

HCL Informix 15.0.0

Troubleshooting Informix



Contents

- Chapter 1. Troubleshooting HCL Informix®.....3**
 - Troubleshooting and support..... 11
 - Techniques for troubleshooting problems..... 11
 - Correcting assertion failures..... 13
- Index..... 17**

Chapter 1. Troubleshooting HCL Informix®

Several troubleshooting techniques, tools, and resources are available for resolving problems that you encounter in your HCL Informix® database server environment.

Use this general approach for troubleshooting problems:

1. Determine if the problem is a known software defect...

Check the [Support Portal](#) for Informix®.

- Search for keywords related to your problem to find problem reports, workarounds, patches, or fix packs.

2. If the problem is *not* a known software defect, diagnose the problem...

Check the Informix® online log.

The log file is in `INFORMIXDIR/tmp/online.log` by default, or in the location set by `MSGPATH`. The log contains messages about the database server status and errors.

- Verify that normal operations are logging successfully (startup, checkpoints, shut down).
- Look for messages about problems that occurred (for example, `process aborted abnormal`).
- Run the `finderr` command with the message number to find information about that message and follow the recommendations (for example, restart the database server, run a particular command or utility).

Monitor the database server.

- If the alarm program is configured, examine event alarms and take appropriate actions.

Search Informix® community resources, such as blogs, forums, wikis.

To avoid future problems, review the troubleshooting resources and take these preventive steps:

- Review online logs regularly.
- Run and analyze system statistics regularly.
- Keep up to date with product fix packs.
- Validate changes on a test system before changing the production system.
- Follow backup and restore practices.
- Set up event alarms for error handling.

Frequently asked questions (FAQs) about Informix 15.0.0

These topics provide short answers to some frequently asked questions about installing, upgrading, and administering HCL Informix®. Many answers include links to topics that contain more information.

Questions about installing HCL Informix®

This topic provides short answers to some frequently asked questions about installing Informix®.

- [What are the system requirements and prerequisites for my operating system? on page 4](#)
- [Can I create a server instance after installation is complete? on page 4](#)
- [Can I install client tools in a different directory from the one in which I installed Informix? on page 4](#)
- [Must I install different versions of client tools in different directories? on page 4](#)
- [Which .Net provider should I install? on page 4](#)
- [What are non-root installations and where can I find information about them? on page 4](#)

What are the system requirements and prerequisites for my operating system?

For details about the requirements for different operating systems, see:

the HCL Informix® Machine Notes® for your specific operating system contain information about system requirements and prerequisites.

Can I create a server instance after installation is complete?

You can create a server instance by manually creating or editing the required configuration files. However, it is faster and much easier to create and configure an instance during installation. See *Installing Informix and client products quickly with defaults (UNIX and Linux)* or *Installing with the GUI typical setup (Windows™)*.

Can I install client tools in a different directory from the one in which I installed Informix®?

Yes, if you set the INFORMIXDIR environment variable to point to the directory in which you installed HCL Informix® Client Software Development Kit, or the other client tools. To access the Informix® database, the INFORMIXDIR, INFORMIXSERVER, INFORMIXSQLHOSTS, and PATH environment variables must be set on the client application. Depending on your application, you might need to set more variables. For information about the environment variables that you might need to set client applications, see information about Informix® environment variables in the *Installation Guide*.

Must I install different versions of client tools in different directories?

If your applications are built with different versions of HCL Informix® Client Software Development Kit, you should install the separate versions in separate directories.

Which .Net provider should I install?

Starting with HCL Informix® Version 11.70.xC5, use the IBM® DB2® .NET Provider with your Informix® installation. The IBM® DB2® .NET Provider is part of the HCL Data Server Driver.

What are non-root installations and where can I find information about them?

A non-root installation on UNIX™ or does not require system administration privileges or user and group informix accounts. This can simplify the installation and deployment of Informix® as an embedded solution, but the absence of root privileges limits the availability of product features.

Questions about upgrading HCL Informix®

This topic provides short answers to some frequently asked questions about upgrading to the new version of HCL Informix®.

- [What can I do to avoid problems when upgrading to the new version of HCL Informix? on page 5](#)
- [What data should I collect when upgrading to a new version of Informix? on page 5](#)
- [Must I remove in-place alter operations before upgrading or reverting between product versions? on page 5](#)
- [Can I upgrade HDR clusters without downtime? on page 5](#)
- [Must I use the new onconfig.std file after upgrading? on page 6](#)
- [What should I do if my queries are slower after an upgrade? on page 6](#)

What can I do to avoid problems when upgrading to the new version of HCL Informix®?

Plan for upgrading to the new version and complete all crucial steps. To avoid problems when upgrading to a new version of Informix®:

1. Read the Release Notes® and Machine Notes® before you begin upgrading the database server.
2. Follow the guidelines in the [Preparing for migration on page](#) and [Migrating to Version 15.0.0 on page](#) topics in the *HCL Informix® Migration Guide*.



Important: Do not cut corners to try to save time. Not doing adequate planning can cause problems later.

What data should I collect when upgrading to a new version of Informix®?

Before you upgrade to a newer version of Informix®, gather diagnostic information, especially if you have large, complex applications. Diagnostic information is useful to verify database server behavior after upgrading.

For information about the data that you should collect, see the [Preparing for migration on page](#) and [Pre-migration checklist of diagnostic information on page](#) topics.

Must I remove in-place alter operations before upgrading or reverting between product versions?

Removing outstanding in-place alter operations is not required before upgrading, but it is a best practice.

If you upgraded, ran in-place alter operations in the new version, and must revert, you must ensure that the outstanding in-place alter operations are resolved. If the reversion process lists all of the tables that have outstanding in-place alter operations if reversion does not complete successfully. You must resolve in-place alter operations on each of the tables in the list before you can revert to the older database server.

For more information, see [Resolve outstanding in-place alter operations on page](#) .

Can I upgrade HDR clusters without downtime?

You can upgrade a high-availability cluster online by following one of several approaches for transforming the primary server and one of the secondary servers, either the remote standalone (RS) secondary server or the HDR secondary server, to

form a pair of Enterprise Replication (ER) servers. ER supports this transformation because it allows replication between heterogeneous server versions.

Must I use the new `onconfig.std` file after upgrading?

Each version of the database server includes a new `onconfig.std` file. The file contains new configuration parameters and might contain new default values for some existing configuration parameters. In addition, some configuration parameters used with earlier versions of the database server might not be used with the new version.

After upgrading, use the new `onconfig.std` file. You can customize it as necessary to match the configuration of the prior version of your database. Do not use the old `onconfig.std` file with the new version of the server.

For information about the new and changed configuration parameters in each new release, see [Configuration parameter changes by version on page](#) .

What should I do if my queries are slower after an upgrade?

If your queries are slower after an upgrade, find out what changed that affects your configuration and adjust your configuration and queries as necessary. For details, see [Tune the new version for performance and adjust queries on page](#) .

Questions about administering and using HCL Informix®

These topics provide short answers to some frequently asked questions about administering and using HCL Informix®.

Starting the database server, user information, and privileges:

- [Where can I find information about what to do if the server does not start? on page 7](#)
- [What can I do if I get a message that the password is incorrect or user informix is not known on the database server? on page 7](#)
- [How can I tell which users have DBA privileges on a database? on page 7](#)

Configuration:

- [Where can I find information about the `oncfg` file? on page 8](#)
- [How can I create a customized `onconfig` file? on page 8](#)
- [How much temporary space is needed for `oncheck` sorting? on page 8](#)

Backup and restore:

- [Where can I find information about ON-Bar return codes? on page 8](#)
- [Where can I find information about solving errors that might occur during backup and restore? on page 8](#)
- [Can ON-Bar start continuous log backup even if all logical logs are 100 percent full? on page 8](#)

Assertion failures:

- [How do I correct an assertion failure? on page 8](#)
- [Can I prevent assertion failures? on page 9](#)

Other administrative tasks and questions:

- [How can I plan responses to severity 4 and 5 event alarms? on page 9](#)
- [What can I compress? on page 9](#)
- [How do I estimate compression ratios? on page 9](#)
- [What can I do to prevent automatic update statistics \(AUS\) from running? on page 9](#)
- [Must I run update statistics manually to set the level of statistics before using automatic update statistics \(AUS\)? on page 10](#)
- [What is the purpose of star-join directives and where can I find information about them? on page 10](#)
- [How do I get the onpladm utility to start jobs in the background on the Windows operating system? on page 10](#)
- [What are the benefits of and what are some examples of defragmenting partitions? on page 10](#)
- [How do you use fragment-level statistics? on page 10](#)

Where can I find information about what to do if the server does not start?

If an **oninit** command encounters an error, the database server returns an error message and a return code. The message text describes why the database server did not start. For a list of **oninit** return codes, the text of the messages that the server returns, and the action you can take to solve each problem, see the [Return codes for the oninit utility on page](#) topic in the *HCL® Informix® Administrator's Reference*.

For example, the message for **oninit** return code 175 tells you that the server failed to initialize the root dbspace. If you go to the [Return codes for the oninit utility on page](#) topic and scroll to return code 175, you learn that you should check the root dbspace related parameters in the server configuration (**onconfig**) file to make sure that the path for the root dbspace is valid.

What can I do if I get a message that the password is incorrect or user informix is not known on the database server?

If you try to connect to Informix® using DB-Access on a Windows™ operating system, and you receive error 951, you can troubleshoot the problem by following the guidelines in this IBM® Technote, which is available from the IBM® Support portal: <http://www.ibm.com/support/docview.wss?uid=swg21200068>

How can I tell which users have DBA privileges on a database?

From DB-Access or your application, run this query:

```
select username,usertype from sysusers;
```

The output shows user names (for example, public and informix) followed by one of the following codes:

- D = DBA privilege
- C = Connect privilege
- R = Resource privilege

For information about the database privileges, see [Database-level privileges on page](#) .

Where can I find information about the oncfg file?

This **oncfg** file contains information about dbspaces, chunks, and logical-logs that are used during a whole system restore. For information about this file, see the *Facts about the ONCFG file* IBM® Technote that is available from the IBM® Support portal: <http://www.ibm.com/support/docview.wss?uid=swg21106660>

If you are looking for information about the `onconfig.std` file, which contains configuration parameters, see [onconfig Portal: Configuration parameters by functional category on page](#) in the *HCL® Informix® Administrator's Reference*.

How can I create a customized onconfig file?

You can customize the default `onconfig.std` file by using the **genoncfg** utility, which is in the `$INFORMIXDIR/bin` directory. For more information, see [The genoncfg Utility on page](#) in the *HCL® Informix® Administrator's Reference*.

How much temporary space is needed for oncheck sorting?

The **oncheck** utility requires sort space when examining an index. The amount of sort space required is the same as the amount that is needed to build the index. For information about calculating the amount of temporary space needed, see [Estimating temporary space for index builds on page](#) in the . If you receive the error "no free disk space for sort," you must estimate the amount of temporary space needed and make that space available.

Where can I find information about ON-Bar return codes?

ON-Bar return codes are accompanied by messages in the ON-Bar activity log. For a complete list of ON-Bar return codes, see the [ON-Bar return codes on page](#) topic in the *HCL® Informix® Backup and Restore Guide*.

Where can I find information about solving errors that might occur during backup and restore?

The Knowledge Collection: Informix® Backup and Restore Utilities IBM® Technote (<http://www.ibm.com/support/docview.wss?uid=swg21404291>) contains links to technotes on backup and restore utilities, such as the ON-Bar and **ontape** utilities.

These technotes contain answers to many common questions, information about solving some errors that might occur, and sample commands. These Technotes supplement the information in the *HCL® Informix® Backup and Restore Guide*.

Can ON-Bar start continuous log backup even if all logical logs are 100 percent full?

The ON-Bar utility can start a continuous backup of the logical logs even when the logical logs are full. However, a better practice is to make a regular backup of the logical logs and set the alarm program (`$INFORMIXDIR/etc/alarmprogram.sh`) to back up each log as soon as the log becomes full. For more information, see [Starting a continuous logical-log file backup on page](#) .

How do I correct an assertion failure?

You can correct some types of assertion failures yourself. For others, you must contact IBM® Software Support. For more information, see [Correcting assertion failures on page 13](#).

Can I prevent assertion failures?

Yes, you can prevent many assertion failures by performing routine administrative tasks and thoroughly testing your applications. However, you should be prepared for assertion failures and other problems. For more information, see [Prevent and prepare for assertion failures on page 15](#).

How can I plan responses to severity 4 and 5 event alarms?

Event alarms have a severity level, which specifies the seriousness of the event on a scale from 1 to 5. A level 5 alarm, which is the most serious event alarm, indicates that the database server failed. For more information, including details on how to respond to server failures, see the [Event alarm IDs on page](#) and [Severity 5 event alarms on page](#) topics in the *HCL® Informix® Administrator's Reference*.

What can I compress?

You can compress and uncompress the following items:

- Row data in tables and fragments and simple large objects in dbspaces. You can also choose to compress or uncompress only row data or only simple large objects in dbspaces. For more information, see [table or fragment arguments: Compress data and optimize storage \(SQL administration API\) on page](#)
- Detached B-tree indexes. For more information, see [index compress repack shrink arguments: Optimize the storage of B-tree indexes \(SQL administration API\) on page](#). To create a new compressed B-tree index, use the COMPRESSED keyword with the CREATE INDEX statement. For more information, see [CREATE INDEX statement on page](#).

How do I estimate compression ratios?

You can estimate the percentage of space that you can save if you compress data in tables or table fragments.

To estimate the compression benefit, run the **admin()** or **task()** function with the **estimate_compression** argument. For more information, including examples, see the topic in the *HCL® Informix® Administrator's Guide*.

If you use the HCL OpenAdmin Tool (OAT) for Informix®, see the help topic on estimating compression.

What can I do to prevent automatic update statistics (AUS) from running?

To prevent AUS from running, disable both the Auto Update Statistics Evaluation task and the Auto Update Statistics Refresh task, as follows:

1. Update the value of the **tk_enable** column of the **ph_task** table to **F** where the value of the **tk_name** column is Auto Update Statistics Evaluation.
2. Update the value of the **tk_enable** column of the **ph_task** table to **F** where the value of the **tk_name** column is Auto Update Statistics Refresh.

For more information and examples of using this method, see [Disabling AUS on page](#) in the .

You can also use the HCL OpenAdmin Tool (OAT) for Informix® to disable the AUS process. See the help topic on configuring and enabling automatic update statistics.

Must I run update statistics manually to set the level of statistics before using automatic update statistics (AUS)?

No. The AUS maintenance system automatically identifies and runs the necessary update statistics. However, if you manually update statistics, AUS does not run the statistics at a lower level, resolution, confidence level, or sampling size. For more information, see [Automatic statistics updating on page](#) in the .

What is the purpose of star-join directives and where can I find information about them?

You can use star-join directives to specify how the optimizer joins tables that have a star schema.

Many data warehouse databases use a star schema, which consists of a fact table and a number of dimensional tables. In a typical star join, the fact table joins with all dimensional tables on a foreign key. For more information, see [Star-Join Directives on page](#) in the *HCL® Informix® Guide to SQL: Syntax*.

How do I get the onpladm utility to start jobs in the background on the Windows™ operating system?

If you want **onpladm** utility jobs to start in the background, set the INTERACTIVE_DESKTOP_OFF environment variable to 1. Otherwise, the utility starts each **onpload** job in a new command window because, by default, the INTERACTIVE_DESKTOP_OFF environment variable is set to 0. For more information, see the topics about using the **onpladm** utility on Windows™ in the .

What are the benefits of and what are some examples of defragmenting partitions?

Because a frequently updated table can become fragmented over time, leading to performance degradation, the ability to merge non-contiguous extents can improve performance. Defragmenting a table or partition brings data rows closer together and avoids partition header page overflow problems.

For more information about and examples of defragmenting partitions, see [defragment argument: Dynamically defragment partition extents \(SQL administration API\) on page](#) and the IBM® developerWorks® article about the defragmenter (<http://www.ibm.com/developerworks/data/library/techarticle/dm-1011informixdefragmenter/index.dital>).

How do you use fragment-level statistics?

Fragment-level statistics affect the way UPDATE STATISTICS MEDIUM and HIGH operations gather data and generate column distributions on fragmented tables. For information about the circumstances under which you might want to use fragment-level statistics and details about setting up your system to use fragment-level statistics, see the IBM® developerWorks® article about taking advantage of fragment-level statistics (<http://www.ibm.com/developerworks/data/library/techarticle/dm-1104fragmentstats/index.dital>).

Also see relevant reference topics such as [Fragment-level statistics on page](#) and [Statistics options of the CREATE TABLE statement on page](#) .

Troubleshooting and support

To isolate and resolve problems with your HCL Informix® products, you can use the troubleshooting and support information, which contains instructions for using the problem-determination resources that are provided with your IBM® products.

Techniques for troubleshooting problems

Troubleshooting is a systematic approach to solving a problem. The goal of troubleshooting is to determine why something does not work as expected and how to resolve the problem. Certain common techniques can help with the task of troubleshooting.

The first step in the troubleshooting process is to describe the problem completely. Problem descriptions help you and the IBM® technical-support representative know where to start to find the cause of the problem. This step includes asking yourself basic questions:

- What are the symptoms of the problem?
- Where does the problem occur?
- When does the problem occur?
- Under which conditions does the problem occur?
- Can the problem be reproduced?

The answers to these questions typically lead to a good description of the problem, which can then lead you to a problem resolution.

What are the symptoms of the problem?

When starting to describe a problem, the most obvious question is “*What is the problem?*” This question might seem straightforward; however, you can break it down into several more-focused questions that create a more descriptive picture of the problem. These questions can include:

- Who, or what, is reporting the problem?
- What are the error codes and messages?
- How does the system fail? For example, is it a loop, hang, crash, performance degradation, or incorrect result?

Where does the problem occur?

Determining where the problem originates is not always easy, but it is one of the most important steps in resolving a problem. Many layers of technology can exist between the reporting and failing components. Networks, disks, and drivers are only a few of the components to consider when you are investigating problems.

The following questions help you to focus on where the problem occurs to isolate the problem layer:

- Is the problem specific to one platform or operating system, or is it common across multiple platforms or operating systems?
- Is the current environment and configuration supported?

- Do all users have the problem?
- (For multi-site installations.) Do all sites have the problem?

If one layer reports the problem, the problem does not necessarily originate in that layer. Part of identifying where a problem originates is understanding the environment in which it exists. Take some time to completely describe the problem environment, including the operating system and version, all corresponding software and versions, and hardware information. Confirm that you are running within an environment that is a supported configuration; many problems can be traced back to incompatible levels of software that are not intended to run together or have not been fully tested together.

When does the problem occur?

Develop a detailed timeline of events leading up to a failure, especially for those cases that are one-time occurrences. You can most easily develop a timeline by working backward: Start at the time an error was reported (as precisely as possible, even down to the millisecond), and work backward through the available logs and information. Typically, you need to look only as far as the first suspicious event that you find in a diagnostic log.

To develop a detailed timeline of events, answer these questions:

- Does the problem happen only at a certain time of day or night?
- How often does the problem happen?
- What sequence of events leads up to the time that the problem is reported?
- Does the problem happen after an environment change, such as upgrading or installing software or hardware?

Responding to these types of questions can give you a frame of reference in which to investigate the problem.

Under which conditions does the problem occur?

Knowing which systems and applications are running at the time that a problem occurs is an important part of troubleshooting. These questions about your environment can help you to identify the root cause of the problem:

- Does the problem always occur when the same task is being performed?
- Does a certain sequence of events need to happen for the problem to occur?
- Do any other applications fail at the same time?

Answering these types of questions can help you explain the environment in which the problem occurs and correlate any dependencies. Remember that just because multiple problems might have occurred around the same time, the problems are not necessarily related.

Can the problem be reproduced?

From a troubleshooting standpoint, the ideal problem is one that can be reproduced. Typically, when a problem can be reproduced you have a larger set of tools or procedures at your disposal to help you investigate. Consequently, problems that you can reproduce are often easier to debug and solve.

However, problems that you can reproduce can have a disadvantage: If the problem is of significant business impact, you do not want it to recur. If possible, re-create the problem in a test or development environment, which typically offers you more flexibility and control during your investigation.

- Can the problem be re-created on a test system?
- Are multiple users or applications encountering the same type of problem?
- Can the problem be re-created by running a single command, a set of commands, or a particular application?

Correcting assertion failures

An assertion failure occurs when the database server cannot continue normal processing and must shut down. You can correct some of the problems that cause assertion failures, such as disk issues. For other problems that cause assertion failures, you must contact HCL Informix® Support.

About this task

If your production server is currently down, contact HCL Informix® Support immediately.

To correct an assertion failure:

1. Look at the online log for the assertion failure message.

The location of the online log is specified by the MSGPATH configuration parameter. Default:

- UNIX™: \$INFORMIXDIR/tmp/online.log
- Windows™: %INFORMIXDIR%\online.log

2. Perform any action mentioned in the assertion failure message. For example, you might need to run an oncheck command or increase the size of the physical or logical log buffer.

The following table lists the common types of assertion failures and what you can do to correct the underlying problems.

Table 1. Common types of assertion failures

Error	Cause	Solutions
bfcheck	Corrupted data page, index, or partition page	<ul style="list-style-type: none"> ◦ Run the oncheck command in the assertion failure message. ◦ If an index is corrupted, drop and recreate the index. ◦ If a partition page is corrupted, you might need to restore the data.
I/O write chunk	Not enough shared memory or a hardware problem	<ul style="list-style-type: none"> ◦ Look at your operating system logs for hardware problems or problems with other software. Sometimes other software can

Error	Cause	Solutions
		trigger an Informix® assertion failure. Fix any hardware or other software problems.
		◦ Free up or increase shared memory for the database server.
No exception handler	Memory corruption	Look at your operating system logs for hardware problems or problems with other software. Fix any hardware or other software problems.
sigkill	The oninit daemon was stopped.	Restart the database server.
segv	◦ Memory corruption	◦ Look at your operating system logs for hardware problems or problems with other software. Fix any hardware or other software problems.
sigbus	◦ An ill-behaved C user-defined routine	◦ If you use C user-defined routines, check your code for problems. Make sure that your user-defined routines follow the guidelines for well-behaved routines.
	◦ A defect in the database server	◦ If you are not running the latest Informix® fix pack, check the latest Fixed and known defects file to see whether this issue is fixed. Upgrade to the latest Informix® fix pack.

- Restart the database server. Some assertion failures can be solved by the database server during startup.
- If the database server does not start or the problem recurs, collect the following files for IBM® Software Support in addition to the standard list of information to gather.

Table 2. Files for support

File	Location
assertion failure files: <code>af.xxx</code> , where xxx is the hexadecimal number associated with the failure	Specified by the DUMPDIR configuration parameter. Default: <ul style="list-style-type: none"> ◦ UNIX™: <code>\$INFORMIXDIR/tmp</code> ◦ Windows™: <code>%INFORMIXDIR%\tmp</code>
<code>sqlhosts</code> file	Specified by the INFORMIXSQLHOSTS environment variable. Default: <ul style="list-style-type: none"> ◦ UNIX™: <code>\$INFORMIXDIR/etc</code> ◦ Windows™: <code>%INFORMIXDIR%\etc</code>
<code>onconfig</code> file	Specified by the ONCONFIG environment variable. Default: <ul style="list-style-type: none"> ◦ UNIX™: <code>\$INFORMIXDIR/etc</code> ◦ Windows™: <code>%INFORMIXDIR%\etc</code>
online log file	Specified by the MSGPATH configuration parameter. Default:

File	Location
	<ul style="list-style-type: none"> ◦ UNIX™: \$INFORMIXDIR/tmp/online.log ◦ Windows™: %INFORMIXDIR%\online.log
shared memory dump files: shmem.xxx, where xxx is the hexadecimal number associated with the failure	<p>Specified by the DUMPPDIR configuration parameter. Default:</p> <ul style="list-style-type: none"> ◦ UNIX™: \$INFORMIXDIR/tmp ◦ Windows™: %INFORMIXDIR%\tmp <p>Generated only if the DUMPSHMEM configuration parameter is set to 1 or 2.</p>

5. Contact HCL Informix® Support.

Prevent and prepare for assertion failures

You can perform routine administrative tasks that can prevent assertion failures or help prepare for recovering from assertion failures.

To prevent assertion failures on Informix®, perform the following tasks regularly:

Table 3. Preventing assertion failures

Prevention task	Purpose
Test your applications thoroughly in a realistic environment before releasing your applications into production.	<p>Applications can trigger assertion failures.</p> <p>Make sure that you have allocated enough resources: for example, physical and logical logs and shared memory.</p>
Check and repair the consistency of your data.	Diagnosing and repairing consistency problems
Run the following oncheck commands to check your system:	quickly can prevent major problems like data loss and assertion failures.
<ul style="list-style-type: none"> • oncheck -cD: Data pages • oncheck -cI: Index pages • oncheck -cr: Reserved pages • oncheck -ce: Chunk-free list • oncheck -cc: System catalog tables • oncheck -pe: Physical information about chunks 	
Upgrade to the latest major and fix pack release.	Assertion failures can be caused by product defects. You can avoid many product defects by upgrading to the most current release.

Table 3. Preventing assertion failures**(continued)**

Prevention task	Purpose
Update statistics. When you change a table or indexing schema, update statistics to update the dictionary cache. By default, statistics are updated automatically by the Scheduler. You can customize automatic statistics updating for your system.	Queries that use stale statistics can use too many resources and cause assertion failures. Stale statistics can also have a negative affect on query performance. Always update statistics after upgrading.

To prepare for assertion failures on Informix®, perform the following tasks regularly:

Table 4. Preparing for assertion failures

Preparation task	Purpose
Back up your data and logical logs. You can use the ON-Bar utility or the ontape utility to back up your data and logical logs. The ON-Bar utility requires a storage manager.	If data loss occurs, you can recover your data and recent transactions.
Back up your database and storage space schema information. Run the following commands to back up your schema information:	In a catastrophic failure, you can recreate your system.
<ul style="list-style-type: none"> • <code>dbschema -ss -d database_name > dbschema_ss_database_name.out:</code> Saves the schema information about the specified database in a file. • <code>dbschema -c > dbschema_c.out:</code> Saves the commands to reproduce storage spaces, chunks, physical logs, and logical logs in a file. 	

Related links

The oncheck utility

Automatic statistics updating

The ph_alert table

Backup and restore

Syntax of the dbschema command

Index

A

- Assertion failures
 - Assertion failure
 - files 13
 - correcting 13
 - preventing 15

F

- FAQ 3
- Frequently asked questions 3

T

- Troubleshoot
 - assertion failures 13, 15
- Troubleshooting
 - FAQ 3
 - frequently asked questions 3