIXSERVER g_server1,g_server2
  SLA oltp DBSERVERS=${list}
  SLA webapp DBSERVERS=${list}
  SLA report DBSERVERS=${list}
}
```

The Enterprise Replication Connection Manager must define a replicate set that includes all replicates that are generated by the sec2er command. You can see the list of replicates by running the following command:

```
cdr list repl
```

You create a replicate set by running the following command:

```
cdr def replset replset_name repl1 repl2 ...
```

In the above example, repl1 and repl2 are replicates created by the sec2er command.

- b. Halt the **cm1** Connection Manager instance and load the cm2 instance.

Performing the above step ensures that client connections are redirected to **group\_2** (because **server2** belongs to **group\_2**).

Here is a sample `sqlhosts` file:

#dbservername	nettype	hostname	servicename	options
g_server1	group	-	-	i=10
ids1	onsoctcp	host1	port1	g=g_server1
g_server2	group	-	-	i=20
ids2	onsoctcp	host2	port2	g=g_server2
oltp	onsoctcp	host1	port3	
webapp	onsoctcp	host1	port4	
report	onsoctcp	host1	port5	

6. On the primary server (**server1**) use the `onmode -k` command to take the database server to offline mode.

```
onmode -k
```

7. You can now clone the upgraded server to the other secondary servers in your cluster.

At this point in the upgrade process, the high-availability cluster is running on the upgraded server.



8. Shut down the cm2 Connection Manager instance and start the cm1 instance.
9. On the former secondary server (**server2**) run the following command to stop Enterprise Replication:

```
cdr stop
```

## Errors and warnings generated by the sec2er command

The sec2er command checks several conditions before converting a primary and secondary server pair to an ER system. The following conditions are checked by the sec2er command; ER conversion will take place only if the following conditions are met:

- The group definition must use the `i=` option.
- The CDRID option must be the same on both the primary and secondary servers (The CDRID is the unique identifier for the database server in the Options field of the `sqlhosts` file).
- The `sqlhosts` files on the primary server must match the `sqlhosts` file on the secondary server. The sec2er command checks only the lines in the `sqlhosts` files that must to match to support ER.
- The database must not contain a typed table without a primary key.
- User-defined types (UDT) must have ER support.
- Tables must not be protected with label-based access control (LBAC).
- A secondary server must be defined.
- An sbspace for a stable queue must exist.
- You must be running Informix® version 11.00 or later.

The following warnings might occur. Warnings do not always indicate a problem but should be addressed. A warning is generated if any of the following are true:

- The CDR\_SERIAL configuration parameter is not set.
- The values for the CDR\_SERIAL configuration parameter are the same on both the primary and secondary servers. Identical values can cause conflicts.
- The database has sequence generators. Because sequence generators are not replicated, if you replicate tables using sequence objects for update, insert, or delete operations, the same sequence values might be generated on different servers at the same time, leading to conflicts.
- The database is not logged.
- A table is not logged.
- The DBSPACE is more than 70-percent full.

If the `cdr start sec2er` command fails or is interrupted, you might see a message that is similar to this message:

```
ERROR: Command cannot be run on pre-11.70 instance if ER is already running
```

If you receive this error, and you do not have the Connection Manager running, remove replication by running the `cdr delete server` command on both servers and then run the `cdr start sec2er` command again. If you have the Connection Manager running, then perform the following steps:

1. Shut down the Connection Manager.
2. Shut down ER using one of the following commands:
  - `cdr delete server group`
  - `cdr delete server server`

3. Start the Connection Manager.
4. Restart the **sec2er** command:

```
cdr start sec2er
```

## Reverting clusters

If you have a high-availability cluster, you must complete additional tasks when you revert from the new version of . You must revert only the primary database server.

### About this task

The server automatically removes secondary servers during reversion. After reversion on the primary server is complete, you must recreate all HDR, RS, and SD secondary servers in a high-availability cluster.

Prerequisites:

- Determine if you can revert. See information in [Review the database schema prior to reversion on page 61](#).
- Complete the steps in [Preparing to migrate, upgrade, or revert clusters on page 27](#).
- Perform all reversion operations as user **informix**.

When you revert clusters, be sure to stop and start the servers in the cluster in the order shown in the following procedure.

To revert high-availability clusters, follow steps 1 through 9. After step 9, you can either complete step 10 which in turn completes reverting of high-availability cluster for you or alternatively, you can skip step 10 and follow steps 11 through 13. These are the two alternate paths available to follow after step 9.

1. Stop the Connection Manager by issuing the `oncmsm -k connection_manager_name` command.
2. If you are using a High-availability Data Replication (HDR) secondary server as a backup server in case of contingencies:
  - a. Quiesce the primary server by issuing an `onmode -sy` command to prevent user connections to the server.
  - b. Force a checkpoint by issuing an `onmode -c` command on the primary server.
3. Stop the servers in the cluster and perform the following tasks in the following order:
  - a. If you have remote standalone (RS) servers, stop them by issuing the `onmode -ky` command.
  - b. If you have shared disk (SD) servers, stop them by issuing the `onmode -ky` command.
  - c. If you have a High-availability Data Replication (HDR) secondary server, stop it by issuing the `onmode -ky` command.

- d. Revert the standard server by issuing an `onmode -b target_IDS_version` command.
  - e. Verify that reversion was successful and the server was stopped. If the reversion was not successful, check the message log for error messages, take appropriate action, and restart reversion.
4. On each server, set the `INFORMIXDIR` environment variable to the full path name for the target installation.
  5. Ensure that all of the necessary configuration files are available in the target installation.
  6. Perform any additional database server reversion tasks, as described in [Reverting from Version 15.0.0 on page 65](#).
  7. Start the primary server by issuing an `oninit` command.
  8. Start the Connection Manager by issuing an `oncmsm` command
  9. Start SD secondary servers by issuing `oninit` commands.
  10. Use the [ifxclone utility](#) to reestablish the RS and HDR servers if necessary. If you have executed step 10 your reversion is complete. As an alternative to `ifxclone` you may execute steps 11-13 below.
  11. Back up all logs. Then use ON-Bar or **ontape** to make a level-0 backup on the primary server to use to reestablish the RS and HDR servers if necessary.
  12. If you have RS secondary servers:
    - a. Add RS entries on the primary server by issuing `onmode -d add RSS rss_server_name` commands.
    - b. Start the RS secondary servers with level-0 restore operations from the level 0 backup that was made on the primary server after reversion.
    - c. On the RS secondary servers, run the `onmode -d RSS primary_server_name` command, and wait for the "RSS secondary server operational" message to appear after each command.
  13. If you have an HDR secondary server:
    - a. Reestablish the HDR pair on the primary server by issuing an `onmode -d primary hdr_secondary_server_name` command.
    - b. Start the HDR secondary server with level-0 restore operations from the level 0 backup that was made on the primary server after reversion.
    - c. On the HDR secondary server, run the `onmode -d secondary primary_server_name` command, and wait for the "HDR secondary server operational" message to appear after each command.

## Results

The reversion of all servers in the cluster is now complete.

## Restoring clusters to a consistent point

You can restore the primary server in a high-availability cluster to a consistent point after a failed upgrade.

### Before you begin

Prerequisites:

- Before you began the upgrade, you must have enabled quick reversion, according to information in [Configuring for recovery of restore point data in case an upgrade fails on page 11](#).
- The server must be offline.

### About this task

To restore the primary server in a cluster to a consistent point after a failed upgrade:

Run the **onrestorept** utility according to [these instructions on page 49](#).

### What to do next

Alternatively, if you backed up your primary server or you prepared for using the RSS or HDR secondary as a backup server before you upgraded and the upgrade fails, you can take other steps to restore the cluster.

## Restoring a cluster from a backup archive

If you backed up the primary server before you migrated or reverted the cluster, you can restore the primary server from the backup archive if migration or reversion fails. After you restore the primary server, you must recreate the other servers in the high-availability cluster.

### Before you begin

Prerequisite: You made a level-0 backup archive of the primary server before migration or reversion.

### About this task

To restore a cluster from a level-0 backup archive:

1. Point your INFORMIXDIR, PATH, and any other relevant environment variables to the directory in which the original version of was installed before you migrated or reverted.
2. Using the level-0 backup archive, perform a full restore of your primary server.
3. Recreate the rest of your high-availability cluster.

## Restoring a cluster from the HDR secondary server

You can restore the primary server from the High-availability Data Replication (HDR) secondary server that you prepared to use as a backup server before you migrated or reverted the cluster. After you restore the primary server, you must recreate the other servers in the high-availability cluster.

### Before you begin

Prerequisite: You prepared the HDR secondary server to use as a contingency backup server, according to information in [Preparing to migrate, upgrade, or revert clusters on page 27](#).

To restore a cluster from the HDR secondary server:

1. Start the HDR secondary server by running an oninit command.
2. Change the secondary server to the primary server by running the onmode -d make primary *hdr\_server\_name* command.
3. If the server is in quiescent mode, change it to multi-user mode by running an onmode -m command.

4. Make a level-0 backup using the ON-Bar or **ontape** utility.
5. Recreate the rest of the high-availability cluster.

## Migrating to 15.0.0

When you migrate to a new version of Informix®, you must complete required migration and post-migration tasks.

## Migrating to the new version of Informix®

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

### Before you begin

- Read the release notes and the machine notes for any new information.
- Complete the steps in [Preparing for migration on page 8](#).

### About this task

To upgrade a Informix® non-root installation, you must run the installation program as the same user who installed the product being upgraded.

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



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

Review and complete all tasks that apply:

1. [Installing the new version of Informix on page 46](#).
2. [Setting environment variables on page 46](#).
3. [Customizing configuration files on page 47](#).
4. [Adding Communications Support Modules on page 48](#).
5. [Installing or upgrading any modules on page 48](#).
6. [Starting the new version of Informix on page 48](#).

### Results

When the migration starts, the `online.log` displays the message `Conversion from version <version number> Started..` The log continues to display start and end messages for all components. When the migration of all components is complete, the message `Conversion Completed Successfully` appears. For more information about this log, see [Migration status messages on page 46](#).

### What to do next

After migration, see [Completing required post-migration tasks on page 54](#) for information about preparing the new server for use.

If the log indicates that migration failed, you can either:

- Install the old database server and restore your database from a level-0 backup.
- Run the `onrestorept` utility to back out of the upgrade and restore files to a consistent state without having to restore from a backup. You can run this utility only if you set the configuration parameters that enable the utility. See [Restoring to a previous consistent state after a failed upgrade on page 49](#).

## Installing the new version of Informix®

Install and configure the new version of Informix®.

If possible, migrate on a database server dedicated to testing your migration before you migrate on your production database server.

Decide whether to upgrade on the same computer, known as an in-place migration, or on a different computer, known as a non-in-place migration and follow the appropriate process.



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

## Migration status messages

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

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

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

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

## Setting environment variables

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

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.

The setting of the **GL\_USEGLU** environment variable must match between the source and target server during migration.

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

## Customizing configuration files

When you initialize the new version of Informix®, which contains a new `onconfig.std` file, use the same configuration that the old database server used. After you observe the performance of new version, you can examine the new file for new configuration parameters that you might want to use and can you start to use the new and changed configuration parameters.

Set the **ALARMPROGRAM** configuration parameter to either nothing or **no\_log.sh** to prevent the generation of errors if the logical log fills during the migration. For more details, see [Starting the new version of Informix on page 48](#). After the migration, change the value of **ALARMPROGRAM** to **log\_full.sh**.

If the **ALARMPROGRAM** configuration parameter is set to the script **alarmprogram.sh**, set the value of **BACKUPLOGS** in **alarmprogram.sh** to `N`.



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

For information about how to configure Informix®, see your *HCL® Informix® Administrator's Guide*. For information about how to tune the configuration parameters, see the *HCL® Informix® Performance Guide*.

## Using SSL/TLS database connections

If you are upgrading database clients that use SSL/TLS connections to Client SDK 4.50.xC4W1 or newer, you may need to migrate their client keystores. For more information, see [Configuring a client for SSL connections on page 48](#).

To perform keystore migration:

1. If your database client installation is co-located with the database server installation, the database client continues to use GSKit as encryption library. In this case, keystore migration is not necessary.
2. If your database client uses a stand-alone installation of Client SDK 4.50.xC4W1 or newer, and you choose to install OpenSSL for the encryption library, rather than GSKit, the following applies:
  - a. Ensure to have an appropriate version of OpenSSL installed before you install Client SDK 4.50.xC4W1 or newer.
  - b. If your client keystore has the GSKit-proprietary format "CMS" (file extension "\*.kdb"), then this keystore needs to be converted to a **PKCS#12** keystore. As the CMS format is GSKit-specific, you need the GSKit command "gsk8capicmd" (or "gsk7capicmd") in order to convert the keystore.

Use a command like:

```
gsk8capicmd -keydb -convert -db KEYSTOREFILE.kdb -stashed  
-old_format cms -new_db KEYSTOREFILE.p12 -new_pw PASSWORD  
-new_format pkcs12 -pqc false
```

- c. Using the stand-alone client installation, create a stash file with the keystore password to use with OpenSSL. Use the new utility "onkstash" contained with Client SDK 4.50.xC4W1 (or newer) to stash the keystore password:

```
onkstash KEYSTOREFILE.p12 PASSWORD
```

## Adding Communications Support Modules

Communication Support Modules - ENCCSM and SPWDSCM are removed in Informix Server 14.10.xC9. These modules are not included in the distribution.

In Version 15.0.0, you should use Transport Layer Security (TLS)/Secure Sockets Layer (SSL) instead.

## Installing or upgrading any modules

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

Register the modules after you initialize the database server.

When you install Informix®, the built-in extensions, such as TimeSeries, are installed and registered automatically. You do not need to perform any actions to upgrade these modules, nor do you need to unload and load any data during migration.



**Important:** If the sysadmin database does not exist or the Scheduler is turned off, automatic registration does not occur. You must register each extension that you want to use by running the SYSBldPrepare() function.

## Starting the new version of Informix®

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

**Prerequisite:** If you installed Informix® as user **root**, you must switch to user **informix** before starting the server.



**Important:** Informix® 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, Informix® stops and indicates that you must back up the logical logs. After you back up the logical logs, the database server can finish building the **sysmaster** database. You cannot use to back up the logical logs because the database has not been converted yet. If you have ALARMPROGRAM set to **log\_full.sh** in the ONCONFIG configuration file, errors are generated as each log file fills during the migration. Set the value of ALARMPROGRAM to either nothing or **no\_log.sh**





so that these errors are not generated. If your logical log does fill up during the migration, you must back it up with **ontape**, the only backup tool you can use at this point. Issue the **ontape -a** command.

Start the new version of Informix® for the first time by running **oninit** command on UNIX™ or by using the **Service** control application on Windows™.

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

Once conversion has completed successfully you will see the following in the message log:

```

* * * * *
* * * * *
* * * * *
Conversion Completed Successfully
* * * * *
* * * * *
* * * * *
```

At that point your Version 15.0.0 instance will be in [compatibility mode on page 50](#). In this mode the server uses the legacy infrastructure, which has a version of 0. The **onstat** banner displays the infrastructure version and indicates whether you are in compatibility mode:

```

IBM Informix Dynamic Server Version 15.0.0.0 -- On-Line -- Up 00:50:06 -- 715340 Kbytes
2024-11-08 07:55:29 -- Infrastructure Version: 0 (Compatibility mode)
```

While in compatibility mode, your legacy applications may be tested and they should work as they did in the prior version. Should reversion to that prior version become necessary, this may be done from compatibility mode using the **onmode -b** command.

A production system may be run in compatibility mode, but most Version 15.0 features will not be enabled. To enable all Version 15.0 features you must [move your instance to native mode on page 35](#).

## Restoring to a previous consistent state after a failed upgrade

If the **CONVERSION\_GUARD** configuration parameter is enabled and an upgrade fails, you can run the **onrestorept** command to undo the changes that are made during the upgrade and restore to a consistent state.

### Before you begin

The following prerequisites must be met:

- The directory that is specified by the **RESTORE\_POINT\_DIR** configuration parameter contains a subdirectory with the files that were stored during the failed upgrade attempt. The name of this subdirectory will be your instance's **SERVERNUM** setting, e.g. "15".
- The server must be offline before you run the **onrestorept -y** command.

## About this task

Note that with your environment set up for your new instance, you will execute the *previous Informix version's onrestorept utility*. The following steps assume that an OLDINFORMIXDIR environment variable has been set to your previous version's informix directory.

To restore the server to a previous consistent state after a failed upgrade:

Run the \$OLDINFORMIXDIR/bin/onrestorept -y command, which displays prompts while the command runs. If you do not specify -y, you must respond to every prompt that appears.



**Important:** Do not start the server until the \$OLDINFORMIXDIR/bin/onrestorept utility finishes running. Starting the server before the \$OLDINFORMIXDIR/bin/onrestorept utility finishes running can damage the server, requiring the server to be restored from a backup copy.

## What to do next

After you restore the server to a consistent state, you can resume work in the source version of the server or find and fix the problem that caused the upgrade to fail.



**Important:** To start Enterprise Replication after the \$OLDINFORMIXDIR/bin/onrestorept utility restores the database server to a consistent state, you must use the cdr cleanstart command.

Before you attempt another upgrade, run the \$OLDINFORMIXDIR/bin/onrestorept -c command to remove the files in the directory that is specified in the RESTORE\_POINT\_DIR configuration parameter. After a successful upgrade, the server automatically deletes restore point files from that directory.

## Compatibility Mode And Native Mode

Informix Version 15.0.0 supports larger tables, chunks, and pages. This support is dependent on the enablement of a new infrastructure that makes low-level disk changes. These changes are so fundamental and extensive that reverting them in-place is not practical. In order to satisfy the potential need for reversion following a migration from an older version, Informix 15.0.0 requires explicit enablement of the new infrastructure by the administrator.

Immediately after the successful completion of an in-place migration from an older Informix server version, an Informix 15.0.0 instance is in *compatibility mode*. While in this mode the legacy infrastructure, referred to as Infrastructure Version 0, is in effect. The purpose of compatibility mode is to allow the testing of applications against the Version 15.0.0 server while still supporting reversion if necessary. The drawback of compatibility mode is that many Version 15.0.0 features are disabled.

Enabling all Version 15.0.0 features requires [switching your instance from compatibility mode to native mode on page 35](#). Once an instance is in native mode it cannot be reverted in-place to an older Informix version. Nor can it be returned to compatibility mode without the physical restoration of an archive. In this case *no logical restore* may be performed, otherwise the recovered instance will be switched to native mode by a record in the logical logs.

## Methods of determining the current mode

### Using onstat command

- This onstat output indicates the instance is in compatibility mode.

```
onstat -
IBM Informix Dynamic Server Version 15.0.0.0 -- On-Line -- Up 18:23:33 -- 2998160 Kbytes
2024-11-27 10:09:34 -- Infrastructure Version: 0 (Compatibility mode)
```

- This onstat output indicates the instance is in native mode.

```
onstat -

IBM Informix Dynamic Server Version 15.0.0.0 -- On-Line -- Up 00:00:44 -- 2998160 Kbytes
2024-11-27 10:23:57 -- Infrastructure Version: 1
```

### Using oncheck command

- This oncheck -pr output indicates the instance is in compatibility mode.

```
oncheck -pr

Validating IBM Informix Dynamic Server reserved pages

Validating PAGE_PZERO...

Identity                      IBM Informix Dynamic Ser
                              ver Copyright 2001, 2024
                              IBM Corporation

Database system state         0
Database system flags         0x9
64-bit server
BigChunk page flags are not in use
Version 1 infrastructure has not been enabled
Page Size                     2048 (b)
Date/Time created              11/27/2024 10:11:27
Version number of creator      50
UID of rootdbs creator         200
Index Page Logging             OFF
HA Disk Owner                  <null>
...
...
```

- This oncheck -pr output indicates the instance is in native mode. Instead of “Version 1 infrastructure enabled” oncheck -pr will display “**Version 1 infrastructure only**” for an instance that was initialized from scratch in Version 15.0.0, as opposed to migrated from an earlier version.

```
oncheck -pr

Validating IBM Informix Dynamic Server reserved pages

Validating PAGE_PZERO...
```

Identity	IBM Informix Dynamic Server Copyright 2001, 2024 IBM Corporation
Database system state	0
Database system flags	0x309
64-bit server	
BigChunk page flags are not in use	
Version 1 infrastructure enabled	
Page Size	2048 (b)
Date/Time created	11/27/2024 10:23:22
Version number of creator	50
UID of rootdbs creator	200
Index Page Logging	OFF
HA Disk Owner	<null>

## Using SQL

- When in compatibility mode, this flag value will **not** contain 0x1000000.

```
dbaccess sysmaster -<<END
select hex(value) from sysshmhdr where name = "sysflags";
END

Database selected.

(expression)

0x00000000000001041

1 row(s) retrieved.

Database closed.
```

- When in native mode, this sysflags value **will** contain 0x1000000.

```
dbaccess sysmaster -<<END
select hex(value) from sysshmhdr where name = "sysflags";
END

Database selected.

(expression)

0x0000000001001041

1 row(s) retrieved.

Database closed.
```

## Using Message log

When in compatibility mode, the online log will contain messages like this:

```

...
...
10:11:28 On-Line Mode
10:11:28 (Compatibility mode)
...
...
10:17:37 Quiescent Mode
10:17:37 (Compatibility mode)

```

## Moving Your Migrated Version 15.0.0 Instance from Compatibility Mode to Native Mode

To enable all version 15.0.0 features one must move the instance into *native mode*.

In-place migration to Version 15.0.0 is a two step process. The first step converts critical disk structures and places your instance in [compatibility mode on page 50](#). While in this mode your legacy applications may be tested, and they should behave as they did in the prior version. Reversion to that prior version, should it become necessary, is allowed while in compatibility mode. New applications may also be tested. Be aware, however, that while in compatibility mode many Version 15.0.0 features such as larger tables, chunks, and page sizes are not supported. To enable all Version 15.0.0 features one must move the instance into *native mode*, after which it is no longer possible to revert the instance in-place. Nor is it possible to return to compatibility mode without physically restoring an archive.

Moving your instance from compatibility mode to native mode may be done one of several ways:

1. Take a whole system level-0 backup using onbar, and include the -IV1 argument on the command line, as in:

```
onbar -b -w -L 0 -IV1
```

2. Take a level-0 backup using ontape, and include the -IV1 argument on the command line, as in:

```
ontape -s -L 0 -IV1
```

3. Take a "fake" archive using onbar, and include the -IV1 argument on the command line, as in:

```
onbar -b -F -IV1
```

4. Take a "fake" archive using ontape, and include the -IV1 argument on the command line, as in:

```
onmode -wf TAPEDEV=/dev/null
```

```
ontape -s -L 0 -IV1
```

5. Through SQL, take a "fake" archive using the sysadmin "archive fake" command, and include the optional IV1 keyword, as in:

```
DATABASE sysadmin;
```

```
EXECUTE FUNCTION task("archive fake IV1");
```

### Results

Taking a real archive during this step is recommended. Physically restoring this archive, should it become necessary, would return the instance to compatibility mode provided no logs are then rolled forward.

The act of moving a primary instance from compatibility mode to native mode generates a log record that is then propagated to all HDR, RSS, and SDS secondaries, moving them into native mode as well.

Once in native mode all Version 15.0.0 features are enabled.

## Completing required post-migration tasks

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

To complete post-migration tasks:

1. If you **do not** use : [For ON-Bar, copy the sm\\_versions file on page 54.](#)
2. [Finish preparing earlier versions of 12.10 databases for JSON compatibility on page 54](#)
3. [Optionally update statistics on your tables after migrating on page 56.](#)
4. [Review client applications on page 56.](#)
5. [Verify the integrity of migrated data on page 58.](#)
6. [Back up Informix after migrating to the new version on page 59.](#)
7. [Tune the new version for performance and adjust queries on page 59.](#)
8. If you use specific features, you might have to perform additional post-migration tasks:
  - [Migrating with Enterprise Replication on page 23](#)
  - [High-availability cluster migration on page 25](#)
  - [Register modules on page 59](#)

### What to do next

Repeat the migration and post-migration procedures for each instance of Informix® Version 15.0.0 that you plan to run on the computer.



**Important:** Do not connect applications to a database server instance until the migration has completed successfully. If a serious error occurs during the migration, you might need to revert to the previous version of the server, restore from a level-0 backup, and then correct the problem before restarting the migration tasks.

## For ON-Bar, copy the sm\_versions file

After migration, if you plan to use a storage manager other than the HCL Informix® Primary Storage Manager, copy your previous `sm_versions` file to use with . does not use the `sm_versions` file.

If you are using other storage managers, copy your previous `sm_versions` file from the old `$INFORMIXDIR/etc` directory to the new `$INFORMIXDIR/etc` directory.

## Finish preparing earlier versions of 12.10 databases for JSON compatibility

To make databases that were created with earlier versions of Informix® 12.10 JSON compatible, you must complete some post-migration steps.

1. Prepare any databases that were created in 12.10.xC1 for JSON compatibility.

If you upgraded to 12.10.xC4 or later fixpacks, skip this step because all databases are made JSON compatible during conversion.

If you upgraded to 12.10.xC3 or 12.10.xC2, complete these steps:

- a. Run the appropriate script on the upgraded server as user **informix** or as a user with DBA privileges.

- Informix® 12.10.xC3: `convTovNoSQL1210.sql`
- Informix® 12.10.xC2: `convTovNoSQL1210X2.sql`

#### Example

For example, in Informix® 12.10.xC3, to make the **db\_name** database JSON compatible, you would run the following command as user **informix** or as a user with DBA privileges:

#### UNIX™

```
dbaccess -e db_name $INFORMIXDIR/etc/convTovNoSQL1210.sql
```

#### Windows™

```
dbaccess -e db_name %INFORMIXDIR%\etc\convTovNoSQL1210.sql
```

- b. Configure and start the wire listener.

2. If you used the wire listener in 12.10.xC2 or 12.10.xC3, after you upgrade to 12.10.xC4 or later fixpacks you must drop and recreate any indexes that you created on your collections. You do not need to drop and recreate the index that is automatically created on the `_id` field of a collection.
3. If you had binary JSON (BSON) DATE fields in your documents in 12.10.xC2, you must load the data from the external table that you created before you upgraded to a later 12.10 fix pack.

On the upgraded server:

- a. Rename the table that contains the BSON column with DATE fields.

For example, use the ALTER table statement to rename the table from **datetab** to **datetab\_original**.

- b. Create a new table that has the original table name.

For example:

```
create table datetab(j int, i bson);
```

- c. Load data into the new table from the external table that you created in your 12.10.xC2 server. During the load, convert the data from JSON format to BSON format.

For example:

```
insert into datetab select j, i::json::bson from ext_datetab;
```

- d. Verify that the data loaded successfully.
  - e. Drop the original, renamed table (for example, **datetab\_original**) after you are certain that the data loaded successfully into the new table (for example, **datetab**).
4. If you used the wire listener in 12.10.xC2 with a database that has any uppercase letters in its name, after you upgrade to a later fix pack you must update your applications to use only lowercase letters in the database name.

## Optionally update statistics on your tables after migrating

Optionally run UPDATE STATISTICS on your tables and on UDRs that perform queries, if you have performance problems after migrating to the new version of Informix®.

An unqualified UPDATE STATISTICS statement includes no additional clauses:

```
UPDATE STATISTICS;
```

By default, an unqualified UPDATE STATISTICS statement updates the statistics in LOW mode for every permanent table in the database, including the system catalog tables.

You can run an UPDATE STATISTICS statement that includes only a FOR ROUTINE clause that specifies no routine name:

```
UPDATE STATISTICS FOR ROUTINE;
```

Running this statement reoptimizes the DML statement execution plans for every SPL routine in the database that operates on local tables.

Similarly, you can substitute the keyword FUNCTION for ROUTINE in the previous example to reoptimize execution plans only for SPL routines that return at least one value, or you can substitute the keyword FUNCTION for PROCEDURE to reoptimize execution plans only for SPL routines that return no value. In these cases, the database server does not update the statistics in the system catalog tables.

You do not need to run UPDATE STATISTICS statements on C or Java™ UDRs.

## Review client applications

After you migrate a database server on the same operating system or move the database server to another compatible computer, review the client applications to take necessary actions related to migration. If necessary, recompile or modify client applications.

### General guidelines

Post migration, review the client applications and `sqlhosts` file or registry-key connections. Recompile or modify client applications if necessary. Verify that the client-application version you use is compatible with your database server version. If necessary, update the `sqlhosts` file or registry key for the client applications with the new database server information.



## Migrating ODBC applications

If you have a 64-bit ODBC application that was compiled and linked with a version of HCL Informix® Client Software Development Kit (Client SDK) that is prior to version 4.10, you must recompile the application after migrating. The SQLLEN and SQLULEN data types were changed to match the Microsoft™ 64-bit ODBC specification. Be sure to analyze any functions that take either of these data types to ensure that the correct type passes to the function. This step is crucial if the type is a pointer. Also note that in the ODBC specification, some parameters that were previously SQLINTEGER and SQLUINTEGER were changed to SQLLEN or SQLULEN.

## Large Row Identifiers and Their Implications for Old and New Clients

Immediately after the successful completion of an in-place migration, an Informix 15.0.0 instance is in [compatibility mode on page 50](#), meaning that most new features are disabled. Enabling all Version 15.0.0 features requires switching your instance from [compatibility mode to native mode on page 35](#).

If you perform a complete disk initialization using Informix Version 15.0.0, the instance will begin and always remain in native mode.

Tables created by Informix Version 15.0.0 in native mode use 8-byte row identifiers by default. The new rowid size and format allow these new tables to grow much larger than they did prior to Version 15.0.0. Tables that make use of large rowids are called [LARGE tables](#).

Because rowids are embedded in the [sqlca](#) structure, which is passed back and forth between clients and the server, applications must be prepared for a new version of sqlca if they want to take advantage of the larger rowids.

New applications compiled and linked using ClientSDK version 15.0.0 are already compatible with large rowids. Applications that communicate with the server using the Version 15.0.0 ODBC or JDBC drivers will also recognize large rowids along with the new sqlca structure.

If you want a legacy application to make use of large tables, before connecting it to a Version 15.0.0 instance in native mode you should recompile and relink the application with the [version 15 header files and libraries](#). Note that because of the changes to internal structures the Version 15.0.0 libraries have new names.

Version 15.0.0 also supports tables that use legacy 4-byte rowids. A table migrated from an earlier version of the server, for example, will always use 4-byte rowids. These are called SMALL tables. All tables created while the instance is in compatibility mode will be in small table format. Small tables may also be [explicitly created](#) while in native mode.

Applications compiled and linked with prior versions of the CSDK will continue to work with the Version 15.0.0 server, but only with small tables. If they attempt to create or read a large table they will receive a new error number:

```
-21569 This client cannot create or open tables that use large rowids.
```

Even when in native mode, all system catalogs, and all tables in system databases such as sysmaster, sysadmin, sysutils, sysuser, and sysha are in “small table” format so that both new and old clients may access them.

## ROWID Columns

When a new application selects the 'rowid' hidden column from a non-fragmented table, the type of that column is BIGINT, whether the table is large or small. A legacy application that cannot access large tables will continue to see this 'rowid' column as an INT type.

Note that a large internal rowid, when displayed in decimal format, will always be a very large negative number, e.g. -9223372036854710271.

When an application selects the 'rowid' hidden column from a large, fragmented table created using the WITH ROWIDS clause, the type of that column is BIGSERIAL. When a small fragmented table is created using the WITH ROWIDS clause, the type of the hidden 'rowid' column is SERIAL.

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

## Verify the integrity of migrated data

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

You can also verify the integrity of the reserve pages, extents, system catalog tables, data, indexes, and smart large objects, as [Table 6: Commands for verifying the data integrity on page 58](#) shows.

**Table 6. Commands for verifying the data integrity**

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

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

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



**Important:** If the value of the MAX\_FILL\_PAGES configuration parameter is 1, you must run oncheck -cD for all tables that include variable length data types (VARCHAR, NVARCHAR, and LVARCHAR), and take the suggested corrective action to avoid warnings about resetting the bitmap mode.

## Back up Informix® after migrating to the new version

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

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



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

## Tune the new version for performance and adjust queries

After backing up the new server, you can tune the database server to maximize performance. If your queries are slower after the upgrade, there are steps you can take to adjust your configuration and queries.

If your queries are slower after an upgrade, find out what changed that affects your configuration and adjust your configuration and queries as necessary:

- Compare the default values in the new `onconfig.std` file to values in your previous installation, and make any necessary adjustments.
- If you created sample queries for comparison, use them to analyze the performance differences between the old and new database servers and to determine if you need to adjust any configuration parameters or the layout of databases, tables, and chunks.
- If you changed your applications, check to see if the changes led to slower performance.
- If you changed your hardware, operating system, or network settings, determine if you need to adjust any related settings or environment variables.
- Make sure that you ran necessary update statistics after upgrading:
  - [Drop data distributions if necessary when upgrading on page](#)
  - [Optionally update statistics on your tables after migrating on page 56](#)

## Register modules

You must register any modules that you installed.

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

## Reverting from Version 15.0.0

As long as your Version 15.0.0 instance is running in [compatibility mode on page 50](#), you can revert to the version of the database server from which you migrated. When you run the reversion utility, you specify the target server for reversion

and then Informix® checks your database. If necessary, Informix® might tell you to drop new objects, before automatically converting your data into the target server.

If Informix® cannot revert a database, Informix® prevents reversion.

There are multiple methods of determining whether an instance is in compatibility mode.

#### 1. Run onstat -

```
IBM Informix Dynamic Server Version 15.0.0.0 -- On-Line -- Up 00:00:42 -- 715340 Kbytes
2024-11-08 14:54:12 -- Infrastructure Version: 0 (Compatibility mode)
```

#### 2. View the message log

```
The server will write messages like the following to the log during startup or a graceful shutdown:
08:37:23 On-Line Mode
08:37:23 (Compatibility mode)
...
...
08:46:41 Shutdown Mode
08:46:43 Quiescent Mode
08:46:43 (Compatibility mode)
```

#### 3. Run oncheck -pr

```
...
...
Database system flags          0x9
64-bit server
BigChunk page flags are not in use
Version 1 infrastructure has not been enabled
...
...
```

#### 4. Execute a sysmaster query

```
database sysmaster;
select hex(value) from sysshmhdr where name = "sysflags";
```

The hexadecimal 8-byte integer returned by that query will not contain 0x1000000 if the instance is in compatibility mode.

Normally, reversion takes only a few minutes.

If you created databases or stored procedures, or altered tables in Version 15.0.0, reversion time is longer, because you must prepare your database and data for reversion. In general the more DDL work you complete in the new version, the more time consuming the reversion. See [Preparing to revert on page 60](#) before you revert.

If you did not create databases or stored procedures and did not alter tables in Version 15.0.0 and you did not complete much work using the new server, you can run the reversion utility and modify the values of the configuration parameters. See [Reverting from Version 15.0.0 on page 65](#).

## Preparing to revert

If you used Version 15.0.0, you must prepare your system for reversion to the pre-migration version of the database server.

Review and complete all tasks that apply:

1. [Review the database schema prior to reversion on page 61.](#)
2. [Check and configure available space for reversion on page 62.](#)
3. [Save copies of the current configuration files on page 62.](#)
4. [Save system catalog information on page 62.](#)
5. [Verify the integrity of the Version 15.0.0 data on page 62.](#)
6. [Back up Version 15.0.0 on page 63.](#)
7. Export or save your data.
8. To prevent reversion issues:
  - a. [Resolve outstanding in-place alter operations on page 14.](#)
  - b. If you have empty tables with no extents, drop those tables.
  - c. Defragment partitions by using the SQL administration API.
9. [Remove unsupported features on page 65.](#)
10. [Remove new extensions on page 65.](#)

## Results

### What to do next

If you use high-availability clusters or Enterprise Replication, you must complete extra tasks before you can revert to the pre-migration version of Informix®.

## Review the database schema prior to reversion

You must review your database schema to determine if reversion is possible. You can revert from Version 15.0.0 to the database server from which you migrated, if you have not added any extensions to the Version 15.0.0 database server and you are not reverting from a newly created instance.

To review the database schema to determine if reversion is possible:

1. Run the dbschema utility command.

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

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

2. Determine if the schema file contains SQL statements that the earlier database server does not support.
3. Determine if any new SPL routines have been created in Informix® Version 15.0.0 or if any routines were imported using dbimport.
4. Determine if tables or indexes using expression fragmentation had expressions changed or new fragments added.
5. Identify any new triggers, procedures, or check constraints.

See #unique\_79 for limitations on reversion to previous databases and prerequisite steps you must take before you revert.

## Check and configure available space for reversion

You must be sure you have enough space for reversion to the source database server.

The tblspace **tblspace** pages can be allocated in non-root chunks. If the root chunk is full and tblspace **tblspace** pages were allocated in non-root chunks, make sure you have enough space in the root chunk of the target database server.

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

For information about space requirements for Informix® Version 15.0.0, see [Checking and configuring available space on page 9](#).

## Save copies of the current configuration files

Save copies of the ONCONFIG and **concsn.cfg** files for when you migrate to Informix® Version 15.0.0 again.

Informix® uses the **concsn.cfg** file to configure CSMs.

## Save system catalog information

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

Run the following command to unload the system catalog tables:

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

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

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

## Verify the integrity of the Version 15.0.0 data

Verify the integrity of your Version 15.0.0 data, if you did not do this after you migrated.

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

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

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

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

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

## Back up Version 15.0.0

Before you begin the reversion, make a complete level-0 backup of Informix® Version 15.0.0.

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

## Resolve outstanding in-place alter operations

Resolving outstanding in-place alter operations is a requirement before *converting* to Version 15.0.0. If it becomes necessary to *revert* to a previous version you must also resolve any new outstanding in-place alter operations first.

An in-place alter operation is *outstanding* when data pages still exist with the prior definition, a state that can be detected with the oncheck -pT command. An in-place alter operation is *new* if the ALTER statement is executed in the higher database version.

If the reversion process detects new outstanding in-place alter operations, reversion fails and the message log will contain a list of all tables whose in-place alter operations must be resolved before the reversion will succeed.

If you are reverting from version 12.10.xC4 or later, you can remove in-place alter operations by running the admin( ) or task( ) SQL administration command with the table update\_ipa or fragment update\_ipa argument. You can include the parallel option to run the operation in parallel. For example, the following statement removes in-place alter operations in parallel from a table that is named **auto**:

```
EXECUTE FUNCTION task('table update_ipa parallel','auto');
```

If you are reverting from an earlier version of 12.10, you can resolve outstanding in-place alter operations by running sample UPDATE statements. Sample UPDATE statements force any outstanding in-place alter operations to complete by updating the rows in the affected tables. To generate a sample UPDATE statement, create an UPDATE statement in which a column in the table is set to its own value. This forces the row to be updated to the latest schema without changing column values. Because the database server always alters rows to the latest schema, a single pass through the table that updates all rows completes all outstanding in-place alter operations.

The sample UPDATE statement differs from a standard UPDATE statement because it does not change the data. A standard UPDATE statement usually changes the value of the affected row.

For example, to create a sample update, specify:

```
UPDATE tab1 SET col1=col1 WHERE 1=1 ;
```

You must ensure that the column selected is a numeric data type (for example, INTEGER or SMALLINT) and not a character data type.

If a table is large, a single update of the whole table can cause a long transaction. To avoid a long transaction, update the table in pieces, by ranges of some column, with this statement:

```
... WHERE {id_column} BETWEEN {low_value} AND {step_value}
```

For example, specify:

```
UPDATE tab1 SET col1=col1 WHERE col1 BETWEEN 1 AND 100;
UPDATE tab1 SET col1=col1 WHERE col1 BETWEEN 101 AND 200;
```

Ensure that the UPDATE statements include the entire data set.

If the table is replicated with Enterprise Replication, the database server replicates all updated rows unnecessarily. To avoid replication, update the table as follows:

```
BEGIN WORK WITHOUT REPLICATION;
...
COMMIT WORK;
```

When all the pending in-place alter operations are resolved, run the oncheck -pT command again for each table. In the output of the command, check information in the `version` section. The number of data pages should match with current version. Also, all other table versions should have `count=0` for the number of data pages that the version is accessing.

For example, if you run the oncheck -pT testdb:tab1 command after outstanding in-place alter operations are resolved, you might see information similar to the information in this segment of sample output:

TBLspace Report for testdb:root.tab1

```
Physical Address      1:860
Creation date         06/23/2011 14:23:08
TBLspace Flags        800801      Page Locking
                        TBLspace use 4 bit bit-maps

Maximum row size      29
Number of special columns 0
Number of keys        0
Number of extents     1
Current serial value   1
Current SERIAL8 value  1
Current BIGSERIAL value 1
Current REFID value    1
Pagesize (k)          2
First extent size     8
Next extent size      8
Number of pages allocated 8
Number of pages used   4
Number of data pages   3
  << Number of data pages used is 3 >>
Number of rows        6
Partition partnum     1048981
Partition lockid      1048981
```

Extents

Logical Page	Physical Page	Size	Physical Pages
0	1:1895	8	8

TBLspace Usage Report for testdb:root.tab1

Type	Pages	Empty	Semi-Full	Full	Very-Full
Free	4				
Bit-Map	1				
Index	0				
Data (Home)	3				



```

Total Pages                8

Unused Space Summary

    Unused data slots                177

Home Data Page Version Summary

    Version                        Count

    3 (oldest)                     0
    << Other version should show data page count=0>>
    4                              0
    << Other version should show data page count=0>>
    5 (current)                     3
    << Current should always match the number of data pages>>

```

## Remove unsupported features

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



**Important:** Do not remove objects that you did not create, such as the boot scripts (`boot90.sql` and `boot901.sql`) in the system catalog because the reversion utility uses them.

For a list of features that you must remove before reversion, see [Review the database schema prior to reversion on page 61](#).

## Remove new extensions

When or SQL Registration are used to register an extension in a database, the `ifxmngr` module, which manages extensions, is registered first. If you need to revert from Version 15.0.0, and you ran against a database, you must remove all extensions.

To remove the extensions, you must use to unregister all modules and then run the following BladeManager command:

```
unprep database_name
```

## Reverting from Version 15.0.0

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

### Before you begin

Prerequisite: Complete the steps in [Preparing to revert on page 60](#). Preparation includes determining if reversion is possible and preparing your database for reversion.

### About this task

Review and complete all tasks that apply:

1. [Run the reversion utility on page 66](#).
2. [Restore original configuration parameters on page 66](#).

3. [Restore original environment variables on page 67.](#)
4. [Remove any Communications Support Module settings on page 67.](#)
5. [Recompile Java user-defined routines on page 67.](#)
6. [Reinstall and start the earlier database server on page 67.](#)
7. [Optionally update statistics on your tables after reverting on page 68.](#)
8. [Verify the integrity of the reverted data on page 69.](#)
9. [Back up the database server after reversion on page 69.](#)
10. [Return the database server to online mode on page 69.](#)
11. If you use high-availability clusters, perform additional tasks that are described in [Reverting clusters on page 42.](#)

## Results



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

## Run the reversion utility

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



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

Informix® Version 15.0.0 must be running when you run the reversion utility. If the reversion utility detects and lists any remaining features that are specific to Informix® Version 15.0.0, you must remove those features before reversion can complete.

To see a list of all of the versions to which you can revert using an `onmode -b` command, type `onmode -b`.

To run the reversion utility, type

```
onmode -b version_number
```

For examples of the `onmode -b` command, see [Syntax of the onmode -b command on page 134.](#)

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

When the reversion is complete, Informix® is offline. The reversion utility drops the Informix® Version 15.0.0 system catalog tables and restores compatibility so that you can access the data with the earlier version of the database server. The reversion utility does not revert changes made to the layout of the data that do not affect compatibility.

## Restore original configuration parameters

Replace the Informix® Version 15.0.0 ONCONFIG configuration file with the ONCONFIG file that you saved before you migrated. Alternatively, you can remove configuration parameters that the earlier database server does not support.

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

## Restore original environment variables

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

## Remove any Communications Support Module settings

If your Informix® Version 15.0.0 instance used CSMs, edit the **sqlhosts** file to remove any **csn** option settings that are not supported in the older database server.

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

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



**Note:** Support for Communication Support Module (CSM) is removed starting Informix Server 14.10.xC9 . You should use Transport Layer Security (TLS)/Secure Sockets Layer (SSL) instead.

## Recompile Java™ user-defined routines

After you revert and before you start the earlier server, recompile Java™ user-defined routines (UDRs) that were compiled with a Java™ development kit version that is earlier than or equal to the version included with the previous server.

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

- If your application uses external JAR and class files (for example, JAR and class files that are listed in JVPCLASSPATH), recompile the files.
- If your application uses JAR files installed in the server (for example, through the `install_jar()` support function), you must remove the old JAR file (by using the `remove_jar()` support function) and reinstall the recompiled JAR file in the database.

## Reinstall and start the earlier database server

Reinstall and configure the earlier version of the database server.

Refer to the instructions in your *HCL® Informix® Installation Guide* and your *HCL® Informix® Administrator's Guide*.



**Before you start the server:** If you have time series data and you reverted to Informix® 12.10.xC1, you must install the patch `patch-idsdb00493404.tar` on Informix® 12.10.xC1. You can obtain the patch from the [IBM® Support Portal](#).

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



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

## Add JSON compatibility to databases that were created in 12.10.xC1

Databases that were created in Informix® 12.10.xC1 are not JSON compatible. In 12.10.xC2 or 12.10.xC3 you can run a script to make those databases JSON compatible. You do not have to run the script after you upgrade to 12.10.xC4 or later fixpacks because all databases are made JSON compatible during conversion.

### Before you begin

This procedure requires Informix® 12.10.xC2 or 12.10.xC3.

You must run the script as user **informix** or as a user with DBA privileges.

Run the appropriate script against a database that was created in Informix® 12.10.xC1 to make the database JSON compatible.

### Choose from:

- Informix® 12.10.xC3: `convTovNoSQL1210.sql`
- Informix® 12.10.xC2: `convTovNoSQL1210X2.sql`

### Example

For example, to make the **db\_name** database JSON compatible in Informix® 12.10.xC3, you would run the following command as user **informix** or as a user with DBA privileges:

#### UNIX™

```
dbaccess -e db_name $INFORMIXDIR/etc/convTovNoSQL1210.sql
```

#### Windows™

```
dbaccess -e db_name %INFORMIXDIR%\etc\convTovNoSQL1210.sql
```

## Optionally update statistics on your tables after reverting

Optionally run UPDATE STATISTICS on your tables and on UDRs that perform queries, if you have performance problems after reverting to the previous version of the database server or to a database server on a different operating system.

An unqualified UPDATE STATISTICS statement includes no additional clauses:

```
UPDATE STATISTICS;
```

By default, an unqualified UPDATE STATISTICS statement updates the statistics in LOW mode for every permanent table in the database, including the system catalog tables.

You can run an UPDATE STATISTICS statement that includes only a FOR ROUTINE clause that specifies no routine name:

```
UPDATE STATISTICS FOR ROUTINE;
```

Running this statement reoptimizes the DML statement execution plans for every SPL routine in the database that operates on local tables.

Similarly, you can substitute the keyword `FUNCTION` for `ROUTINE` in the previous example to reoptimize execution plans only for SPL routines that return at least one value, or you can substitute the keyword `FUNCTION` for `PROCEDURE` to reoptimize execution plans only for SPL routines that return no value. In these cases, the database server does not update the statistics in the system catalog tables.

You do not need to run `UPDATE STATISTICS` statements on C or Java™ UDRs.

## Verify the integrity of the reverted data

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

Follow the steps in [Verifying the integrity of the data on page 17](#).

## Back up the database server after reversion

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

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



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

## Return the database server to online mode

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

Then users can access the data.

## Register modules

You must register any modules that your databases require. Built-in database extensions automatically revert to the version included in the database server to which you reverted.

If you are reverting to version 11.70.xC1 or later, you do not need to register built-in database extensions because they are registered automatically.

For information on registering modules, see the *HCL® Informix® DataBlade® Module Installation and Registration Guide*.

# Migration of data between database servers

## Migrating database servers to a new operating system

When you migrate to a new operating system, you must choose a tool for migrating your data, you might need to make some adjustments to your tables, and you must review environment-dependent configuration parameters and environment variables.

## Choosing a tool for moving data before migrating between operating systems

If you are migrating between different operating systems, you must choose a method for exporting and importing data. The tool that you choose for exporting and importing data depends on how much data you plan to move.

All these methods deliver similar performance and enable you to modify the schema of the database. The tools that you can use include:

- The **dbexport** and **dbimport** utilities, which you can use to move an entire database
- The UNLOAD and LOAD statements, which move selected columns or tables (The LOAD statement does not change the data format.)
- The **dbload** utility, which you can use to change the data format
- The **onunload** utility, which unloads data in page-sized chunks, and the **onload** utility, which moves data to an identical database server on a computer of the same type
- Enterprise Replication, which you can use to transfer data between Informix® on one operating system and Informix® on a second operating system.

For an overview of all of these data-migration tools, a comparison of tools, and information about which versions of the database server do not support all of the tools, see [Data-migration tools on page 3](#).

## Adjusting database tables for file-system variations

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

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

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

## Moving data to a database server on a different operating system

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

### About this task

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

1. Save a copy of the current configuration files.
2. Use or **ontape** to make a final level-0 backup.  
For more information, refer to your *HCL® Informix® Backup and Restore Guide*.
3. Choose one of the following sets of migration utilities to unload the databases:

- **dbexport** and **dbimport**
  - UNLOAD, **dbschema**, and LOAD
  - UNLOAD, **dbschema**, and **dbload**
4. Bring the source database server offline.
  5. Install and configure the target database server. If you are migrating to Windows™, also install the administration tools.
  6. Bring the target database server online.
  7. Use **dbimport**, LOAD, or **dbload**, or external tables to load the databases into the target database server, depending on which utility you used to export the databases.
  8. Make an initial level-0 backup of the target database server.
  9. Run UPDATE STATISTICS to update the information that the target database server uses to plan efficient queries.

## Results

### Adapting your programs for a different operating system

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

Certain database server configuration parameters and environment variables are environment-dependent.

For details, see the information about configuration parameters in the *HCL® Informix® Administrator's Guide* and the *HCL® Informix® Administrator's Reference* and the information about environment variables in the *HCL® Informix® Administrator's Guide* and the *HCL® Informix® Guide to SQL: Reference*.

### Ensuring the successful creation of system databases

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

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

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

## Data migration utilities

### External tables

External tables are a fast and versatile method for moving data between database servers. External tables are the fastest method for loading data into a RAW table with no indexes.

You run the CREATE EXTERNAL TABLE statement to define an external table that is not part of your database to unload data from your database. You run INSERT ... SELECT statements to load data from the external table into your database.

You can unload and load data in the internal Informix® representation. All Informix® data types are supported. You can define a value to be interpreted as a NULL when you load or unload data from an external table. You can specify the date and currency format and replace missing values with column defaults.

You define a named pipe to copy data from one Informix® instance to another without writing the data to an intermediate file. You can monitor the I/O queues to determine whether you have enough FIFO virtual processors. If necessary, you can add more FIFO virtual processors to improve performance. You can specify to run high-speed transfers in parallel.

Rows that have conversion errors during a load are written to a reject file on the server that performs the conversion.

## The dbexport and dbimport utilities

The dbexport and dbimport utilities import and export a database and its schema to disk or tape.

The dbexport utility unloads an entire database into text files and creates a schema file. You can unload the database and its schema file either to disk or tape. If you prefer, you can unload the schema file to disk and unload the data to tape. You can use the schema file with the dbimport utility to re-create the database schema in another HCL Informix® environment, and you can edit the schema file to modify the database that dbimport creates.

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



### Attention:

When you import a database, use the same environment variable settings and configuration parameter settings that were used when the database was created.

- If any environment variables or configuration parameters that affect fragmentation, constraints, triggers, or user-defined routines are set differently than they were when these database objects were created originally, the database that is created by the dbimport utility might not be an accurate reproduction of the original.
- Incompatible settings are likely to occur if you move data from an earlier version of the database server to a newer version. Over time, some configuration parameters or environment variables are deprecated, or their default values are changed. For example, assume that attached indexes were created by default in the original database. In the current version of the database server, detached indexes are created by default. If you want to maintain the original behavior, you can set the **DEFAULT\_ATTACH** environment variable to 1 before you run the dbimport utility.

Also, the dbimport operation might fail when you attempt to import from a higher server version to a lower server version if the database schema changed between versions. For example, the **am\_expr\_pushdown** column was added to the **sysams** system catalog table in Informix® 11.70. The dbimport operation will fail if you attempt to import a database from Informix® 12.10 that contains the **am\_expr\_pushdown** column into a database from Informix® 11.50 that is missing that column. In that case, you must review the messages in the `dbimport.out` file, which is in your





current directory. After you address the issues that caused the dbimport operation to fail, run the dbimport command again.

Requirements or limitations apply in the following cases:

#### Compressed data

The dbexport utility uncompresses compressed data. You must recompress the data after you use the dbimport utility to import the data.

#### Date values

Use four-digit years for date values. The date context for an object includes the date that the object was created, the values of the **DBCENTURY** and **GL\_DATE** environment variables, and some other environment variables. If the date context during import is not the same as when these objects were created, you might get explicit errors, you might not be able to find your data, or a check constraint might not work as expected. Some of these problems do not generate errors.



**Tip:** By default, the dbexport utility exports dates in four-digit years unless environment variables are set to values that would override that format.

#### High-availability clusters

You cannot use the dbexport utility on HDR secondary servers or shared disk (SD) secondary servers.

The dbexport utility is supported on a remote standalone (RS) secondary server only if the server is set to stop applying log files. Use the STOP\_APPLY configuration parameter to stop application of log files.

The dbimport utility is supported on all updatable secondary servers.

#### Label-based access control (LBAC)

When you export data that is LBAC-protected, the data that is exported is limited to the data that your LBAC credentials allow you to read. If your LBAC credentials do not allow you to read a row, that row is not exported, but no error is returned. To export all the rows, you must be able to see all the rows.

#### NLSCASE mode

Whether the NLSCASE mode of your source database is SENSITIVE or INSENSITIVE, you can reduce the risk of case-sensitivity issues by always migrating to a target database that has the same NLSCASE mode as the source database. For tables that include columns with case-variant NCHAR and NVARCHAR data values (for example, 'IBM', 'ibm', 'lbm'), you might encounter the following differences after migration:

- ORDER BY and sorting operations can produce a different ordering of qualifying rows in query results, compared to the result of the same query before migration.
- Unique indexes and referential constraints with which the data were compliant before the migration might have integrity violations in the new database, if any index or constraint key column contains case-variant forms of the same character string.
- Queries with predicates that apply conditional operators to NCHAR or NVARCHAR values might return different results from the same data after migration.

### Nondefault database locales

If the database uses a nondefault locale and the **GL\_DATETIME** environment variable has a nondefault setting, you must set the **USE\_DTEENV** environment variable to the value of **1** so that localized DATETIME values are processed correctly by the dbexport and dbimport utilities.

### SELECT triggers on tables

You must disable SELECT triggers before you export a database with the dbexport utility. The dbexport utility runs SELECT statements during export. The SELECT statement triggers might modify the database content.

### Virtual tables for the Informix® MQ extension


The MQCreateVtiRead(), MQCreateVtiReceive(), and MCQCreateVtiWrite() functions create virtual tables, and map them to the appropriate IBM® WebSphere® MQ message queue. When the dbexport utility unloads data, it removes the messages from WebSphere® MQ queues. Before you use the dbexport utility, drop any MQ virtual tables. After you load the database with the dbimport utility, you can create the tables in the target database by using the appropriate functions.

## Syntax of the dbexport command

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

```
dbexport {[{ | -c | -d | -no-data-tables tablename | -no-data-tables-accessmethods accessmethods | -nw | -q |
<Destination Options> (explicit id) | -ss | -si | -x } ] database | [{ -v | -version } ] }
```

Element	Purpose	Key Considerations
<b>-c</b>	Makes dbexport complete exporting unless a fatal error occurs	<b>References:</b> For details on this option, see <a href="#">dbexport errors on page 77</a> .
<b>-d</b>	Makes dbexport export simple-large-object descriptors only, not simple-large-object data	
<b>-q</b>	Suppresses the display of error messages, warnings, and generated SQL data-definition statements	None.

Element	Purpose	Key Considerations
<b>-ss</b>	Generates database server-specific information for all tables in the specified database	<b>References:</b> For details on this option, see <a href="#">dbexport server-specific information on page 77</a> .
<b>-si</b>	Excludes the generation of index storage clauses for non-fragmented tables  The <b>-si</b> option is available only with the <b>-ss</b> option.	<b>References:</b> For details on this option, see <a href="#">dbexport server-specific information on page 77</a> .
<b>-X</b>	Recognizes HEX binary data in character fields	None.
<b>-no-data-tables</b>	Prevents data from being exported for the specified tables. Only the definitions of the specified tables are exported.	Accepts a comma-separated list of names of tables for which data will not be exported.   <b>Default behavior:</b> Only the definition of the <b>tsinstanceTable</b> table is exported, not the data. The data and definitions of all other tables are exported.
<b>-no-data-tables-access-methods</b>	Prevents data from being unloaded using the specified access methods.	Accepts a comma-separated list of names of access methods. Tables using those access methods are not unloaded.  <b>Default value:</b>  ts_rts_vtam, ts_vtam  Tables using ts_rts_vtam and ts_vtam access methods are not unloaded.
<b>-nw</b>	Generates the SQL for creating a database without the specification of an owner	None.
<b>-V</b>	Displays the software version number and the serial number	None.
<b>-version</b>	Extends the <b>-V</b> option to display additional information about the build operating system, build number, and build date	None.
<i>database</i>	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.

Element	Purpose	Key Considerations
		<b>References:</b> If you want to use more than the simple name of the database, see <a href="#">Database Name</a> on page .

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 *HCL® Informix® GLS User's Guide*, dbexport can handle foreign characters in data and export the data from databases. For more information, refer to [Database renaming on page 86](#).

You can set the IFX\_UNLOAD\_EILSEQ\_MODE environment variable to enable dbexport to use character data that is invalid for the locale specified in the environment.

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

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

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



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

## Example

### Example

The following command exports the table definitions but no data for all the tables in the customer database.

```
dbexport -no-data-tables -no-data-tables-accessmethods customer
```

## Example

### Example

The following command generates the schema and data for the `customer` database without the specification of an owner:

```
dbexport customer -nw
```

## Termination of the dbexport utility

You can stop the **dbexport** utility at any time.

To cancel **dbexport**, press your **Interrupt** key.

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

## dbexport errors

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

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

- **dbexport** is unable to open the specified tape.
- **dbexport** finds bad writes to the tape or disk.
- Invalid command parameters were used.
- **dbexport** cannot open the database or there is no system permission for doing so.
- A subdirectory with the name specified during invocation already exists

## dbexport server-specific information

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

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

## dbexport destination options

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

### Destination options

```
" [ { -o directory | -t device -b blocksize -s tapesize [ -f pathname ] } ] "
```

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 name of the path where you want the schema file stored, if you are storing the data files on tape.	<p>The path name can be a complete path name or a file name. If only a file name is given, the file is stored in the current directory.</p> <p>If you do not use the <b>-f</b> option, the SQL source code is written to the tape.</p>

Element	Purpose	Key Considerations
<code>-o directory</code>	Specifies the directory on disk in which dbexport creates the <i>database.exp</i> directory.  This directory holds the data files and the schema file that dbexport creates for the database.	The specified directory must exist.
<code>-s tapesize</code>	Specifies, in kilobytes, the amount of data that you can store on the tape.	To write to the end of the tape, set the value to 0.  If you do not specify 0, the maximum size is 2 097 151 KB.
<code>-t device</code>	Specifies the path name of the tape device where you want the text files and, possibly, the schema file stored.	You cannot 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 database servers on UNIX™ or , the command is:

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

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

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

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

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

### Windows™ Only

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

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

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

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

## Exporting time series data in rolling window containers

The `dbexport` utility exports time series data except any data that is in the dormant window of rolling window containers.

### About this task

The active window in rolling window containers is re-created after the time series data is loaded into a container.

The dormant window is not exported. To export the data from the dormant window, you must move the data into the active window.

To export time series data in the dormant window of a rolling window container:

1. If necessary, increase the size of the active window by running the `TSContainerManage` function. The size of the active window must be large enough to fit all the intervals in the dormant window can fit into the active window.
2. Move the intervals in the dormant window into the active window by running the `TSContainerManage` function.
3. Export the data by running the `dbexport` utility.

## Contents of the schema file that `dbexport` creates

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

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

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 *HCL® Informix® Guide to SQL: Reference*.

- The starting values of `SERIAL` columns
- The dormant window interval values for time series rolling window containers

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

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

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



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

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

## Syntax of the dbimport command

The dbimport command imports previously exported data into another database.

```
dbimport [{ | -c | -D | -nv | -q | -x } ] <Input-File Location> (explicit id) <Create Options> (explicit id) [{ -v | -version } ]
database
```

Element	Purpose	Key Considerations
-c	Completes importing data even when certain nonfatal errors occur	For more information, see <a href="#">dbimport errors and warnings on page 82</a> .
-D	Specifies a default extent size of 16 KB for the first and subsequent extents during the import operation, if the extent sizes are not specified in the CREATE TABLE statement.	<p>This option is ignored if the extent sizes are specified in the CREATE TABLE statement.</p> <p>Default values help to ensure that enough space is available in the dbspace that is designated for the import operation. This option prevents the automatic calculation of extent sizes during the import operation, and is useful especially in the following situations:</p> <ul style="list-style-type: none"> <li>• When importing tables that contain columns with large maximum row sizes, such as LVARCHAR columns.</li> <li>• When importing data after the dbexport command was run without the -ss option. The -ss option specifies server-specific information about extent sizes.</li> </ul>
-nv	While the dbimport -nv command is running, tables with foreign-key constraints that	By bypassing the checking of referential constraints, this option can reduce migration time for very large tables



Element	Purpose	Key Considerations
	ALTER TABLE ADD CONSTRAINT creates in enabled or filtering mode are not checked for violations, as if you had also specified NOVALIDATE	that already conform to their foreign-key constraints. The NOVALIDATE mode does not persist after the ALTER TABLE ADD CONSTRAINT statement has completed.
<b>-q</b>	Suppresses the display of error messages, warnings, and generated SQL data-definition statements	None.
<b>-V</b>	Displays the software version number and the serial number	None.
<b>-version</b>	Extends the <b>-V</b> option to display additional information about the build operating system, build number, and build date	None.
<b>-X</b>	Recognizes HEX binary data in character fields	None.
<i>database</i>	Declares the name of the database to create	If you want to use more than the simple name of the database, see <a href="#">Database Name</a> on page .

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

- All input files are on disk.
- All input files are on tape.
- The schema file is 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 an imported Informix® database server:

- Specify the dbspace where the database will reside
- Create an ANSI-compliant database with unbuffered logging
- Create a database that supports explicit transactions (with buffered or unbuffered logging)
- Create an unlogged database
- Create a database with the NLS case-insensitive property for NCHAR and NVARCHAR strings.
- Process all ALTER TABLE ADD CONSTRAINT and SET CONSTRAINTS statements in the .sql file of the exported database that define enabled or filtering referential constraints so that any foreign-key constraints that are not specified as DISABLED are in ENABLED NOVALIDATE or in FILTERING NOVALIDATE mode.



**Note:** If you specify the -nv option, the .sql file of the exported database is not modified, but any foreign-key constraints that ALTER TABLE ADD CONSTRAINT or SET CONSTRAINTS statements enable are processed without



checking each row of the table for violations. The `ENABLED`, or `FILTERING WITH ERROR`, or `FILTERING WITHOUT ERROR` constraint mode specifications are implemented instead as the `ENABLED NOVALIDATE`, or `FILTERING WITH ERROR NOVALIDATE` or `FILTERING WITHOUT ERROR NOVALIDATE` modes. After the foreign-key constraints have been enabled without checking for violations, their modes automatically revert to whatever the `.sql` file specified so that subsequent DML operations on the tables enforce referential integrity.

The user who runs the `dbimport` utility 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 *HCL® Informix® GLS User's Guide* describes, `dbimport` can import data into database server versions that support .

## Termination of the `dbimport` utility

You can stop the **`dbimport`** utility at any time.

To cancel the `dbimport` utility, press your `Interrupt` key .

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

## `dbimport` errors and warnings

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

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

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

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

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

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

## dbimport input-file location options

The input-file location specifies the location of the `database.exp` directory, which contains the files that the dbimport utility imports.

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

### dbimport input-file location

```
"[{ -i directory | -t device [ -b blocksize -s tapesize ] [ -f pathname ] }]"
```

Element	Purpose	Key Considerations
<code>-b blocksize</code>	Specifies, in kilobytes, the block size of the tape device	<p>If you are importing from tape, you must use the same block size that you used to export the database.</p> <p>If you do not use the <code>-b</code> option, the default block size is 1.</p>
<code>-f pathname</code>	Specifies where dbimport can find the schema file to use as input to create the database when the data files are read from tape	<p>If you use the <code>-f</code> option to export a database, you typically use the same path name that you specified in the <code>dbexport</code> command. If you specify only a file name, dbimport looks for the file in the <code>.exp</code> subdirectory of your current directory.</p> <p>If you do not use the <code>-f</code> option, the SQL source code is written to the tape.</p>
<code>-i directory</code>	Specifies the complete path name on disk of the <code>database.exp</code> directory, which holds the input data files and schema file that dbimport uses to create and load the new database. The directory name must be the same as the database name.	This directory must be the same directory that you specified with the <code>dbexport -o</code> option. If you change the directory name, you also rename your database.
<code>-s tapesize</code>	Specifies, in kilobytes, the amount of data that you can store on the tape	<p>To read to the end of the tape, specify a tape size of 0.</p> <p>If you are importing from tape, you must use the same tape size that you used to export the database. The maximum size is 2 097 151 KB.</p> <p>If you do not use the <code>-s</code> option, the default value is 0 (read to the end of the tape).</p>
<code>-t device</code>	Specifies the path name of the tape device that holds the input files	You cannot specify a remote tape device.

### Examples showing input file location on UNIX™ or Linux™

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

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

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

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

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

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

### Examples showing input file location on Windows™

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

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

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

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

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

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

### dbimport create options

The **dbimport** utility supports options for creating a database, specifying a dbspace for that database, defining logging options, and optionally specifying ANSI/ISO-compliance or NLS case-insensitivity (or both) as properties of the database.

#### Create options

```
" [-adbspace] [-1 { [buffered] | [-ansi] }] [-ci] "
```

Element	Purpose	Key Considerations
-ansi	Creates an ANSI/ISO-compliant database in which the ANSI/ISO rules for transaction logging are enabled. Otherwise, the	If you omit the -ansi option, the database uses explicit transactions.

Element	Purpose	Key Considerations
	database uses explicit transactions by default.	<b>Additional Information:</b> For more information about ANSI/ISO-compliant databases, see the <i>HCL® Informix® Guide to SQL: Reference</i> .
<b>-ci</b>	Specifies the NLS case-insensitive property. Otherwise, the database is case-sensitive by default.	<b>Additional Information:</b> See the <i>HCL® Informix® Guide to SQL: Syntax</i> and <i>HCL® Informix® Guide to SQL: Reference</i> descriptions of the NLS case-insensitive property.
<b>-d dbspace</b>	Specifies the dbspace where the database is created. .	If this is omitted, the default location is the <b>root</b> dbspace
<b>-l</b>	Establishes unbuffered transaction logging for the imported database. If the <b>-l</b> flag is omitted, the database is unlogged,	<b>References:</b> For more information, see <a href="#">Database-logging mode on page 86</a> .
<b>-l buffered</b>	Establishes buffered transaction logging for the imported database. If <b>-l</b> is included but <b>buffered</b> is omitted, the database uses unbuffered logging.	<b>References:</b> For more information, see <a href="#">Database-logging mode on page 86</a> .

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

### Example showing dbimport create options (UNIX™ or Linux™)

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

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

The new database is ANSI/ISO-compliant.

The next example similarly imports the **stores\_demo** database from the **/usr/informix/port/stores\_demo.exp** directory. The imported database uses buffered transaction logging and explicit transactions. The **-ci** flag specifies *case insensitivity* in queries and in other operations on columns and character strings of the NCHAR and NVARCHAR data types:

```
dbimport -c stores_demo -i /usr/informix/port -l buffered -ci
```

The **-ansi** and **-ci** options for database properties are not mutually exclusive. You can specify an ANSI/ISO-compliant database that is also NLS case-insensitive, as in the following example of the **dbimport** command:

```
dbimport -c stores_demo -i /usr/informix/port -l -ansi -ci
```

### Example showing dbimport create options (Windows™)

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

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

The imported database is ANSI/ISO-compliant and is case-sensitive for all built-in character data types.

## Database-logging mode

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

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

- ANSI-compliant database with unbuffered logging
- Unbuffered logging
- Buffered logging
- No logging

For more information, see [dbimport create options on page 84](#).

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

- Omitting any of the **-l** options is equivalent to omitting the WITH LOG clause.
- The **-l** option is equivalent to the WITH LOG clause.
- The **-l buffered** option is equivalent to the WITH BUFFERED LOG.
- The **-l ansi** option is equivalent to the WITH LOG MODE ANSI clause, and implies unbuffered logging.

## 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.

You can use the RENAME DATABASE statement to change the database name.

### Alternative ways to change the database name

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

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

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

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

To change the database name to a new name on Windows™:

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

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

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

## Changing the database locale with dbimport

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

To change the locale of a database:

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

## Simple large objects

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

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

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

### Windows™ Only

Windows™ sets the **TMP** and **TEMP** environment variables in the command prompt sessions, by default.

However, if the **TMP**, **TEMP**, and **DBTEMP** environment variables are not set, **dbimport** places the temporary files for the simple large objects in the **\tmp** directory.



**Attention:** If a table has a CLOB or BLOB in a column, you cannot use **dbexport** to export the table to a tape. If a table has a user-defined type in a column, using **dbexport** to export the table to a tape might yield unpredictable



results, depending on the export function of the user-defined type. Exported CLOB sizes are stored in hex format in the unload file.

## The dbload utility

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

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

You cannot use the **dbload** utility on secondary servers in high-availability clusters.

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

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

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

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

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

## Syntax of the dbload command

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

```
dbload [{ | -a database | -c command file | -l error log file } [{ | -r | -k | -e errors [-p] | -i ignore rows | -n commit interval | -x } ] [-S] [{ -v | -version } ] }
```

Element	Purpose	Key Considerations
-c command file	Specifies the file name or path name of a <b>dbload</b> command file	<b>References:</b> For information about building the command file, see <a href="#">Command file for the dbload utility on page 92</a> .



Element	Purpose	Key Considerations
<b>-d database</b>	Specifies the name of the database to receive the data	<b>Additional Information:</b> To use more than the simple name of the database, see <a href="#">Database Name on page</a> .
<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 <a href="#">Bad-row limit during a load operation on page 91</a> .
<b>-i ignore rows</b>	Specifies the number of rows to ignore in the input file	<b>References:</b> For more information, see <a href="#">Rows to ignore during a load operation on page 91</a> .
<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 <a href="#">Table locking during a load operation on page 90</a>.</p> <p>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 file name or path name of an error log file	<p>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>	Specifies the commit interval in number of rows  The default interval is 100 rows.	<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>HCL® 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 <a href="#">Bad-row limit during a load operation on page 91</a> .
<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 <a href="#">Table locking during a load operation on page 90</a>.</p> <p>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>

Element	Purpose	Key Considerations
<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.
<b>-V</b>	Displays the software version number and the serial number	None.
<b>-version</b>	Extends the <b>-V</b> option to display additional information about the build operating system, build number, and build date	None.
<b>-X</b>	Recognizes HEX binary data in character fields	None.



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

## Example

### dbload command example

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

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

## Table locking during a load operation

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

If you do not specify the **-k** option, the tables specified in the command file are locked in shared mode. When tables are locked in shared mode, the database server still must 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 need to acquire exclusive locks on rows or pages as it inserts rows during the load operation.

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

## Rows to ignore during a load operation

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

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

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

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

## Bad-row limit during a load operation

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

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

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

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

## Termination of the dbload utility

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

## Name and object guidelines for the dbload utility

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

**Table 7. Name and object guidelines for the dbload utility**

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

**Table 7. Name and object guidelines for the dbload utility (continued)**

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

## Command file for the dbload utility

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

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

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

- Delimiter form
- Character-position form

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

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

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

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

## Delimiter form of the FILE and INSERT statements

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

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

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

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

### Syntax for the delimiter form

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

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

**FILE***filename***DELIMITER ' c ' nfields**

Element	Purpose	Key Considerations
<i>c</i>	Specifies the character as the field delimiter for the specific input file	<p>If the delimiter specified by <i>c</i> appears as a literal character anywhere in the input file, the character must be preceded with a backslash (\) in the input file. For example, if the value of <i>c</i> is specified as a square bracket ([) , 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.</p> <p>You can specify any printable character, as defined by current locale, the tab character <b>TAB</b> (<b>CTRL-I</b>), or a blank space (ASCII 32) as the delimiter symbol. You cannot specify non-printable character, a hexadecimal character, or a backslash character.</p>
<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 95](#).

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

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

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

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

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

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

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

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

```
FILE "testtbl2.unl" DELIMITER '|' 1;
    INSERT INTO testtbl
        (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.

```
INSERT INTO [ owner. ] table [ ( column ) ] [ <VALUES clause> (explicit id) ] ;
```

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

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

## How to write a dbload command file in delimiter form

Command files must contain required elements, including delimiters.

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

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

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

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

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

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

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

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

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

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

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



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

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

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

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

## Character-position form of the FILE and INSERT statements

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

The examples in this topic are based on an input data file, **cust\_loc\_data**, which contains the last four columns (**city**, **state**, **zipcode**, and **phone**) of the **customer** table. Fields in the input file are padded with blanks to create data rows in which the location of data fields and the number of characters are the same across all rows. The definitions for these fields are CHAR(15), CHAR(2), CHAR(5), and CHAR(12), respectively. [Figure 2: A Sample Data File on page 96](#) displays the character positions and five example data rows from the **cust\_loc\_data** file.

Figure 2. A Sample Data File

	12	3
	1234567890123456789012345678901234	
Sunnyvale	CA94086408-789-8075	
Denver	CO80219303-936-7731	
Blue Island	NY60406312-944-5691	
Brighton	MA02135617-232-4159	
Tempe	AZ85253xxx-xxx-xxxx	

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



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

## Syntax for the character-position form

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

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

```
FILE filename( fieldn [ start [-end] ] [ NULL= 'null string' ] )
```

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

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

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

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

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

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

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

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

```
INSERT INTO [ owner. ] table [ ( column ) ] [ <VALUES clause> (explicit id) ] ;
```

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

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

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

## How to write a dbload command file in character-position form


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

The FILE statement in the following example defines six data fields from the **cust\_loc\_data** table data rows.

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

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

Figure 3. A Sample Data File



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

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

The null strings that are defined for the **phone** and **area\_cd** fields generate the null values in those columns, but they do not affect the values that are stored in the **state\_area** column.

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

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

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

Column	Values from Data Row 1	Values from Data Row 2
<b>col1</b>	Sunnyvale++++++	Tempe+++++++
<b>col2</b>	null	null
<b>col3</b>	CA	AZ
<b>col4</b>	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
<b>col1</b>	408	null
<b>col2</b>	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

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

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

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

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

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



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

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

You can use **dbload** with named row types, unnamed row types, sets, and lists.

## Using the dbload utility with named row types

The procedure for using the **dbload** utility with named row types is somewhat different than the procedure for using **dbload** with other complex data types, because named row types are actually user-defined data types.

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

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

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

**To load data for a named row type (or for any user-defined data type)**

1. Use the UNLOAD statement to unload the table to an input file. In this example, the input file sees the named row type as six separate fields:

```
Brown, James|13 First St.|San Francisco|CA|94070|01/04/1940|
Karen Smith|1820 Elm Ave #100|Fremont|CA|94502|01/13/1983|
```

2. Use the **dbschema** utility to capture the schema of the table and the row type. You must use the **dbschema -u** option to pick up the named row type.

```
dbschema -d stores_demo -u person_t > schema.sql
dbschema -d stores_demo -t person > schema.sql
```

3. Use to re-create the **person** table in the new database.

For detailed steps, see [Use dbschema output as DB-Access input on page 119](#).

4. Create the **dbload** command file. This **dbload** command file inserts two rows into the **person** table in the new database.

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

This **dbload** example shows how to insert new data rows into the **person** table. The number of rows in the INSERT statement and the **dbload** command file must match:

```
FILE person.unl DELIMITER '|' 6;
INSERT INTO person
VALUES ('Jones, Richard', '95 East Ave.',
       'Philadelphia', 'PA',
       '19115',
       '03/15/97');
```

5. Run the **dbload** command:

```
dbload -d newdb -c uds_command -l errlog
```



**Tip:** To find the number of fields in an unloaded table that contains a named row type, count the number of fields between each vertical bar (|) delimiter.

## Using the dbload utility with unnamed row types

You can use the **dbload** utility with unnamed row types, which are created with the ROW constructor and define the type of a column or field.

In the following example, the **devtest** table contains two columns with unnamed row types, **s\_name** and **s\_address**. The **s\_name** column contains three fields: **f\_name**, **m\_init**, and **l\_name**. The **s\_address** column contains four fields: **street**, **city**, **state**, and **zip**.

```
CREATE TABLE devtest
(
  s_name ROW(f_name varchar(20), m_init char(1), l_name varchar(20)
not null),
```

```
s_address ROW(street varchar(20), city varchar(20), state char(20),
zip varchar(9)
);
```

The data from the **devtest** table is unloaded into the **devtest.unl** file. Each data row contains two delimited fields, one for each unnamed row type. The ROW constructor precedes each unnamed row type, as follows:

```
ROW('Jim','K','Johnson')|ROW('10 Grove St.','Eldorado','CA','94108')|
ROW('Maria','E','Martinez')|ROW('2387 West Wilton
Ave.','Hershey','PA','17033')|
```

This **dbload** example shows how to insert data that contains unnamed row types into the **devtest** table. Put double quotes around each unnamed row type or the insert will not work.

```
FILE devtest.unl DELIMITER '|' 2;
INSERT INTO devtest (s_name, s_address)
VALUES ("row('Stephen', 'M', 'Wu')",
"row('1200 Grand Ave.', 'Richmond', 'OR', '97200')");
```

## Using the dbload utility with collection data types

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

### SET data type example

The SET data type is an unordered collection type that stores unique elements. The number of elements in a SET data type can vary, but no nulls are allowed.

The following statement creates a table in which the **children** column is defined as a SET:

```
CREATE TABLE employee
(
    name char(30),
    address char(40),
    children SET (varchar(30) NOT NULL)
);
```

The data from the **employee** table is unloaded into the **employee.unl** file. Each data row contains four delimited fields. The first set contains three elements (**Karen**, **Lauren**, and **Andrea**), whereas the second set contains four elements. The SET constructor precedes each SET data row.

```
Muriel|5555 SW Merry
Sailing Dr.|02/06/1926|SET{'Karen','Lauren','Andrea'}|
Larry|1234 Indian Lane|07/31/1927|SET{'Martha',
'Melissa','Craig','Larry'}|
```

This **dbload** example shows how to insert data that contains SET data types into the **employee** table in the new database. Put double quotes around each SET data type or the insert does not work.

```
FILE employee.unl DELIMITER '|' 4;
INSERT INTO employee
VALUES ('Marvin', '10734 Pardee', '06/17/27',
"SET{'Joe', 'Ann'}");
```

## LIST data type example

The LIST data type is a collection type that stores ordered, non-unique elements; that is, it allows duplicate element values.

The following statement creates a table in which the **month\_sales** column is defined as a LIST:

```
CREATE TABLE sales_person
(
    name CHAR(30),
    month_sales LIST(MONEY NOT NULL)
);
```

The data from the **sales\_person** table is unloaded into the **sales.unl** file. Each data row contains two delimited fields, as follows:

```
Jane Doe|LIST{'4.00','20.45','000.99'}|
Big Earner|LIST{'0000.00','00000.00','999.99'}|
```

This **dbload** example shows how to insert data that contains LIST data types into the **sales\_person** table in the new database. Put double quotes around each LIST data type or the insert does not work.

```
FILE sales_person.unl DELIMITER '|' 2;
INSERT INTO sales_person
VALUES ('Jenny Chow', "{587900, 600000}");
```

You can load multisets in a similar manner.

## The dbschema utility

The **dbschema** utility displays the SQL statements (the *schema*) that are necessary to replicate database objects.

You can also use the **dbschema** utility for the following purposes:

- To display the distributions that the UPDATE STATISTICS statement creates.
- To display the schema for the Information Schema views
- To display the schema for creating objects such as databases, tables, sequences, synonyms, storage spaces, chunks, logs, roles, and privileges
- To display the distribution information that is stored for one or more tables in the database
- To display information about user-defined data types and row types

After you obtain the schema of a database, you can redirect the **dbschema** output to a file that you can use with DB-Access.

The **dbschema** utility is supported on all updatable secondary servers.

The **dbschema** utility is also supported on read-only secondary servers. However, the **dbschema** utility displays a warning message when running on these servers.



**Attention:** Use of the **dbschema** utility can increment sequence objects in the database, creating gaps in the generated numbers that might not be expected in applications that require serialized integers.

## Object modes and violation detection in dbschema output

The output from the **dbschema** utility shows object modes and supports violation detection.

The **dbschema** output shows:

- The names of not-null constraints after the not-null specifications.

You can use the output of the utility as input to create another database. If the same names were not used for not-null constraints in both databases, problems could result.

- The object mode of objects that are in the disabled state. These objects can be constraints, triggers, or indexes.
- The object mode of objects that are in the filtering state. These objects can be constraints or unique indexes.
- The violations and diagnostics tables that are associated with a base table (if violations and diagnostics tables were started for the base table).

For more information about object modes and violation detection, see the SET, START VIOLATIONS TABLE, and STOP VIOLATIONS TABLE statements in the *HCL® Informix® Guide to SQL: Syntax*.

## Guidelines for using the dbschema utility

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



**Global Language Support:** You must disable SELECT triggers and correctly set GLS environment variables before using the **dbschema** utility.

When the GLS environment variables are set correctly, as the *HCL® Informix® GLS User's Guide* describes, the **dbschema** utility can handle foreign characters.

## Syntax of the dbschema command

The **dbschema** command displays the SQL statements (the *schema*) that are necessary to replicate a specified database object. The command also shows the distributions that the UPDATE STATISTICS statement creates.



```
dbschema { <Table options> <Database options> { <UDT options> } | [ { -v | -version } ] } <Storage, space, and log options>
<No owner option>
```

UDT options

```
" { (explicit id) | [ { -u all } ] | [ { -ua | -ui } udt_name ] } "
```

Table options

```
" { [ <Tables, Views, or Procedures> (explicit id) ] [ <Synonyms> (explicit id) ] [ <Privileges> (explicit id) ] | -hd { all |
<Table Name> (explicit id) } } "
" [ -r { role | all } (explicit id) ] "
```

Database options

```
" [ -ss ] "
" [ -seq { sequence | all } ] "
" -adatabase "
" [ -wpassword ] "
```

Storage space and log options

```
" [ -c [ -ns ] file_name (explicit id) ] "
```

No owner option

```
" [ -nw ] "
```

Element	Purpose	Additional Information
all	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.

Element	Purpose	Additional Information
<b>-c</b> <i>file_name</i>	Generates commands to reproduce storage spaces, chunks, physical logs, and logical logs.	<p>If you use the <b>-c</b> element without the <b>-ns</b> element, the database server generates SQL administration API commands.</p> <p>If you use the <b>-c</b> element and also use the <b>-ns</b> element, the database server generates <b>onspaces</b> or <b>onparams</b> commands.</p>
<b>-d</b> <i>database</i>	Specifies the database to which the schema applies. The <i>database</i> can be on a remote database server.	To use more than the simple name of the database, see <a href="#">Database Name on page</a> .
<i>filename</i>	Specifies the name of the file that contains the <b>dbschema</b> output.	If you omit a file name, <b>dbschema</b> sends the output to the screen. If you specify a file name, <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.	If you specify the ALL keyword for the table name, the distributions for all the tables in the database are displayed.
<b>-it</b>	<p>Sets the isolation type for <b>dbschema</b> while <b>dbschema</b> queries catalog tables. Isolation types are:</p> <p>DR = Dirty Read  CR = Committed Read  CS = Cursor Stability  CRU = Committed Read with RETAIN UPDATE LOCKS  CSU = Cursor Stability with RETAIN UPDATE LOCKS  DRU = Dirty Read with RETAIN UPDATE LOCKS  LC = Committed Read, Last Committed  RR = Repeatable Read</p>	
<b>-l</b>	Sets the lock mode to wait <i>number of</i> seconds for <b>dbschema</b> while <b>dbschema</b> queries catalog tables.	None.
<b>-ns</b>	Generates <b>onspaces</b> or <b>onparams</b> utility commands to reproduce storage spaces, chunks, physical logs, and logical logs.	The <b>-c</b> element must precede the <b>-ns</b> element in your command.

Element	Purpose	Additional Information
<b>-nw</b>	Generates the SQL for creating an object without the specification of an owner.	The <b>-nw</b> element is also a <b>dbexport</b> command option.
<b>-q</b>	Suppresses the database version from the header.	This optional element precedes other elements.
<b>-r</b>	Generates information about the creation of roles.	For details, see <a href="#">Role creation on page 115</a> .
<b>-seq</b> <i>sequence</i>	Generates the DDL statement to define the specified <i>sequence</i> object	None.
<b>-ss</b>	Generates server-specific information	This option is ignored if no table schema is generated.
<b>-si</b>	Excludes the generation of index storage clauses for non-fragmented tables	This option is available only with the <b>-ss</b> option.
<b>-sl</b> <i>length</i>	Specifies the maximum length, in bytes, of unformatted CREATE TABLE and ALTER TABLE statements.	None.
<b>-u all</b>	Prints the definitions of all user-defined data types, including all functions and casts defined over the types.	Specify <b>-u all</b> to include all the user-defined data types in the list of distributions.
<b>-ua</b> <i>udt_name</i>	Prints the definition of a user-defined data type, including functions and casts defined over an opaque or constructor type.	None.
<b>-ui</b> <i>udt_name</i>	Prints the definition of a user-defined data type, including type inheritance.	None.
<b>-V</b>	Displays the software version number and the serial number	None.
<b>-version</b>	Extends the <b>-V</b> option to display additional information about the build version, host, operating system, build number and date, and the GLS version.	None.
<b>-w</b> <i>password</i>	Specifies the database password, if you have one.	None.

You must be the DBA or have the Connect or Resource privilege for the database before you can run **dbschema** on it.

#### Example

## Example

The following command generates the schema with all the tables or sequence objects in the `customer` database, but without the specification of an owner:

```
dbschema -d customer all -nw
```

## Database schema creation

You can create the schema for an entire database or for a portion of the database.

Use the **dbschema** utility options 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.
- 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.

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 network

The **dbschema -d** option creates and displays the schema for databases on a UNIX™ or network.

You can specify a database on any accessible database server.

The following command displays the schema for the **stores\_demo** database on the **finland** database server on the UNIX™ or system console:

```
dbschema -d //finland/stores_demo
```

## Changing the owner of an object

You can edit **dbschema** output to change the owner of a new object.

The **dbschema** utility uses the *owner.object* convention when it generates any CREATE TABLE, CREATE INDEX, CREATE SYNONYM, CREATE VIEW, CREATE SEQUENCE, CREATE PROCEDURE, CREATE FUNCTION, or GRANT statement, and when it reproduces any unique, referential, or check constraint. As a result, if you use the **dbschema** output to create a new object (table, index, view, procedure, constraint, sequence, or synonym), the owner of the original object owns the new object. If you want to change the owner of the new object, you must edit the **dbschema** output before you run it as an SQL script.

You can use the output of **dbschema** to create a new function if you also specify the path name to a file in which compile-time warnings are stored. This path name is displayed in the **dbschema** output.

For more information about the CREATE TABLE, CREATE INDEX, CREATE SYNONYM, CREATE VIEW, CREATE SEQUENCE, CREATE PROCEDURE, CREATE FUNCTION, and GRANT statements, see the *HCL® Informix® Guide to SQL: Syntax*.

## dbschema server-specific information

The **dbschema -ss** option generates server-specific information. The **-ss** option always generates the lock mode, extent sizes, and the dbspace name if the dbspace name is different from the database dbspace. In addition, if tables are fragmented, the **-ss** option displays information about the fragmentation strategy.

When you specify the **dbschema -ss** option, the output also displays any GRANT FRAGMENT statements that are issued for a particular user or in the entire schema.

The **-si option**, which is available only with the **-ss** option, excludes the generation of index storage clauses for non-fragmented tables.

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 *HCL® Informix® Guide to SQL: Syntax*.

## User-defined and complex data types

The **dbschema -u** option displays the definitions of any user-defined and complex data types that the database contains. The suboption **i** adds the type inheritance to the information that the **dbschema -u** option displays.

The following command displays all the user-defined and complex data types for the **stork** database:

```
dbschema -d stork -u all
```

Output from **dbschema** that ran with the specified option `-u all` might appear as the following example shows:

```
create row type 'informix'.person_t
(
  name varchar(30, 10) not null,
  address varchar(20, 10),
  city varchar(20, 10),
  state char(2),
  zip integer,
  bdate date
);
create row type 'informix'.employee_t
(
```

```
salary integer,
manager varchar(30, 10)
) under person_t;
```

The following command displays the user-defined and complex data types, as well as their type inheritance for the **person\_t** table in the **stork** database:

```
dbschema -d stork -ui person_t
```

Output from **dbschema** that ran with the option `-ui person_t` might appear as the following example shows:

```
create row type 'informix'.person_t
(
  name varchar(30, 10) not null,
  address varchar(20, 10),
  city varchar(20, 10),
  state char(2),
  zip integer,
  bdate date
);
create row type 'informix'.employee_t
(
  salary integer,
  manager varchar(30, 10)
) under person_t;
create row type 'informix'.sales_rep_t
(
  rep_num integer,
  region_num integer,
  commission decimal(16),
  home_office boolean
) under employee_t;
```

## Sequence creation

The `dbschema -seq sequence` command generates information about sequence creation.

The following syntax diagram fragment shows sequence creation.

```
-seq { sequence | all }
```

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

Running **dbschema** with option **-seq** sequitur might produce this output:

```
CREATE SEQUENCE sequitur INCREMENT 10 START 100 NOCACHE CYCLE
```

For more information about the CREATE SEQUENCE statement, see the *HCL® Informix® Guide to SQL: Syntax*.

# Synonym creation

The **dbschema -s** command generates information about synonym creation.

The following syntax diagram fragment shows the creation of synonyms.

**Synonyms**

```

- s
" { ownername | all } "

```

Element	Purpose	Key Considerations
<b>-s ownername</b>	Displays the CREATE SYNONYM statements owned by <i>ownername</i>	None.
<b>-s all</b>	Displays all CREATE SYNONYM statements for the database, table, or view specified	None.

Output from **dbschema** that ran with the specified option **-s alice** might appear as the following example shows:

```
CREATE SYNONYM 'alice'.cust FOR 'alice'.customer
```

For more information about the CREATE SYNONYM statement, see the *HCL® Informix® Guide to SQL: Syntax*.

# Table, view, or procedure creation

Several **dbschema** options generate information that shows the creation of tables, views, and procedures.

The following syntax diagram shows the creation of tables, views, and procedures.

**Tables, Views, or Procedures:**

```

[ -t { table | view | all } ]
[ -t { table | view | { p | f } all | all } ]

```

Element	Purpose	Key Considerations
<b>-f all</b>	Limits the SQL statement output to those statements that replicate all functions and procedures	None.
<b>-f function</b>	Limits the SQL statement output to only those statements that replicate the specified function	None.

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

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

## Table information

The `dbschema -ss` command retrieves information about fragmented tables, the lock mode, and extent sizes.

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

```
{ TABLE "sallyc".t1 row size = 8 number of columns = 1 index size = 0 }
create table "sallyc".t1
(
  c1 integer
) fragment by expression
(c1 < 100 ) in db1 ,
((c1 >= 100 ) AND (c1 < 200 ) ) in db2 ,
remainder in db4
extent size 16 next size 16 lock mode page;
revoke all on "sallyc".t1 from "public";
```

The following **dbschema** output shows information about partitions in partition-fragmented tables.

DBSCHEMA Schema Utility grant dba to "sqlqa";

```
{ TABLE "sqlqa".t1 row size = 24 number of columns = 2 index size = 13 }
create table "sqlqa".t1
(
  c1 integer,
  c2 char(20)
)
fragment by expression
partition part_1 (c1 = 10 ) in dbs1 ,
partition part_2 (c1 = 20 ) in dbs1 ,
partition part_3 (c1 = 30 ) in dbs1 ,
partition part_4 (c1 = 40 ) in dbs1 ,
```



```
partition part_5 (c1 = 50 ) in dbs1
extent size 16 next size 16 lock mode page;
```

## Storage space, chunk, and log creation

The `dbschema -c` command generates SQL administration API commands for reproducing storage spaces, chunks, logical logs, and physical logs. If you use the `dbschema -c -ns` command, the database server generates **onspaces** or **onparams** utility commands for reproducing storage spaces, chunks, physical logs, and logical logs.

For example:

- Run the following command to generate a file named `dbschema1.out` that contains the commands for reproducing the storage spaces, chunks, physical logs, and logical logs in SQL Admin API format:

```
dbschema -c dbschema1.out
```

- Run the following command to generate a file named `dbschema2.out` that contains the commands for reproducing the storage spaces, chunks, physical logs, and logical logs in **onspaces** and **onparams** utility format:

```
dbschema -c -ns dbschema2.out
```

Optionally, specify **-q** before you specify **-c** or **-c -ns** to suppress the database version when you run the command. For example, specify:

```
dbschema -q -c -ns dbschema3.out
```

## Sample output for the creation of storage spaces, chunks, and logs

The output of the **dbschema -c** or **dbschema -c -ns** commands contain all of the SQL administration API or **onspaces** and **onparams** utility commands that you can use to reproduce storage spaces, chunks, and logs.

### Example of output in SQL administration API format

```
# Dbspace 1 -- Chunk 1
EXECUTE FUNCTION TASK ('create dbspace', 'rootdbs',
  '/export/home/informix/data/rootdbs1150fc4', '200000',
  '0', '2', '500', '100')

# Dbspace 2 -- Chunk 2
EXECUTE FUNCTION TASK ('create dbspace', 'datadbs1',
  '/export/home/informix/data/datadbs1150fc4', '5000000',
  '0', '2', '100', '100')

# Dbspace 3 -- Chunk 3
EXECUTE FUNCTION TASK ('create dbspace', 'datadbs2',
  '/export/home/informix/data/datadbs2150fc4', '5000000',
  '0', '2', '100', '100')

# Dbspace 4 -- Chunk 4
EXECUTE FUNCTION TASK ('create dbspace', 'datadbs3',
  '/export/home/informix/data/datadbs3_1150fc4', '80000',
  '16', '8', '400', '400')
EXECUTE FUNCTION TASK ('start mirror', 'datadbs3',
  '/export/home/informix/data/datadbs3_1150fc4', '80000',
  '16', '/export/home/informix/data/mdatadbs3_1150fc4', '16')
```

```

# Dbspace 5 -- Chunk 5
EXECUTE FUNCTION TASK ('create tempdbspace', 'tempdbs',
  '/export/home/informix/data/tempdbs_1150fc4', '1000',
  '0', '2', '100', '100')

# Dbspace 6 -- Chunk 6
EXECUTE FUNCTION TASK ('create sbospace', 'sospace',
  '/export/home/informix/data/sospace_1150fc4',
  '1000', '0')

# Dbspace 6 -- Chunk 7
EXECUTE FUNCTION TASK ('add chunk', 'sospace',
  '/export/home/informix/data/sospace_1_1150fc4',
  '1000', '0')

# Dbspace 7 -- Chunk 8
EXECUTE FUNCTION TASK ('create blobspace', 'blobdbs',
  '/export/home/informix/data/blobdbs_1150fc4',
  '1000', '0', '4')

# External Space 1
EXECUTE FUNCTION TASK ('create extspace', 'extspace',
  '/export/home/informix/data/extspac_1150fc4')

# Physical Log
EXECUTE FUNCTION TASK ('alter plog', 'rootdbs', '60000')

# Logical Log 1
EXECUTE FUNCTION TASK ('add log', 'rootdbs', '10000')

```

## Example of output in onspaces and onparams utility format

```

# Dbspace 1 -- Chunk 1
onspaces -c -d rootdbs -k 2 -p
/export/home/informix/data/rootdbs1150fc4
-o 0 -s 200000 -en 500 -ef 100

# Dbspace 2 -- Chunk 2
onspaces -c -d datadbs1 -k 2 -p
/export/home/informix/data/datadbs1150fc4
-o 0 -s 5000000 -en 100 -ef 100

# Dbspace 3 -- Chunk 3
onspaces -c -d datadbs2 -k 2 -p
/export/home/informix/data/datadbs2150fc4
-o 0 -s 5000000 -en 100 -ef 100

# Dbspace 4 -- Chunk 4
onspaces -c -d datadbs3 -k 8
-p /export/home/informix/data/datadbs3_1150fc4
-o 16 -s 80000 -en 400 -ef 400
-m /export/home/informix/data/mdatadbs3_1150fc4 16

# Dbspace 5 -- Chunk 5
onspaces -c -d tempdbs -k 2 -t -p
/export/home/informix/data/tempdbs_1150fc4 -o 0 -s 1000

```

```
# Dbspace 6 -- Chunk 6
onspaces -c -S sbspace -p
/export/home/informix/data/sbspace_1150fc4
-o 0 -s 1000 -Ms 500

# Dbspace 7 -- Chunk 7
onspaces -c -b blobdbbs -g 4 -p
/export/home/informix/data/blobdbbs_1150fc4 -o 0 -s 1000

# External Space 1
onspaces -c -x extspace -l
/export/home/informix/data/extspac_1150fc4

# Logical Log 1
onparams -a -d rootdbbs -s 10000
```

Role creation

The dbschema -r command generates information on the creation of roles.

The following syntax diagram shows the creation of roles.

Roles

```
"[-r{role | all}]"
```

Element	Purpose	Key Considerations
-r role	Displays the CREATE ROLE and GRANT statements that are needed to replicate and grant the specified role.	You cannot specify a list of users or roles with the -r option. You can specify either one role or all roles.
-r all	Displays all CREATE ROLE and GRANT statements that are needed to replicate and grant all roles.	None

The following **dbschema** command and output show that the role **calen** was created and was granted to **cathl**, **judith**, and **sallyc**:

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

DBSCHEMA Schema Utility
Software Serial Number RDS#N0000000
create role calen;

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

# Privileges

The `dbschema -p` command generates information on privileges.

The following syntax diagram fragment shows privileges information.



Element	Purpose	Key Considerations
-p user	Displays the GRANT statements that grant privileges to <i>user</i> , where <i>user</i> is a user name or role name. Specify only one user or role	You cannot specify a specific list of users with the <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

You can generate **dbschema** information about the grantor of a GRANT statement.

In the **dbschema** output, the AS keyword indicates the grantor of a GRANT statement. The following example output indicates that **norma** issued the GRANT statement:

```
GRANT ALL ON 'tom'.customer TO 'claire' AS 'norma'
```

When the GRANT and AS keywords appear in the **dbschema** output, you might need to grant privileges before you run the **dbschema** output as an SQL script. Referring to the previous example output line, the following conditions must be true before you can run the statement as part of a script:

- User **norma** must have the Connect privilege to the database.
- User **norma** must have all privileges WITH GRANT OPTION for the table **tom.customer**.

For more information about the GRANT, GRANT FRAGMENT, and REVOKE FRAGMENT statements, see the *HCL® Informix® Guide to SQL: Syntax*.

## Displaying privilege information for a role

You can generate **dbschema** information about the privileges that were granted for a particular role.

A *role* is a classification with privileges on database objects granted to the role. The DBA can assign the privileges of a related work task, such as an engineer, to a role and then grant that role to users, instead of granting the same set of privileges to every user. After a role is created, the DBA can use the GRANT statement to grant the role to users or to other roles.

For example, issue the following **dbschema** command and to display privileges that were granted for the **calen** role.

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

An example of information the **dbschema** utility displays is:

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

## Distribution information for tables in dbschema output

The **dbschema -hd** command with the name of the table retrieves the distribution information that is stored for a table in a database. If you specify the ALL keyword for the table name, the distributions for all the tables in the database are displayed.

During the **dbimport** operation, distribution information is created automatically for leading indexes on non-opaque columns. Run the UPDATE STATISTICS statement in MEDIUM or HIGH mode to create distribution information about tables that have the following types of indexes:

- Virtual Index Interface (VII) or function indexes
- Indexes on columns of user-defined data types
- Indexes on columns of built-in opaque data types (such as BOOLEAN or LVARCHAR)

Output from the **dbschema** utility shows distribution information if you used the SAMPLING SIZE keywords when UPDATE STATISTICS in MEDIUM or HIGH mode ran on the table.

For information about using the UPDATE STATISTICS statement, see the *HCL® Informix® Guide to SQL: Syntax*.

The output of **dbschema** for distributions is provided in the following parts:

- Distribution description
- Distribution information
- Overflow information

Each section of **dbschema** output is explained in the following sections. As an example, the discussion uses the following distribution for the fictional table called **invoices**. This table contains 165 rows, including duplicates.

You can generate the output for this discussion with a call to **dbschema** that is similar to the following example:

```
dbschema -hd invoices -d pubs_stores_demo
```

## Example of dbschema output showing distribution information

The **dbschema** output can show the data distributions that have been created for the specified table and the date when the UPDATE STATISTICS statement that generated the distributions ran.

The follow example of **dbschema** output shows distribution information.

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

```
High Mode, 10.000000 Resolution
```

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

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

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

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

## Description of the distribution information in the example

The first part of the sample **dbschema** output describes which data distributions have been created for the specified table. The name of the table is stated in the following example:

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

The output is for the **invoices** table, which is owned by user `cathl`. This data distribution describes the column **invoice\_num**. If a table has distributions that are built on more than one column, **dbschema** lists the distributions for each column separately.

The **dbschema** output also shows the date when the UPDATE STATISTICS statement that generated the distributions ran. You can use this date to tell how outdated your distributions are.

The last line of the description portion of the output describes the mode (MEDIUM or HIGH) in which the distributions were created, and the resolution. If you create the distributions with medium mode, the confidence of the sample is also listed. For example, if the UPDATE STATISTICS statement runs in HIGH mode with a resolution of `10`, the last line appears as the following example shows:

```
High Mode, 10.000000 Resolution
```

## Distribution information in dbschema output

The distribution information in **dbschema** output describes the bins that are created for the distribution, the range of values in the table and in each bin, and the number of distinct values in each bin.

Consider the following example:

```

      (              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 in dbschema output

The last portion of the **dbschema** output shows values that have many duplicates.

The number of duplicates of indicated values must be greater than a critical amount that is determined as approximately 25 percent of the resolution times the number of rows. If left in the general distribution data, the duplicates would skew the distribution, so they are moved from the distribution to a separate list, as the following example shows:

```

--- OVERFLOW ---

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

```

For this example, the critical amount is  $0.25 * 0.10 * 165$ , or `4.125`. Therefore, any value that is duplicated five or more times is listed in the overflow section. Two values in this distribution are duplicated five or more times in the table: the value `56` is duplicated five times, and the value `63` is duplicated six times.

## Use dbschema output as DB-Access input

You can use the **dbschema** utility to get the schema of a database and redirect the **dbschema** output to a file. Later, you can import the file into and use to re-create the schema in a new database.

## Inserting a table into a dbschema output file

You can insert CREATE TABLE statements into the **dbschema** output file and use this output as input.

The following example copies the CREATE TABLE statements for the customer table into the **dbschema** output file, **tab.sql**:

```
dbschema -d db -t customer > tab.sql
```

Remove the header information about **dbschema** from the output file, **tab.sql**, and then use 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 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.

### About this task

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 to create a new database.
4. Use 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.

### Results

The databases **db** and **testdb** differ in name but have the same schema.

## The LOAD and UNLOAD statements

You can use the SQL LOAD and UNLOAD statements to move data. The LOAD statement is moderately fast and easy to use, but it only accepts specified data formats. You usually use the LOAD statement with data that is prepared with an UNLOAD statement.

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

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™, , or Windows™ utilities to load those files onto tape.



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**.

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

## Syntax of the UNLOAD statement

The UNLOAD statement unloads selected rows from a table into a text file.

```
UNLOAD TO ' filename ' [ DELIMITER ' delimiter ' ] <SELECT Statement> (explicit id)
```

Element	Purpose	Key Considerations
<i>delimiter</i>	Character to use as delimiter	<b>Requirements:</b> See <a href="#">Syntax for the delimiter form on page 93</a>
<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 [UNLOAD statement on page](#) .

## Syntax of the LOAD statement

The LOAD statement appends rows to an existing table of a database.

```
(explicit id unique_166_Connect_42_note010) unique_166_Connect_42_note010 LOAD FROM ' filename ' [ DELIMITER ' delimiter ' ] INSERT INTO {  
<Table Name> | <Synonym Name> | <View Name> } [ ( column ) ]
```

Element	Purpose	Key Considerations
<i>column</i>	The name of a column to receive data from <i>filename</i>	Must be a column in the specified table or view.
<i>delimiter</i>	Character to use as delimiter	See <a href="#">Syntax for the delimiter form on page 93</a>
<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 [LOAD statement on page](#) .

## Load and unload statements for locales that support multibyte code sets

For locales that support multibyte code sets, be sure that the declared size (in bytes) of any column that receives character data is large enough to store the entire data string.

For some locales, this can require up to 4 times the number of logical characters in the longest data string.

### Load and unload operations for nondefault locales and GL\_DATETIME or DBTIME environment variables

In nondefault locales, operations that load or unload DATETIME or INTERVAL values can be sensitive to the settings of the **GL\_DATETIME**, **DBTIME**, and **USE\_DTEENV** environment variables.

If the database uses a nondefault locale and the **GL\_DATETIME** or **DBTIME** environment variable has a nondefault setting, you must set the **USE\_DTEENV** environment variable to the value of 1 before you can process localized DATETIME or INTERVAL values correctly

- with the LOAD or UNLOAD statements of **DB-Access**,
- or with the **dbimport** or **dbexport** utilities,
- or in DML operations on objects that the CREATE EXTERNAL TABLE statement defined.

### The onunload and onload utilities

The **onunload** and **onload** utilities provide the fastest way to move data between computers that use the same database server on the same platform.

For example, your site purchases a more powerful UNIX™ computer to allow faster access for users. You need to transfer existing databases to the new database server on the new computer. Use **onunload** to unload data from the first database server and then use **onload** to load the data into the second database server. Both database servers must have the same version number, or they must have compatible version numbers. You can move an entire database or selected tables only, but you cannot modify the database schema.

The **onunload** utility can unload data more quickly than either **dbexport** or the UNLOAD statement because **onunload** copies the data in binary format and in page-sized units. The **onload** utility takes a tape or a file that the **onunload** utility creates and re-creates the database or the table.

The **onunload** and **onload** utilities are faster than **dbimport**, **dbload**, or LOAD but are much less flexible and do not let you modify the database schema or move from one operating system or database server version to another.

### Guidelines for when to use the onunload and onload utilities

You can use **onunload** and **onload** only when certain conditions are met.

You can use only **onunload** and **onload** if your answer to each of the following questions is *yes*. If your answer is *no*, you cannot use **onunload** and **onload**.

Use onunload and onload only if your answer to each question is yes	My answer
Is the target database server on the same hardware platform?	
Do you want to move to another database server of the same version?	
Do you want to keep the existing database schema without modifying it?	

Use <b>onunload</b> and <b>onload</b> only if your answer to each question is yes	My answer
Do you want to move an entire database or an entire table?	
Are the page images compatible?	
Are the numeric representations the same?	

### When you cannot use the **onunload** and **onload** utilities

Because the data is written in page-sized units, you cannot use **onunload** and **onload** to move data between UNIX™ or and Windows™ because they use different page sizes. For example, the page size is 2 KB on some UNIX™ systems and 4 KB on Windows™.

Additionally, you cannot use **onunload** and **onload**:

- To move data between GLS and non-GLS databases.
- To move compressed data from one database to another.

You must uncompress data in compressed tables and fragments before you use the **onload** and **onunload** utilities.

- To move external tables or databases that contain external tables.

You must drop all the external tables before you use the **onunload** utility.

- To move tables and databases that contain extended or smart-large-object data types

### Requirements for using the **onload** and **onunload** utilities

The **onload** and **onunload** utilities have limitations. You can use these utilities only to move data between database servers of the same version on the same operating system. You cannot modify the database schema, logging must be turned off, and the utilities can be difficult to use.

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

- The original database and the target database must be from the same version of the database server. You cannot use the **onload** and **onunload** utilities to move data from one version to another version.
- You cannot use **onload** and **onunload** to move data between different types of database servers.
- The **onload** command must have the same scope as the corresponding **onunload** command that unloaded the same table or tables that **onload** references. You cannot, for example, use **onunload** to unload an entire database, and then use **onload** to load only a subset of the tables from that database.
- Do not use **onload** and **onunload** to move data if the database contains extended or smart-large-object data types.
- Because the tape that **onload** reads contains binary data that is stored in disk-page-sized units, the computers where the original database resides (where you use **onunload**) and where the target database will reside (where you use **onload**) must have the same page size, the same representation of numeric data, the same byte alignment for structures and unions.
- You cannot use **onload** and **onunload** to move data between non- and locales.

- You cannot use **onload** and **onunload** on servers in high-availability clusters.
- You cannot use **onload** and **onunload** if you compressed tables or fragments.

You can use **onunload** and **onload** to move data between databases if the NLS and locales are identical.

If the page sizes are different, **onload** fails. If the alignment or numeric data types on the two computers are different (for example, with the most significant byte as last instead of first, or different float-type representations), the contents of the data page could be misinterpreted.

### How the onunload and onload utilities work

The **onunload** utility, which unloads data from a database, writes a database or table into a file on tape or disk. The **onload** utility loads data that was created with the **onunload** command into the database server.

The **onunload** utility unloads the data in binary form in disk-page units, making this utility more efficient than **dbexport**.

You can use the **onunload** utility to move data between computers that have the same version of the database server.



**Important:** You cannot use the **onload** and **onunload** utilities to move data from one version of a database server to another or between different types of database servers. In addition, the **onload** command must have the same scope as the corresponding **onunload** command that unloaded the same table or tables that **onload** references. You cannot, for example, use **onunload** to unload an entire database, and then use **onload** to load only a subset of the tables from that database.

The **onload** utility creates a database or table in a specified dbspace. The **onload** utility then loads it with data from an input tape or disk file that the **onunload** utility creates.

During the load, you can move simple large objects that are stored in a blobspace to another blobspace.

### Syntax of the onunload command

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

```
onunload [ < -FILE option > (explicit id) ] [ <Destination Parameters> (explicit id) ] database [ : [ owner. ] table ] [ { -v | -version } ]
```

Element	Purpose	Key Considerations
<i>database</i>	Specifies the name of a database	<p><b>Additional Information:</b> The database name cannot be qualified by a database server name (<i>database@dbservername</i>).</p> <p><b>References:</b> Syntax must conform to the Identifier segment; see <a href="#">Identifier on page</a> .</p>
<i>owner.</i>	Specifies the owner of the table	<p><b>Additional Information:</b> The owner name must not include invalid characters.</p>

Element	Purpose	Key Considerations
		<b>References:</b> For path name syntax, see your operating-system documentation.
<i>table</i>	Specifies the name of the table	<b>Requirement:</b> The table must exist.  <b>References:</b> Syntax must conform to the Identifier segment; see <a href="#">Identifier on page</a> .

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 *HCL® Informix® Administrator's Reference*.)

The **-V** option displays the software version number and the serial number. The **-version** option extends the **-V** option to display additional information about the build operating system, build number, and build date.

## onunload destination parameters

The **onunload** utility supports tape or file destination options.

The following syntax diagram fragment shows **onunload** destination parameters

### Destination parameters

```
" { | -l | -b blocksize | -s tapesize | -t source } (explicit id) "
```

Element	Purpose	Key Considerations
<b>-b <i>blocksize</i></b>	Specifies in kilobytes the block size of the tape device	<b>Requirement:</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>Requirement:</b> The <i>tapesize</i> must be an integer. To write to the end of the tape, specify a tape size of 0.  If you do not specify 0, then the maximum <i>tapesize</i> is 2 097 151 KB.

Element	Purpose	Key Considerations
		<b>Additional Information:</b> This option overrides the default value in TAPESIZE or LTAPESIZE.
<b>-t source</b>	Specifies the path name 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. The path name must be a valid path name.

## Constraints that affect onunload

When you use the **onunload** utility, you must be aware of constraints that affect how you load the data on the **onunload** tape.

The following constraints apply to **onunload**:

- You must load the data on the **onunload** tape into a database or table that your database server manages.
- You cannot use **onunload** and **onload** if the databases contain extended data types.
- You must load the tape that **onunload** writes onto a computer with the same page size and the same representation of numeric data as the original computer.
- You must read the file that **onunload** creates with the **onload** utility of the same version of your database server. You cannot use **onunload** and **onload** to move data from one version to another.
- When you unload a complete database, you cannot modify the ownership of database objects (such as tables, indexes, and views) until after you finish reloading the database.
- When you unload and load a table, **onunload** does not preserve access privileges, synonyms, views, constraints, triggers, or default values that were associated with the original tables. Before you run **onunload**, use the **dbschema** utility to obtain a listing of the access privileges, synonyms, views, constraints, triggers, and default values. After you finish loading the table, use **dbschema** to re-create the specific information for the table.

## Privileges for database or table unloading

To unload a database, you must have DBA privileges for the database or be user **informix**. To unload a table, you must either own the table, have DBA privileges for the database in which the table resides, or be user **informix**.

User **root** does not have special privileges with respect to **onunload** and **onload**.

## Tables that are unloaded with a database

If you unload a database, all of the tables in the database, including the system catalog tables, are unloaded.

All triggers, SPL routines, defaults, constraints, and synonyms for all of the tables in the database are also unloaded.

## Data that is unloaded with a table

If you unload a table, **onunload** unloads the table data and information from the **systables**, **systables**, **syscolumns**, **sysindexes**, and **sysblobs** system catalog tables.

When you unload a table, **onunload** does not unload information about constraints, triggers, or default values that are associated with a table. In addition, access privileges that are defined for the table and synonyms or views that are associated with the table are not unloaded.

## Locking during unload operation

During the unload operation, the database or table is locked in shared mode. An error is returned if **onunload** cannot obtain a shared lock.

The **onload** utility creates a database or table in a specified dbspace. The **onload** utility then loads it with data from an input tape or disk file that the **onunload** utility creates.

## Logging mode

The **onunload** utility does not preserve the logging mode of a database. After you load the database with **onload**, you can make a database ANSI compliant or add logging.

For information about logging modes, refer to the *HCL® Informix® Guide to SQL: Syntax*.

During the load, you can move simple large objects that are stored in a blob space to another blob space.

If you do not specify any source-parameter options, **onload** uses the device that is specified as TAPEDEV. The block size and tape size are the values that are specified as TAPEBLK and TAPESIZE, respectively. (For more information about TAPEDEV, TAPEBLK, and TAPESIZE, refer to your *HCL® Informix® Administrator's Guide*.)

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

## Syntax of the onload command

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

```
onload [ < -FILE option > (explicit id) ] [ <Source Parameters> (explicit id) ] { -adb spacedatabase [ : [ owner. ] table ] | <Create Options> (explicit id) } [ { -v | -version } ]
```

Element	Purpose	Key Considerations
<b>-d</b> <i>dbspace</i>	Loads a database or table into the specified dbspace	The tape being loaded must contain the specified database or table.
<i>database</i>	Specifies the name of the database	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 <a href="#">Identifier on page</a> .
<i>owner.</i>	Specifies the owner of the table	The owner name must not include invalid characters.

Element	Purpose	Key Considerations
		<b>References:</b> For path name syntax, refer to your operating-system documentation.
<i>table</i>	Specifies the name of the table	The table must exist.  <b>References:</b> Syntax must conform to the Identifier segment; see <a href="#">Identifier on page</a> .

The **-V** option displays the software version number and the serial number. The **-version** option extends the **-V** option to display additional information about the build operating system, build number, and build date.

## onload source parameters

The **onload** command includes options for specifying information about the tape or file source.

The following syntax diagram fragment shows **onload** source parameters.

### Source parameters

```
" { | -l | -b blocksize | -s tapesize | -t source } (explicit id) "
```

Element	Purpose	Key Considerations
<b>-b <i>blocksize</i></b>	Specifies in kilobytes the block size of the tape device	<b>Requirements:</b> Unsigned integer. Must specify the block size of the tape device.  <b>Additional Information:</b> This option overrides the default value in TAPEBLK or LTAPEBLK.
<b>-l</b>	Directs <b>onload</b> to read the values for tape device, block size, and tape size from the configuration parameters LTAPEDEV, LTAPEBLK, and LTAPESIZE, respectively	<b>Additional Information:</b> If you specify <b>-l</b> , and then <b>-b</b> , <b>-s</b> , or <b>-t</b> , the value that you specify overrides the value in the configuration file.
<b>-s <i>tapesize</i></b>	Specifies in kilobytes the amount of data that the database server can store on the tape	<b>Requirements:</b> Unsigned integer. To write to the end of the tape, specify a tape size of 0.  If you do not specify 0, then the maximum <b>tapesize</b> is 2 097 151 KB.  <b>Additional Information:</b> This option overrides the default value in TAPESIZE or LTAPESIZE.



Element	Purpose	Key Considerations
<b>-t source</b>	Specifies the path name of the file on disk or of the tape device where the input tape is mounted	<p>Must be a legitimate path name.</p> <p><b>Additional Information:</b> This option overrides the tape device that TAPEDEV or LTAPEDEV specifies.</p> <p><b>References:</b> For path name syntax, see your operating-system documentation.</p>

## onload create options

The **onload** command includes information that is used to recreate the database.

The following syntax diagram fragment shows **onload** create options.

### Create options

```
" { | -c oldcnstrnt newcnstrnt | -i oldindex newindex | -fd olddbsp newdbsp | -fi indexname olddbsp newdbsp } "
```

```
" database [ : [ owner. ] table ] "
```

Element	Purpose	Key Considerations
<b>-c oldcnstrnt newcnstrnt</b>	Directs <b>onload</b> to rename the specified constraint.	None.
<b>-i oldindex newindex</b>	Directs <b>onload</b> to rename the table index when it stores the index on disk.	<p><b>Additional Information:</b> Use the <b>-i</b> option to rename indexes during the load to avoid conflict with existing index names.</p> <p><b>References:</b> Syntax must conform to the Identifier segment; see <a href="#">Identifier on page</a> .</p>
<b>-fd olddbsp newdbsp</b>	Moves a data fragment from one dbspace to another.	<p>The new dbspace must exist and must not already contain another data fragment for the table.</p> <p><b>Additional Information:</b> This option is used with parallel data query (PDQ) and table fragmentation.</p>
<b>-fi indexname olddbsp newdbsp</b>	Moves index fragments from one dbspace to another.	<p>The new dbspace must exist and must not already contain another index fragment for the table.</p> <p><b>Additional Information:</b> This option is used with PDQ and table fragmentation.</p>

Element	Purpose	Key Considerations
<i>database</i>	Specifies the name of the database	<b>Requirement:</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 <a href="#">Identifier on page</a> .
<i>owner.</i>	Specifies the owner of the table	<b>Requirement:</b> The owner name must not include invalid characters.  <b>References:</b> For path name syntax, refer to your operating-system documentation.
<i>table</i>	Specifies the name of the table	<b>Requirement:</b> The table must not exist.  <b>References:</b> Syntax must conform to the Identifier segment; see <a href="#">Identifier on page</a> .

If you do not specify any create options for non-fragmented tables, the **onload** utility stores the database or table in the root dbspace.

For fragmented tables, **onunload** preserves the fragmentation expression for later use by **onload**. Thus an imported table is fragmented in the same way as the original table.

You can use the **-c**, **-i**, **-fd**, and **-fi** options in any order and as often as necessary as long as you use unique pairs.

## Constraints that affect onload

The **onload** utility performs faster than the **dbimport**, **dbload**, or **LOAD** methods. In exchange for this higher performance, **onload** has certain constraints.

The **onload** utility has the following constraints:

- The **onload** utility only creates a new database or table; you must drop or rename an existing database or table of the same name before you run **onload**. During execution, the **onload** utility's prompt will ask you if you want to rename blobspaces.
- The **onload** utility places a shared lock on each of the tables in the database during the load. Although you cannot update a table row with the lock in place, the database is available for queries.
- When you load a complete database, the user who runs **onload** becomes the owner of the database.
- The **onload** utility creates a database without logging; you must initiate logging after **onload** loads the database.
- When you use **onload** to load a table into a logged database, you must turn off logging for the database during the operation.
- For fragmented tables, the dbspace assignment is preserved, unless you override it using the **-fn** option.
- For non-fragmented tables, the **onload** utility attempts to store the table in root dbspace if a target dbspace is not specified with the **-d** option. If storing the table in root dbspace or in the dbspace specified with the **-d** option is

not possible due to difference in page sizes, the **onload** utility tries to use a dbspace that has the same dbspace number as the dbspace number of the originally unloaded table. If this dbspace still has a different page size, the load operation will fail.

## Logging during loading

When you use the **onload** utility to create tables from an **onunload** input tape, **onload** can load information only into a database without logging. Thus, before you load a table into an existing, logged database, you must end logging for the database.

You also might want to consider loading during off-peak hours. Otherwise, you might fill the logical-log files or consume excessive shared-memory resources. After you load the table, create a level-0 dbspace backup before you resume database logging.

When you use **onload** to create databases from an **onunload** input tape, the databases that result are not ANSI-compliant and do not use transaction logging. You can make a database ANSI compliant or add logging after you load the database.

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

## Movement of simple large objects to a blobspace

If you load a table that contains simple large objects stored in a blobspace, the **onload** utility 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 the **onload** utility 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 [The dbschema utility on page 103](#) describes, before you run **onunload**.

## Exclusive locking during a load operation

During a load operation, the **onload** utility 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.

## Moving a database between computers with the onunload and onload utilities

You can use the onunload and onload utilities to move a complete database from one computer to another.

To move a database from one computer to another:

1. Make sure that the page size, numeric representations, and byte alignment on structures and unions are the same on both computers.

The page size is 2 KB on certain UNIX™ systems and 4 KB on Windows™. For information about page size, see your *HCL® 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 the TAPEDEV or LTAPEDEV configuration parameter, or specify another tape device. Make sure that the tape device that you specify is available for onunload. However, if you set the TAPEDEV configuration parameter to STDIO, the onunload utility will not be able to unload data.
3. Run the oncheck utility to make sure that your database is consistent.

For information about oncheck, see your *HCL® Informix® Administrator's Reference*.

4. Run the onunload utility to unload the data from the database.

For details on the syntax of the onunload command, see [Syntax of the onunload command on page 124](#).

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.

For details on the syntax of the onload command, see [Syntax of the onload command on page 127](#).

7. Set the logging status for the new database.

For information about logging status, see your *HCL® Informix® Administrator's Guide*.

8. If necessary, change the DBA privileges of the database.
9. Create a level-0 backup of the new database.

## Moving a table between computers with the onunload and onload utilities

You can use the **onunload** and **onload** utilities to move a table from one computer to another.

To move a table from one computer to another:

1. Make sure that the page size, numeric representations, and byte alignment on structures and unions are the same on both computers. (The page size is 2 KB on certain UNIX™ systems and 4 KB on Windows™.)
2. Decide where to store the unloaded data.
3. Run the **oncheck** utility to make sure that your database is consistent.
4. If you want to save the triggers, access privileges, SPL routines, defaults, constraints, and synonyms for the table, run the **dbschema** utility.
5. Run the **onunload** utility.

For details on the syntax of the **onunload** command, see [Syntax of the onunload command on page 124](#).

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.

For details on the syntax of the **onload** command, see [Syntax of the onload command on page 127](#).

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 that are not preserved, and synonyms for the table, run the **dbschema** utility or recreate these objects manually.

Constraints such as primary keys or default values are preserved, even for a single table. Foreign keys, access privileges, SPL routines and synonyms are not preserved.

## Moving a table between dbspaces with the onunload and onload utilities

You can use the **onunload** and **onload** utilities to move a table from one dbspace to another dbspace on the same computer.

To move a table from one dbspace to another dbspace on the same computer:

1. Run the **onunload** utility to unload the table.

For details on the syntax of the **onunload** command, see [Syntax of the onunload command on page 124](#).

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** command.

For details on the syntax of the **onload** command, see [Syntax of the onload command on page 127](#).

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.

## The onmode utility reversion option

You use the **-b** option of the **onmode** utility to revert to the older database server from which you converted.

### What the onmode -b command does

When you convert a database server, several modifications make the format of the databases incompatible with the older version. The **onmode -b** command modifies data so that the earlier version of the database server can access it.

When you convert a database server, several modifications make the format of the databases incompatible with the older version. The **onmode -b** command modifies data so that the earlier version of the database server can access it. In some cases the format of the databases is compatible between versions and the **onmode -b** command is not needed. Type **onmode -b** to see the usage message for options that are available for your database server.

The utility does not revert changes made to the layout of the data that do not affect compatibility.

You must revert the databases before users can access the data with the earlier database server version.

For information about other **onmode** options, see [The onmode utility on page](#) in your *HCL® Informix® Administrator's Reference*.

### Preparation for using the onmode -b command

Before you use the **onmode -b** command, 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 **onmode -b** command includes an implicit **-yuk** command.

Make sure that the **INFORMIXSERVER** environment variable is set to the correct database server.

#### UNIX/Linux Only

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

#### Windows™ Only

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

### Syntax of the onmode -b command

The **onmode -b** command restores your databases to the version of Informix® from which you converted. You cannot use this command to revert to any other version of the server.

When you convert a database server, several modifications make the format of the databases incompatible with the older version. The **onmode -b** command modifies data so that the earlier version of the database server can access it. In some cases the format of the databases is compatible between versions and the **onmode -b** command is not needed.



**Tip:** To see a list of all of the versions to which you can revert, run this command:

```
onmode -b
```

When you see the options that are available, choose the option that is closest to the version that you want.

## Syntax

```
onmode { -b version_number }
```

Element	Purpose
-b <i>version_number</i>	Reverts the database to the specified version.

## The onrestorept utility

You can use the onrestorept utility to restore the database server back to the state that it was in just before the start of the failed upgrade.

The utility depends on these configuration parameters being set correctly before you attempt to upgrade the database server:

- The CONVERSION\_GUARD configuration parameter must be set to **1** or **2** (the default value) for data to be stored during an upgrade. If the conversion guard operations fail (for example, because the server has insufficient space to store the data that is captured during the upgrade), you cannot use the onrestorept utility.
- The RESTORE\_POINT\_DIR configuration parameter must specify a directory where the conversion guard utility can store data files during the upgrade. Those files are needed to restore the database server to a consistent state after a failed upgrade. This directory must be empty before the upgrade starts. After a successful upgrade, the contents of the directory are automatically removed.



**Important:** must be offline when you run the onrestorept utility. Do not start the server until the utility finishes running. Executing the onrestorept utility when the server is online, or starting the server before the utility finishes running, can damage the database and requires you to restore the database server from a backup copy.

To start Enterprise Replication after the onrestorept utility restores the database server to a consistent state, you must use the cdr cleanstart command.

## Syntax of the onrestorept command

The onrestorept command undoes changes made during a failed upgrade, restoring files to the state they were in when you shut down the server.

```
onrestorept [ -v ][ -c ][ -y ]
```

Element	Purpose	Key Considerations
-V	Display the version of the current server and the software serial number.	
-c	After a failed upgrade, delete the files in the directory specified in the RESTORE_POINT_DIR configuration parameter.	<p>Before you begin another upgrade, you must delete these restore point files. Do this before you make another migration attempt, but not before you run the <b>onrestorept</b> utility to recover the files (if possible).</p> <p>If the upgrade was successful, restore point files are automatically deleted and there is no need to run onrestorept -c.</p>
-y	Automatically accept every prompt while the onrestorept command runs.	If you do not specify -y, you must respond to every prompt.

### Example

### Examples

The following command restores files after a failed upgrade:

```
onrestorept
```

The following command removes restore point files after a failed upgrade:

```
onrestorept -c
```



# Index

## A

- Abbreviated years 72
- add chunk argument 20
- admin() functions
  - add chunk argument 20
- ALARMPROGRAM configuration parameter 47, 48

## B

- Backups
  - after upgrading 59
  - before reverting 63
  - before upgrading to a new version 18
  - logical logs 48
  - ON-Bar utility 18, 69
  - ontape utility 18, 69
  - source database 18
- Binary files, loading 122, 123
- BladeManager
  - installing and registering DataBlade modules 59, 69
  - removing new extensions before reversion 65
- Blobspaces
  - moving, with onunload and onload 131, 132
- BSON date fields 13, 54, 68
- migrating 13

## C

- Character-position form of FILE and INSERT statements 96
- Checking available space
  - before migration 9
- Chunks
  - reverting reserve pages 62
- client applications 56
- client compatibility 50, 56
- Clusters
  - Apply fix pack or PID 28
  - migrating to new release 27
  - migrating to new version 33
  - restoring from a backup archive 44
  - restoring from the HDR secondary server 44
  - reverting 27, 42
  - upgrading 36
  - upgrading to a new fix pack 32
  - upgrading to a new PID 32
  - upgrading with conversion 33
- Command file
  - dbload 92
- Communications Support Module
  - configuring after migration 48
  - removing if reverting 67
  - saving before reverting 62
- concsn.cfg file
  - creating entries after migration 48
  - removing if reverting 67
  - saving before reverting 62
- Configuration file
  - customizing after migration 47
  - replacing after reversion 66
  - saving before migration 12
  - saving before reverting 62
- Configuration parameters
  - ALARMPROGRAM 47, 48
  - CONVERSION\_GUARD 49, 135
  - HPL\_DYNAMIC\_LIB\_PATH 47

- RESTORE\_POINT\_DIR 49, 135
- ROOTOFFSET 47
- ROOTPATH 47
- ROOTSIZE 47
- STOP\_APPLY 72
- UPDATABLE\_SECONDARY 72
- USELASTCOMMITTED 72
- CONVERSION\_GUARD configuration parameter 49, 135
- convTovNoSQL1210.sql script 54, 68

## D

- Data integrity 58
- data migration
  - external tables 71
- Data migration
  - constraints 3
  - issues to consider 3
  - overview 3
  - prerequisites 3
  - tools 3
- data types
  - SQLINTEGER 56
  - SQLLEN 56
  - SQLUINTEGER 56
  - SQLULEN 56
- Database server
  - initializing after upgrading 48
  - new version
    - performance tuning 59
  - reverting 60
  - reverting from current version 65
  - starting after upgrading 48
- Database servers
  - migrating 45, 45, 54
  - preparing for migration 8
  - upgrading 45, 45, 54
- Databases
  - ownership, set by onload 131
- DataBlade modules
  - installing after upgrading 48
  - registering 59, 69
- DB-Access
  - input from the dbschema utility 119, 119
- dbexport
  - SELECT triggers, disabling 72
- dbexport utility 72
  - c option 74, 77
  - d option 74
  - nw option 74
  - q option 74
  - si option 74, 77
  - ss option 74, 77
  - V option 74
  - version option 74
  - X option 74
  - destination options 77, 77
  - Interrupt key 76
  - schema output 79
  - syntax 74
- dbimport utility 72
  - c option 80, 82
  - D option 80
  - I option 86
  - nv option 80
  - q option 80
  - V option 80
  - version option 80
- X option 80
  - create options 84
  - database logging mode 86
  - importing from another computer 7
  - input file location options 83, 83
  - Interrupt key 82
  - locale, changing 87
  - renaming a database 86
  - syntax 80
  - using with GLS 80
  - using with NLS 87
- dbload utility
  - c command file option 88
  - d database option 88
  - e errors option 88
  - e option 91
  - i ignore rows option 88
  - i option 91
  - k option 88
  - l error log file option 88
  - p option 88
  - r option 88, 90
  - s option 88
  - V option 88
  - version option 88
  - X option 88
  - compared to LOAD 88
  - creating a command file 92
  - dbload utility
    - n commit interval option 88
  - FILE statement 92
  - guidelines for handling objects 91
  - ignoring rows 91
  - importing from another computer 7
  - INSERT statements 92
    - compared to SQL INSERT statement 97
    - using 93
  - Interrupt key 91
  - number errors to allow 91
  - overview 88
  - speed, increasing 91
  - syntax 88
  - table locking 90
  - writing a command file
    - in character-position form 98
    - in delimiter form 95
- dbschema utility
  - ss option 109
  - u option 109
  - chunk schema 113, 113
  - create schema across a network 108
  - create schema for a database 108
  - distribution information 117
  - example of file for DB-Access 119
  - guidelines 104
  - log schema 113, 113
  - output example 117
  - overview 103
  - owner conventions 108
  - privileges information 116, 116
  - privileges information for a role 116
  - re-creating the schema 120
  - sequence schema 110
  - specifying a table, view, or procedure 111
  - storage space schema 113, 113
  - synonym schema 111
  - syntax 104
  - syntax for role schema 115

- table information 112
- DBSECADM role 72
- dbspaces
  - moving tables to another dbspace 133
- Delimiter form of FILE and INSERT statements 92, 93, 95
- Diagnostic information to gather 18
- Directories
  - installation 46

## E

- Environment variables
  - DB\_LOCALE 87
  - DBCENTURY 72
  - DBTEMP 87
  - GL\_DATE 72
  - GL\_DATETIME 122
  - GL\_USEGLU 46
  - INFORMIXSERVER 46
  - INFORMIXSQLHOSTS 46
  - ONCONFIG 46
  - PATH 46
  - resetting after reversion 67
  - TEMP 87
  - TMP 87
  - USE\_DTENV 122
- Exporting
  - time series data 79
- external tables 71
- Extracting schema information 70

## F

- Fast recovery
  - initiating 16
- Features
  - reviewing 9
- FILE statement
  - character-position form 96
  - delimiter form 92, 93, 95
  - syntax for
    - character-position form 97
    - delimiter form 93
  - with dbload 92
- Fix packs
  - Apply to clusters 28
- Functions, SQL administration API
  - add chunk argument 20

## G

- GL\_DATETIME environment variable 72, 122
- GL\_USEGLU environment variable 46
- Global Language Support (GLS)
  - dbimport utility 80
  - using onload and onunload 123
- GRANT statement
  - role privileges 116

## H

- HEX binary data 80
- HPL\_DYNAMIC\_LIB\_PATH configuration parameter 47

## I

- id\_column 14, 63
- Importing
  - non-
    - Informix
      - data
        - 7
- in dbschema output 116, 116
- In-place alters
  - oncheck -pT command 14, 63

- Sample updates 14, 63
- Index
  - rebuilding 48
- Informix
  - installing 46
- INFORMIXDIR directory 46
- INFORMIXSERVER environment variable 46
- INFORMIXSQLHOSTS environment variable 46
- Initializing
  - after upgrading 48
- INSERT statements
  - character-position form 96
  - delimiter form 92
  - syntax for character-position form 97
  - with dbload 92
- Installation directory 46
- Installing
  - Informix
    - 46

## J

- Java UDRs 67
- JSON compatibility 54, 68

## L

- Label-based access control (LBAC) 72
- large row identifiers 56
- Level-0 backup 18, 59, 69
  - after moving data 132
- LOAD SQL statement
  - for locales that support multibyte code sets 121
  - overview 120
  - syntax 121
- LOAD statement
  - for non-default GL\_DATETIME environment variable settings 122
  - for non-default locales 122
- Loading
  - ASCII files 7
  - binary data 122, 123
  - data 3, 6
- Locking
  - set by onload 131
- Logical log
  - backup 48
  - out of space 48
  - space required for migration 9
- LTAPEDEV configuration parameter
  - onunload/onload 132

## M

- Migrating
  - between 32-bit and 64-bit database servers 19
- Migrating a database
  - constraints 3
  - issues to consider 3
  - overview 3
  - planning for 3
  - prerequisites 3
  - to a new operating system 69, 70
  - tools 3
- Migrating with Enterprise Replication 22, 23
- Migration
  - before you begin 8
  - checklist 18
  - compatibility mode 35, 53
  - diagnostic information that you need before upgrading 18
  - monitoring status with online.log 46
  - native mode 35, 53

- onload utility 124
- onunload utility 124
- planning 8
- preparing for 8
- prerequisites 8
- space requirements 9
- with HDR, RS, and SD servers 25, 27
- with high-availability clusters 25, 27, 32, 33, 36
- with secondary servers 33

## Mode

- checking 18
- modifying
  - client applications 56
- Monitoring migration status 46
- Moving data
  - blobspaces 132
  - constraints 3
  - overview 3
  - using dbexport and dbimport 72
  - using dbload 88
  - using distributed SQL 7
  - using onload and onunload 122, 122
  - using onunload and onload 132, 132, 133
  - when changing operating systems 69, 70, 70, 71
- Multi-node Active Clusters for High Availability (MACH)
  - Clusters
    - Migrating to new release 25

## N

- Native Language Support (NLS)
  - populating with dbimport 87
- native mode 50
- Non-
  - Informix
    - data, importing
      - 7
- NOVALIDATE constraint mode 80

## O

- ON-Bar utility 54
  - backing up
    - after upgrading 59
  - backing up before upgrading 18
- oncheck utility
  - cc database\_name option 17
  - cd database\_name option 17
  - ce option 17
  - cl database\_name option 17
  - cr option 17
  - rebuilding table indexes 48
  - verifying database integrity 17, 58
- ONCONFIG environment variable 46
- ONCONFIG file
  - customizing after migration 47
- oninit utility
  - s option 16
- online.log 46
- onload and onunload utilities 132, 132, 133
- onload utility
  - constraints 130
  - constraints on use 123
  - create options 129
  - handling large objects in a blobspace 131
  - how it works 124
  - logging status 131
  - moving a database 132
  - moving a table 132, 133, 133
  - moving locales 123
  - moving to another dbspace 133

- ownership and privileges 131
- specifying source parameters 128
- syntax 127
- using between computers 122, 122
- onmode -b command 134, 134
- onmode utility
  - b option 66
  - ky option 16
  - sy option 16
  - reverting from the current version 66
  - shutting down 16
  - shutting down the server 16
- onmode-b command 134
  - syntax 134
- onrestorept utility
  - Clusters
    - restoring primary server to a consistent point 43
  - overview 135
  - syntax 135
  - undoing failed upgrade changes 43, 49
- onstat utility 18
- ontape utility
  - a option 48
  - backing up
    - after upgrading 59
    - before upgrading 18
- onunload utility
  - constraints on use 123, 126
  - destination parameters 125
  - how it works 124
  - locking 127
  - logging mode 127
  - moving a database 132
  - moving a table 132, 133, 133
  - moving locales 123
  - moving to another dbspace 133
  - ownership and privileges 126
  - syntax 124
  - unloading tables 126
  - using between computers 122, 122
  - what is included with a
    - database 126
    - table 126
- Operating system
  - adjusting tables after changing operating systems 70
  - moving data to another one 69, 70, 70, 71

## P

- PATH environment variable 46
- Performance tuning
  - adjusting queries after upgrading 59
  - after upgrading 59
- PID
  - Apply to clusters 28
- Planning
  - before moving data 3
  - data migration 3
  - for exporting and importing data 70
  - for migration 8
  - for upgrading your server 8
- Platforms, moving data between
  - compatible computers 7
- post-migration steps 54, 68
- Privileges 116, 116
  - required for onunload 126

## Q

- Queries
  - adjusting after upgrading 59
- Quiescent mode 18

## R

- recompiling
  - client applications 56
- Recompiling Java UDRs 67
- Registering DataBlade modules 59, 69
- restore points 135
- RESTORE\_POINT\_DIR configuration parameter 49, 135
- Reverse migration 60
- Reversion utility 66
- Revert to original version 60
- Reverting
  - backing up before you start 63
  - before using the onmode -b command 134
  - chunk reserve pages 62
  - database schema 61
  - from the new version 59
  - from
    - Version 15.0.0
    - 65
  - remove features 65
  - restrictions for reverting to prior versions 61
  - using the onmode -b command 134, 134
  - with Enterprise Replication 24
  - with HDR, RS, and SD servers 27, 42
  - with high-availability clusters 27, 42
- Rolling upgrade
  - with temporary Enterprise Replication 36
- Rolling upgrades 28
- ROOTOFFSET configuration parameter 47
- ROOTPATH configuration parameter 47
- ROOTSIZE configuration parameter 47
- rowid 50, 56
- Running the reversion utility 66

## S

- Sample updates 14, 63
- Schema
  - create across a network 108
  - create for a database 108
  - display with dbschema 103
- scripts
  - convTovNoSQL1210.sql 54, 68
- Scripts
  - concdr.bat 23
  - concdr.sh 23
- SELECT triggers, disabling with dbexport 72
- Simple large objects
  - moving with onload 131, 131, 132, 132
- Slow queries
  - adjusting after upgrading 59
- sm\_versions file 54
- smi\_unld utility 62
- Space
  - checking availability before migration 9
  - for sysmaster database 9
- SQL administration API functions
  - add chunk argument 20
- SQL statements
  - UPDATE STATISTICS
    - data distributions 117
- sqlhosts file
  - save a copy when migrating 12
- sqlhosts file, UNIX
  - changing name or path 46
  - csm option 48
- Starting
  - after upgrading 48
  - server after reversion 67
- STOP\_APPLY configuration parameter 72

- sysmaster database
  - and logical logs 48
  - space required for migration 9

## T

- TAPEDEV configuration parameter, with onunload and onload 132
- task() functions
  - add chunk argument 20
- Time series data
  - exporting 79
- Transactions
  - checking for open ones 16

## U

- UNLOAD SQL statement
  - for locales that support multibyte code sets 121
  - overview 120
  - syntax 121
- UNLOAD statement
  - for non-default GL\_DATETIME environment variable settings 122
  - for non-default locales 122
- UPDATABLE\_SECONDARY configuration parameter 72
- UPDATE statements
  - sample 14, 63
- UPDATE STATISTICS statement 56
  - data distributions 117
  - using after reversion 68
- Upgrading
  - with 22
- Upgrading your server
  - overview of tasks 45
  - planning 8
  - preparing for 8
  - preparing to undo changes 11
  - prerequisites 8
  - restoring files after a failure 49, 135
- USE\_DTVENV environment variable 72, 122
- USELASTCOMMITTED configuration parameter 72
- USELASTCOMMITTED session environment variable 72
- User-defined data types
  - Preparing Replication 23
- Using SSL/TLS database connections 47
- Utilities
  - dbexport 70, 72, 72
  - dbexport syntax 74
  - dbimport 70, 72
  - dbimport syntax 80
  - dbload 88
  - dbschema 103
  - onload 70, 127
  - onload and onunload 122, 132, 132, 133
  - onunload 70, 124

## V

- Verifying data integrity 58