

HCL Informix 14.10 - Backup and Restore Guide



Contents

Chapter 1. Backup and Restore Guide.....	3
Overview of backup and restore.....	3
Backup and restore concepts.....	3
Plan for backup and restore.....	13
ON-Bar backup and restore system.....	20
Overview of the ON-Bar backup and restore system.....	20
Configure the storage manager and ON-Bar.....	27
Back up with ON-Bar.....	43
Restore data with ON-Bar.....	64
External backup and restore.....	92
Customize and maintain ON-Bar.....	102
ON-Bar catalog tables.....	116
ON-Bar messages and return codes.....	122
ontape backup and restore system.....	132
Configure ontape.....	132
Back up with ontape.....	139
Restore with ontape.....	157
Perform an external backup and restore.....	173
Backup and restore a Remote Secondary Server(RSS).....	179
Backup and Restore Guide.....	
Backup and Restore Guide.....	
Backup and Restore Guide.....	
Informix® Primary Storage Manager.....	180
HCL Informix® Primary Storage Manager.....	181
archecker table level restore utility.....	203
archecker table level restore utility.....	203
Backup and restore configuration parameter reference.....	220
Backup and restore configuration parameters.....	220
Cloud Backup.....	267
Back up to Amazon Simple Storage Service using ON-Bar and the PSM.....	267
Back up to Softlayer using ON-Bar and the PSM.....	269
Appendixes.....	271
Troubleshooting some backup and restore errors.....	271
Migrate data, servers, and tools.....	274
GLS support.....	276
Index.....	278

Chapter 1. Backup and Restore Guide

These topics describe how to use the HCL® Informix® ON-Bar and ontape utilities to back up and restore database server data. These utilities enable you to recover your databases after data is lost or becomes corrupted due to hardware or software failure or accident.

These topics are of interest to the following users:

- Database administrators
- System administrators
- Backup operators
- Technical support personnel

These topics are written with the assumption that you have the following background:

- Some experience with storage managers, which are applications that manage the storage devices and media that contain backups
- A working knowledge of your computer, your operating system, and the utilities that your operating system provides
- Some experience working with relational databases or exposure to database concepts
- Some experience with database server administration, operating-system administration, or network administration

Overview of backup and restore

These topics provide an overview of backup and restore concepts. They also provide information about planning for backup and restore operations.

Backup and restore concepts

HCL Informix® provides two utilities for backing up and restoring database server data. Both utilities back up and restore storage spaces and logical logs. However, they support different features and it is important to know the differences. These topics explain basic backup and restore concepts for HCL® Informix® database servers and compares the ON-Bar and ontape utilities.

ON-Bar backs up and restores storage spaces (dbspaces) and logical file, by using a storage manager, whereas ontape does not use a storage manager.

Recovery system

A *recovery system*, which includes backup and restore systems, enables you to back up your database server data and later restore it if your current data becomes corrupted or inaccessible.

The causes of data corruption or loss can range from a program error to a disk failure to a disaster that damages the entire facility. A recovery system enables you to recover data that you already lost due to such mishaps.

Backup systems

A *backup* is a copy of one or more *dbspaces* (also called *storage spaces*) and logical logs that the database server maintains. You can also back up *blobspaces* and *sbspaces*.


The backup copy is typically written to a *secondary storage* medium such as disk or magnetic tape. Store the media offline and keep a copy off site if possible.

The backup copy is typically written to a *secondary storage* medium such as disk, magnetic tape, or optical disk. Store the media offline and keep a copy off site if possible.

! **Important:** Database backups do not replace ordinary operating-system backups, which back up files other than HCL® Informix® database files.

The following figure illustrates the basic concept of a database backup.

Figure 1. A backup of database server data



You do not always have to back up all the storage spaces. If some tables change daily but others rarely change, it is inefficient to back up the storage spaces that contain the unchanged tables every time that you back up the database server. You need to plan your backup schedule carefully to avoid long delays for backing up or restoring data.

Backup levels

To provide flexibility, the ON-Bar and ontape utilities support three backup levels.

Level 0

Level 0 backs up all used pages that contain data for the specified storage spaces.

You need all these pages to restore the database to the state that it was in at the time that you made the backup.

Level-0 backups can be time-consuming because ON-Bar writes all the disk pages to back up media. Level-1 and level-2 backups might take almost as much time as a level-0 backup because the database server must scan all the data to determine what has changed since the last backup. It takes less time to restore data from level-0, level-1, and level-2 backups than from level-0 backups and a long series of logical-log backups.

Level 1

Level 1 backs up only data that has changed since the last level-0 backup of the specified storage spaces.

All changed table and index pages are backed up, including those pages with deleted data. The data that is copied to the backup reflects the state of the changed data at the time that the level-1 backup began.

A level-1 backup takes less space and might take less time than a level-0 backup because only data that changed since the last level-0 backup is copied to the storage manager.

Level 2

Level 2 backs up only data that has changed since the last level-1 backup of the specified storage spaces.

A level-2 backup contains a copy of every table and index page in a storage space that has changed since the last level-1 backup.

A level-2 backup takes less space and might take less time than a level-1 backup because only data that changed since the last level-1 backup is copied to the storage manager.



Important: If disks and other media are destroyed and need to be replaced, you need at least a level-0 backup of all storage spaces and relevant logical logs to restore data completely on the replacement hardware.

Related reference

[Schedule backups on page 15](#)

Logical-log backup

A *logical-log backup* is a copy to disk or tape of all full logical-log files. The logical-log files store a record of database server activity that occurs between backups.

To free full logical-log files, back them up. The database server reuses the freed logical-log files for recording new transactions. For a complete description of the logical log, see your *HCL® Informix® Administrator's Guide*.



Restriction: Even if you do not specify logging for databases or tables, you need to back up the logical logs because they contain administrative information such as checkpoint records and additions and deletions of chunks. When you back up these logical-log files, you can do warm restores even when you do not use logging for any of your databases.

Manual and continuous logical-log backups

You can manually back up logical logs or you can enable continuous logical-log backup.

A *manual logical-log backup* backs up all the full logical-log files and stops at the current logical-log file. You must monitor your logical logs carefully and start logical-log backups as needed.

To find out if a logical-log file is ready to be backed up, check the flags field of `onstat -l`. After the logical-log file is marked as backed up, it can be reused. When the flags field displays any of the following values, the logical-log file is ready to be backed up:

```
U-----
U-----L
```

The value `U` means that the logical-log file is used. The value `L` means that the last checkpoint occurred when the indicated logical-log file was current. The value `C` indicates the current log. If `B` appears in the third column, the logical-log file is already backed up and can be reused.

```
U-B---L
```

The flag values `U---C-L` or `U---C--` represent the current logical log. While you are allowed to back up the current logical log, doing so forces a log switch that wastes logical-log space. Wait until a logical-log file fills before you back it up.

If you turn on *continuous logical-log backup*, the database server backs up each logical log automatically when it becomes full. If you turn off continuous logical-log backup, the logical-log files continue to fill. If all logical logs are filled, the database server hangs until the logs are backed up. You can start continuous logical log backups by setting the `ALARMPROGRAM` configuration parameter in the `onconfig` file or by running an `ON-Bar` or `ontape` command.

Log salvage

When the database server is offline, you can perform a special logical-log backup, called a *log salvage*. In a log salvage, the database server accesses the log files directly from disk. The log salvage backs up any logical logs that have not yet been backed up and are not corrupted or destroyed.

The log salvage enables you to recover all of your data up to the last available and uncorrupted logical-log file and the last complete transaction.

Save logical-log backups

You should perform frequent logical-log backups and then save the logical-log backups from at least the last two level-0 backups so that you can use them to complete a restore.

Perform frequent logical-log backups for the following reasons:

- To free full logical-log files
- To minimize data loss if a disk that contains logical logs fails
- To ensure that restores contain consistent and the latest transactions

You should save the logical-log backups from the last two level-0 backups because if a level-0 backup is inaccessible or unusable, you can restore data from an older backup. If any of the logical-log backups are also inaccessible or unusable, however, you cannot roll forward the transactions from those logical-log files or from any subsequent logical-log files.

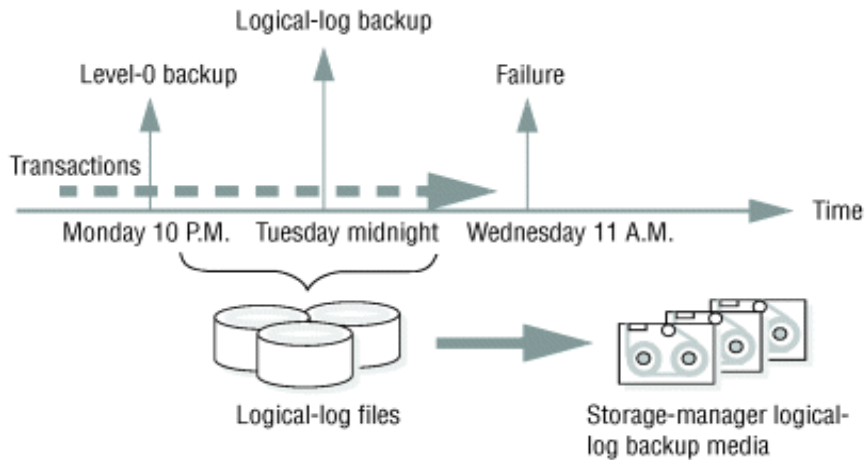


Important: You lose transactions in logical-log files that are not backed up or salvaged.

To illustrate, as the following figure shows, suppose you perform a level-0 backup on Monday at 10 p.m. and then back up the logical logs on Tuesday at midnight. On Wednesday at 11 a.m., you suffer a mishap that destroys your databases. You would be unable to restore the transactions that occurred between midnight on Tuesday and 11 a.m. on Wednesday unless you had continuous logical-log backup setup.

If the disks that contain the storage spaces with the logical logs are damaged, the transactions after midnight on Tuesday might be lost. To restore these transactions from the last logical-log backup, try to salvage the logical logs before you repair or replace the bad disk and then perform a cold restore.

Figure 2. Storage space and logical-log backups



Restore systems

A *restore* recreates database server data from backed-up storage spaces and logical-log files.

A restore recreates database server data that has become inaccessible because of any of the following conditions:

- You need to replace a failed disk that contains database server data.
- A logic error in a program has corrupted a database.
- You need to move your database server data to a new computer.
- A user accidentally corrupted or destroyed data.

To restore data up to the time of the failure, you must have at least one level-0 backup of each of your storage spaces from before the failure and the logical-log files that contain all transactions since these backups.


Physical and logical restores

ON-Bar and ontape restore database server data in two phases. The first phase is the *physical restore*, which restores data from backups of all or selected storage spaces. The second phase is the *logical restore*, which restores transactions from the logical-log backups.

Physical restore

During a physical restore, ON-Bar or ontape restores the data from the most recent level-0, level-1, and level-2 backups. When you suffer a disk failure, you can restore to a new disk only those storage spaces with chunks that resided on the failed disk. The following figure illustrates a physical restore.

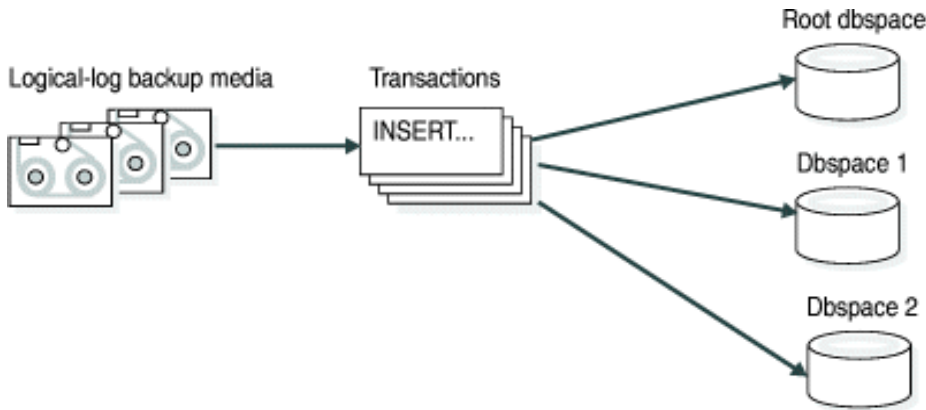
Figure 3. Physical restore



Logical restore

As the following figure shows, the database server *replays* the logical logs to reapply any database transactions that occurred after the last backup. The logical restore applies only to the physically restored storage spaces.

Figure 4. Logical restore



The database server automatically knows which logical logs to restore.

For more information, see [Restore data with ON-Bar on page 64](#) and [Restore with ontape on page 157](#).

Warm, cold, and mixed restores

When you restore data, you must decide whether to do so while the database server is in quiescent, online, or offline mode. The type of restore depends on which of these operating modes the server is in.

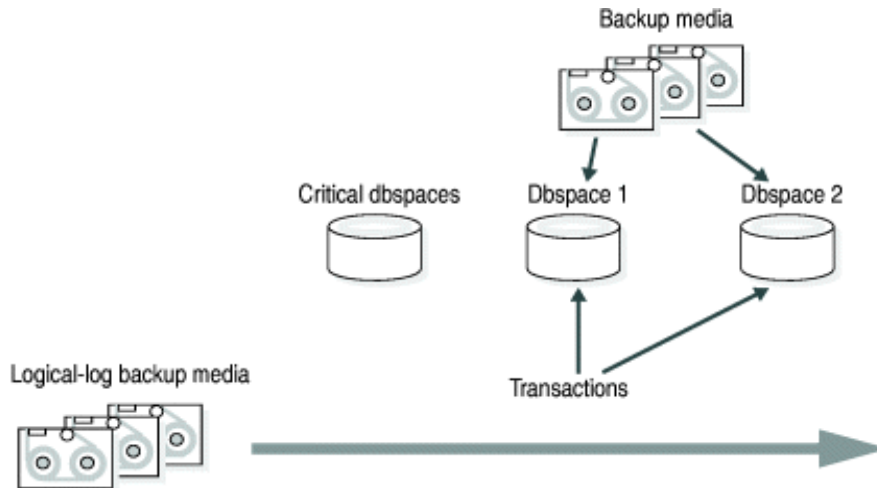
The types of restores are as follows:

- If you restore noncritical dbspaces while the database server is online or quiescent, that process is called a *warm restore*.
- When HCL Informix® is offline, you can perform only a *cold restore*.
- A *mixed restore* is a cold restore of some storage spaces followed by a warm restore of the remaining storage spaces.

Warm restore

As the following figure shows, a warm restore restores noncritical storage spaces. A warm restore consists of one or more physical restores, a logical-log backup, and a logical restore.

Figure 5. Warm restore

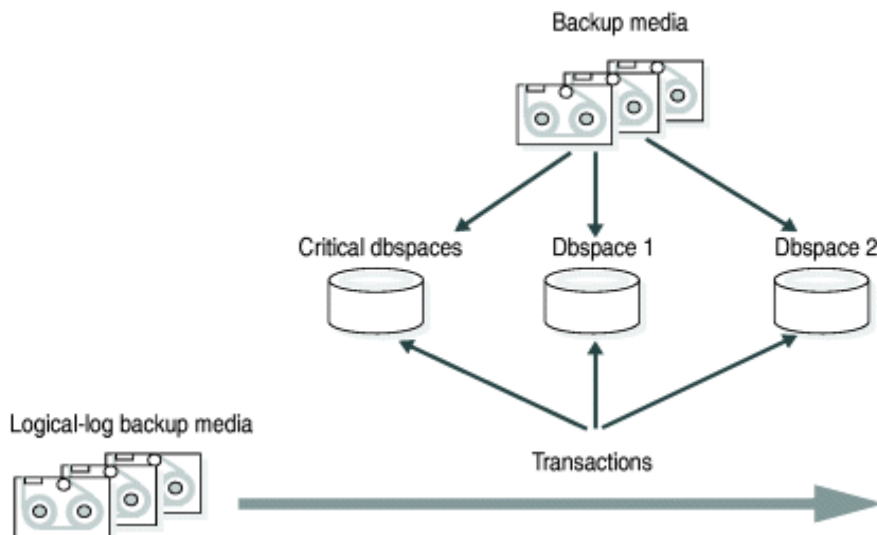


You cannot perform more than one simultaneous warm restore.

Cold restore

As the following figure shows, a cold restore salvages the logical logs, and restores the critical dbspaces (root dbspace and the dbspaces that contain the physical log and logical-log files), other storage spaces, and the logical logs.

Figure 6. Cold restore





You can perform a cold restore onto a computer that is not identical to the one on which the backup was performed by giving any chunk a new path name and offset during the restore.

When restoring a whole-system backup, it is not necessary to restore the logical logs. A whole-system backup contains a snapshot of the entire instance at the moment the backup was performed, which is logically consistent across all dbspaces.

When restoring a standard backup, you must restore the logical logs by performing a logical restore.

A cold restore starts by physically restoring all critical storage spaces, then the noncritical storage spaces, and finally the logical logs. The database server goes into recovery mode after the reserved pages of the root dbspace are restored. When the logical restore is complete, the database server goes into quiescent mode. Use the onmode command to bring the database server online.

 **Tip:** If you mirror the critical dbspaces, you are less likely to have to perform a cold restore after a disk failure because the database server can use the mirrored storage space. If you mirror the logical-log spaces, you are more likely to be able to salvage logical-log data if one or more disks fail.

 **Required:** Cold restores are required for Enterprise Replication servers before resuming replication.

Mixed restores

A mixed restore makes the critical data available sooner, however, the complete restore takes longer because the logical logs are restored and replayed several times, once for the initial cold restore and once for each subsequent warm restore.

The initial set of storage spaces you restore in the cold restore must include all critical storage spaces in the server. To the extent that you do not restore all storage spaces during the initial cold restore and avoid the time necessary to restore them, you can bring the server online faster than if you were to perform a cold restore of the entire server. You can then restore the remaining storage spaces in one or more warm restores.

The storage spaces that you do not restore during the cold restore are not available until after you restore them during a warm restore, although they might not have been damaged by the failure.

Related reference

[Determine failure severity on page 14](#)

Continuous log restore

A *continuous log restore* keeps a second system available to replace the primary system if the primary system for restoring logs fails.

Normal log restore restores all of the available log file backups and applies the log records. After the last available log is restored and applied, the log restore finishes. Transactions that are still open are rolled back in the transaction cleanup phase, then the server is brought into quiescent mode. After the server is quiesced, no more logical logs can be restored.

With continuous log restore, instead of transaction clean up the server is put into log restore suspended state after the last available log is restored. The restore client (ontape or ON-Bar) exits and returns control to you. With the server in this state, you can start another logical restore after additional logical logs become available. As long as you start each log restore as a continuous log restore, you can continue this cycle indefinitely.

One use of continuous log restore is to keep a second system available in case the primary system fails. You can restore logical logs backed up on the primary system on the secondary system as they become available. If the primary system fails, you can restore remaining available logical logs on the secondary system and bring that secondary system online as the new primary system.

Continuous log restore requires much less network bandwidth than High-Availability Data Replication (HDR) and enterprise data replication (ER). Continuous log restore is more flexible than HDR and ER because you can start continuous log restore at any time. As a result, continuous log restore is more robust than HDR or ER in unpredictable circumstances, such as intermittent network availability.

For more information, see [Configuring a continuous log restore by using ON-Bar on page 76](#) and [Configuring continuous log restore with ontape on page 167](#).

Related information

[Configuring a continuous log restore by using ON-Bar on page 76](#)

[Configuring continuous log restore with ontape on page 167](#)

Comparison of the ON-Bar and ontape utilities

This topic contains information to help you compare the ON-Bar and ontape utilities, so you can determine when to use each utility.

ON-Bar

Backs up and restores storage spaces (dbspaces) and logical files, by using a storage manager to track backups and storage media. Use this utility when you need to:

- Select specific storage spaces
- Back up to a specific point in time
- Perform separate physical and logical restores
- Back up and restore different storage spaces in parallel
- Use multiple tape drives concurrently for backups and restores
- Perform imported restores
- Perform external backups and restores

ontape

Logs, backs up, and restores data, and enables you to change the logging status of a database. It does not use a storage manager. Use this utility when you need to:

- Back up and restore data without a storage manager
- Back up without selecting storage spaces
- Change the logging mode for databases



Important: The backup that ontape and ON-Bar produce are not compatible. You cannot create a backup with ontape and restore it with ON-Bar, or vice versa.

The following table compares ON-Bar and ontape.

Table 1. Differences between ON-Bar and ontape

Can the utility	ON-Bar	ontape
Use a storage manager to track backups and storage media?	yes	no
Back up all database server data?	yes	yes
Back up selected storage spaces?	yes	no
Back up logical-log files?	yes	yes
Perform continuous logical-log backups?	yes	yes
Perform continuous logical-log restore?	yes	yes
Back up while the database server is online?	yes	yes
Back up while the database server is in quiescent mode?	yes	yes
Restore all database server data?	yes	yes
Restore selected storage spaces?	yes	yes
Back up and restore storage spaces serially?	yes	yes
Perform cold restores with the database server offline?	yes	yes
Initialize high availability data replication?	yes	yes
Restore data to a specific point in time?	yes	no
Perform separate physical and logical restores?	yes	yes
Back up and restore different storage spaces in parallel?	yes	no
Use multiple tape drives concurrently for backups and restores?	yes	no
Restart a restore?	yes	no
Rename a chunk path name or device during a cold restore?	yes	yes
Perform imported restores?	yes	yes
Perform external backups and restores?	yes	yes
Monitor performance?	yes	no
Change logging mode for databases?	no	yes
Transform data with external programs?	yes	yes

Table 1. Differences between ON-Bar and ontape (continued)

Can the utility	ON-Bar	ontape
Back up to or restore from cloud storage?	no	yes
Encrypt or decrypt a storage space during a restore?	yes	yes

Additional differences:

- Emergency boot files and sysutils database

The ontape utility does not use the **sysutils** database or the emergency boot files.

- Simultaneous sessions

ON-Bar, with HCL Informix® Primary Storage Manager, supports simultaneous sessions.

ON-Bar, with ISM, supports up to four simultaneous sessions per ISM instance. The ontape utility supports two simultaneous sessions, one for physical backup or restore, and one for log backup.

- Device support and storage management

The ontape utility supports remote backup devices on other hosts.

ON-Bar, with the Informix® Primary Storage Manager, supports the export of backup generations into specified directories and devices.

ON-Bar supports different sets of tape drives on various hardware platforms.

You can also use ON-Bar with the IBM Spectrum Protect or third-party storage managers to obtain device support and storage management.

- Changing the logging mode of a database

You cannot change the logging mode for ON-Bar; however you can use the ondblog utility to do this task when using ON-Bar.

You can also use the SQL administration API alternative, ALTER LOGMODE to change the logging mode.

For details about each utility, see [Back up with ON-Bar on page 43](#) and [Back up with ontape on page 139](#).

Related reference

[Configure ontape on page 132](#)

Plan for backup and restore

These topics describe the planning for backup and restore, for example by planning your recovery strategy and backup system.

Plan a recovery strategy

Before you use ON-Bar or ontape, plan your recovery goals.

Types of data loss

The first step in planning a recovery strategy is to determine how much data loss, if any, is acceptable.

The following types of data loss can occur:

- Deletion of the following:
 - Rows, columns, tables, or databases
 - Chunks, storage spaces, or logical logs
- Data corruption or incorrect data created
- Hardware failure (such as a disk that contains chunk files fails or a backup tape that wears out)
- Database server failure
- Natural disaster

Determine failure severity

After you determine your recovery goals, create your recovery plan. The plan should include recovery goals for multiple levels of failure.

The following table shows recovery plans for failures with amounts of lost data.

Table 2. Sample recovery plans

Failure severity	Data loss	Suggested recovery plan
Small	Noncritical data is lost.	Restore of the data can wait until a nonpeak time. Use a warm restore.
Medium	The data that is lost is critical for your business but does not reside in a critical dbspace.	Perform a warm restore of this data as soon as possible.
Large	Critical dbspaces are lost.	Use a mixed restore to restore the critical data right away and a warm restore to restore noncritical data during off-peak hours.
Disaster	All data is lost.	Perform a cold or mixed restore as soon as possible.

Related information

[Warm, cold, and mixed restores on page 8](#)

Data use determines your backup schedule

After you develop your recovery plan, create a backup plan based on how you use your data.

How you use the data determines how you plan your backup schedule, as follows:

- Data usage

How do users use the data?

- Critical dbspaces (root dbspace and dbspaces that contain the physical log and at least one logical-log file)
- Critical business application data
- Long-term data storage for legal or record-keeping reasons
- Data sharing among groups
- Test data

- Transaction Time

How much transaction time can be lost? Also, how long might it take to re-enter lost transactions manually? For example, can you afford to re-enter all transactions that occurred over the past three hours?

- Quantity and Distribution

How much data can you afford to lose? For example, you lost one fourth of your customer profiles, or you lost the Midwest regional sales figures but the West Coast figures are intact.

Ask the following questions to assist in deciding how often and when you want to back up the data:

- Does your business have downtime where the system can be restored?
- If your system is 24x7 (no downtime), is there a nonpeak time where a restore could occur?
- If a restore must occur during a peak period, how critical is the time?
- Which data can you restore with the database server online (warm restore)? Which data must be restored offline (cold restore)?
- How many storage devices are available to back up and restore the data?

Schedule backups

Your recovery strategy should include a schedule of backups. Tailor your backup plan to the requirements of your system. The more often the data changes and the more important it is, the more frequently you need to back it up.

Your backup plan should also specify the backup level.

The following table shows a sample backup plan for a small or medium-sized system.

Table 3. Sample backup plan

Backup level	Backup schedule
Complete backup (level-0)	Saturday at 6 p.m.
Incremental backup (level-1)	Tuesday and Thursday at 6 p.m.

Table 3. Sample backup plan (continued)

Backup level	Backup schedule
Incremental backup (level-2)	Daily at 6 p.m.
Level-0 backup of storage spaces that are updated frequently	Hourly



Important: Perform a level-0 backup after you change the physical schema, such as adding a chunk to a storage space. (See [Preparing to back up data on page 44.](#))

Related information

[Backup levels on page 4](#)

Security requirements for label-based access control

For label-based access control (LBAC), the person who runs ON-Bar or ontape does not require an exemption to security policies or an additional privilege to back up or restore data.

LBAC protection remains intact after you restore data with ON-Bar or ontape.

Plan a backup system for a production database server

To plan for adequate backup protection for your data, analyze your database server configuration and activity and the types of backup media available at your installation.


Also, consider your budget for storage media, disks, computers and controllers, and the size of your network.

Actions after which to perform a level-0 back up

You must perform a level-0 backup of, at minimum, the root dbspace and the modified storage spaces after you perform any of the following actions:

- Add or drop mirroring.
- Move, drop, or resize a logical-log file.
- Change the size or location of the physical log.
- Change your storage-manager configuration.
- Add, move, or drop a dbspace.
- Add, move, or drop a chunk to any type of storage space.
- Add, move, or drop a blobspace or sbospace.

For example, if you add a new dbspace **dbs1**, you see a warning in the message log that asks you to perform a level-0 backup of the root dbspace and the new dbspace. If you attempt an incremental backup of the root dbspace or the new dbspace instead, ON-Bar automatically performs a level-0 backup of the new dbspace.

 **Tip:** Although you no longer need to back up immediately after adding a log file, your next backup should be level-0 because the data structures have changed.

If you create a storage space with the same name as a deleted storage space, perform a level-0 backup twice:

1. Back up the root dbspace after you drop the storage space and before you create the storage space with the same name.
2. After you create the storage space, back up the root dbspace and the new storage space.

Actions before which to perform a level-0 back up

You must perform a level-0 backup of the modified storage spaces before you perform any of the following actions:

- Convert a nonlogging database to a logging database.
- Before you alter a RAW table to type STANDARD. This backup ensures that the unlogged data is restorable before you switch to a logging table type.

Related reference

[onbar -b syntax: Backing up on page 46](#)

Evaluate hardware and memory resources

When planning your backup system, evaluate your hardware and memory resources.

Evaluate the following database server and hardware configuration elements to determine which storage manager and storage devices to use:

- The number of I/O virtual processors
- The amount of memory available and the distribution of processor activity

Also consider temporary disk space needed for backup and restore. The database server uses temporary disk space to store the before images of data that are overwritten while backups are occurring and overflow from query processing that occurs in memory.

When preparing to back up data, make sure that you correctly set the DBSPACETEMP environment variable or parameter to specify dbspaces with enough space for your needs. If there is not enough room in the specified dbspaces, the backup will fail, root dbspace will be used, or the backup will fail after filling the root dbspace.

Evaluate backup and restore time

Several factors, including database server configuration and the size of your database, affect the amount of time that the system needs to back up and restore data.

How long your backup or restore takes depends on the following factors:

- The speed of disks or tape devices

The faster the storage devices, the faster the backup or restore time.

- The number of incremental backups that you want to restore if a disk or system failure requires you to rebuild the database

Incremental backups use less storage space than full backups and also reduce restore time.

- The size and number of storage spaces in the database

Backups: Many small storage spaces take slightly longer to back up than a few large storage spaces of the same total size.

Restores: A restore usually takes as long to recover the largest storage space and the logical logs.

- Whether storage spaces are mirrored

If storage spaces are mirrored, you reduce the chance of having to restore damaged or corrupted data. You can restore the mirror at nonpeak time with the database server online.

- The length of time users are interrupted during backups and restores

If you perform backups and warm restores while the database server is online, users can continue their work but might notice a slower response. If you perform backups and warm restores with the database server in quiescent mode, users must exit the database server. If you perform a cold restore with the database server offline, the database server is unavailable to users, so the faster the restore, the better. An external backup and restore eliminates system downtime.

- The backup schedule

Not all storage spaces need to be included in each backup or restore session. Schedule backups so that you can back up more often the storage spaces that change rapidly than those storage spaces that seldom or never change. Be sure to back up each storage space at level-0 at least once.

- The layout of the tables across the dbspaces and the layout of dbspaces across the disks

When you design your database server schema, organize the data so that you can restore important information quickly. For example, you isolate critical and frequently used data in a small set of storage spaces on the fastest disks. You also can fragment large tables across dbspaces to balance I/O and maximize throughput across multiple disks. For more information, see your .

- The database server and system workload

The greater the workload on the database server or system, the longer the backup or restore time.

- The values of backup and restore configuration parameters

For example, the number and size of data buffers that ON-Bar uses to exchange data with the database server can affect performance. Use the `BAR_NB_XPORT_COUNT` and `BAR_XFER_BUF_SIZE` configuration parameters to control the number and size of data buffers.

Evaluate logging and transaction activity

When planning your backup system, also consider logging and transaction activity.

The following database server usage requirements affect your decisions about the storage manager and storage devices:

- The amount and rate of transaction activity that you expect
- The number and size of logical logs

If you need to restore data from a database server with little transaction activity, define many small logical logs. You are less likely to lose data because of infrequent logical-log backups.

- How fast the logical-log files fill

Back up log files before they fill so that the database server does not hang.

- Database and table logging modes

When you use many nonlogging databases or tables, logical-log backups might become less frequent.

Compress row data

Compressing row data can make backing up and restoring data more efficient.

Compressing row data before backing it up can improve the speed of backing up and restoring and requires less backup media. A smaller size of data results in the following advantages over uncompressed data during backup and restore:


- Backing up is quicker.
- Restoring is quicker.
- The logical logs are smaller.
- The backup image is smaller.

Using an external compression utility to compress a backup image of compressed row data might not reduce the size of the backup image, because already compressed data usually cannot be further compressed. In some cases, the size of the backup image of compressed row data might be larger than the size of the backup image that was compressed by an external utility.

Transform data with external programs

You can use external programs as filter plug-ins to transform data to a different format before a backup and transform it back after the restore.

To compress or transform data, use the `BACKUP_FILTER` and `RESTORE_FILTER` configuration parameters to call external programs.

 **Tip:** If you compress row data before backing it up, compressing the backup image with an external utility might not result in a smaller backup image.

The filter can be owned by anyone, but cannot have write access to non-privileged users. Permission on the filters is the same as that of permission on any other executable file that is called by the HCL Informix® server or Informix® utilities.

ON-Bar backup and restore system

Related reference

[Backup and restore configuration parameters on page 220](#)

Overview of the ON-Bar backup and restore system

ON-Bar consists of various components and it works with a storage manager to back up and restore data.

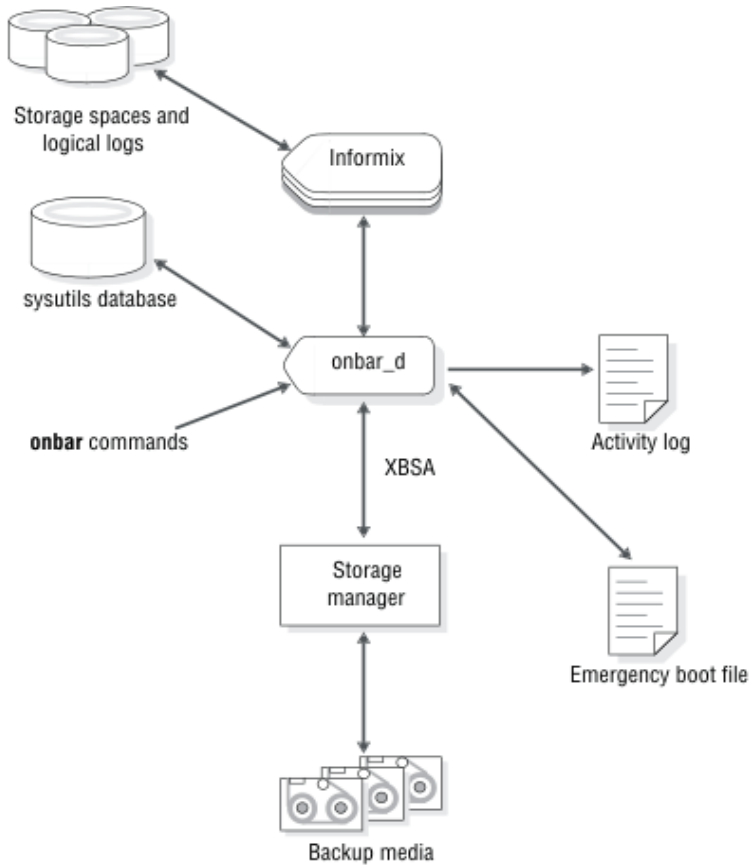
ON-Bar components

ON-Bar components include a command-line utility, catalog tables, an activity log, and an emergency boot file. You use ON-Bar with a storage manager and the XBSA shared library for the storage manager.

The following figure shows the ON-Bar and database server components:

- The storage spaces (dbspaces, blobspaces, and sbspaces) and logical logs from the database server
- The **sysutils** database, which includes ON-Bar catalog tables
- The onbar and the onbar-d command-line utilities
- The XBSA shared library for the storage manager on your system
- The storage media for storing backups
- The ON-Bar activity log
- The ON-Bar emergency boot file

Figure 7. ON-Bar components for Informix®



ON-Bar communicates with both the database server and the storage manager. You use the `onbar` command to start a backup or restore operation. By default, ON-Bar backs up and restores storage spaces in parallel. ON-Bar always processes log files serially.

For a backup session, ON-Bar requests the contents of storage spaces and logical logs from the database server and passes them to the storage manager. The storage manager stores the data on storage media. For a restore session, ON-Bar requests the backed up data from the storage manager and restores it on the database server.

ON-Bar backs up the critical `dbspaces` first, then the remaining storage spaces, and finally the logical logs. The *critical dbspaces* are the **rootdbs** and the `dbspaces` that contain the logical logs and physical log.

ON-Bar also places the following critical files in the archive during backups:

- The `onconfig` file
- UNIX™: The `sqlhosts` file
- The ON-Bar emergency boot file: `ixbar.servernum`
- The server boot file: `oncfg_servername.servernum`

You can restore storage spaces stored in both raw and cooked files. If your system contains primary and mirror storage spaces, ON-Bar writes to both the primary and mirror chunks at the same time during the restore, except for an external restore.

ON-Bar status and error messages are written to the activity log file: `bar_act.log`.

Backup Services API (XBSA)

ON-Bar and the storage manager communicate through the Backup Services Application Programming Interface (XBSA), which enables the storage manager to manage media for the database server. By using an open-system interface to the storage manager, ON-Bar can work with various storage managers that also use XBSA.

Each storage manager develops and distributes a unique version of the XBSA shared library. You must use the version of the XBSA shared library provided with the storage manager. For example, if you use HCL Informix® Primary Storage Manager, you must also use the XBSA shared library provided with ON-Bar. ON-Bar and the XBSA shared library must be compiled the same way (32-bit or 64-bit). For example, if you use ISM, also use the XBSA shared library provided with ON-Bar. ON-Bar and the XBSA shared library must be compiled the same way (32-bit or 64-bit).

ON-Bar uses XBSA to exchange the following types of information with a storage manager:

Control data

ON-Bar exchanges control data with a storage manager to verify that ON-Bar and XBSA are compatible, to ensure that objects are restored in the correct order to the correct instance of the database server, and to track the history of backup objects.

Backup or restore data

During backups and restores, ON-Bar and the storage manager use XBSA to exchange data from specified storage spaces or logical-log files.

ON-Bar uses XBSA transactions to ensure data consistency. All operations included in a transaction are treated as a unit. All operations within a transaction must succeed for objects transferred to the storage manager to be restorable.

Related information

[HCL Informix Primary Storage Manager on page 181](#)

ON-Bar catalog tables

ON-Bar uses the catalog tables in the **sysutils** database to track backup and restore operations. The `onsmsync` utility uses other catalog tables to track its operations.

ON-Bar uses the following catalog tables in the **sysutils** database to track backup and restore operations:

- The **bar_server** table tracks instances of the database server.
- The **bar_object** table tracks backup objects. A *backup object* is a backup of a dbspace, blobspace, sbspace, or logical-log file.

- The **bar_action** table tracks all backup and restore attempts against each backup object, except some log salvage and cold restore events.
- The **bar_instance** table describes each object that is backed up during a successful backup attempt.

The `onsmsync` utility uses and maintains the following tables to track its operations:

- The **bar_ixbar** table contains the history of all unexpired successful backups in all timelines.
- The **bar_syncdeltab** table is normally empty except when `onsmsync` is running.

For a description of the content of these tables, see [ON-Bar catalog tables on page 116](#).

ixbar file: ON-Bar emergency boot file

The emergency boot file is automatically updated after every backup. It contains the information that ON-Bar needs to perform a cold restore.



Important: Do not modify the emergency boot file. Doing so might cause ON-Bar to select the wrong backup as part of a restore, possibly leading to data corruption or system failure.

The file name for the boot file is `ixbar.servernum`, where `servernum` is the value of the `SERVERNUM` configuration parameter.

The ON-Bar emergency boot file is in the `$INFORMIXDIR/etc` directory on UNIX™ and in the `%INFORMIXDIR%\etc` directory on Windows™. You can override the default path and name of the boot file by changing the information specified in the `BAR_IXBAR_PATH` configuration parameter.

bar_act.log file: ON-Bar activity log

ON-Bar writes informational, progress, warning, error, and debugging messages to the ON-Bar activity log, `bar_act.log`.

ON-Bar backup and restore errors do not appear in standard output. If an error occurs when you back up and restore data, check information in the ON-Bar activity log

You can also use the activity log to:

- Monitor backup and restore activities such as, which storage spaces and logical logs were backed up or restored, the progress of the operation, and approximately how long it took.
- Verify whether a backup or restore succeeded.
- Track errors from the `ondblog` utility.
- Track ON-Bar performance statistics

The ON-Bar activity log is in the `/tmp` directory on UNIX™ and in the `%INFORMIXDIR%\etc` directory on Windows™. You specify the location of the ON-Bar activity log with the `BAR_ACT_LOG` configuration parameter.

Related reference

[BAR_ACT_LOG configuration parameter on page 222](#)

[ON-Bar messages and return codes on page 122](#)

[View ON-Bar backup and restore performance statistics on page 114](#)

[onbar -m syntax: Monitoring recent ON-Bar activity on page 53](#)

ON-Bar script

The ON-Bar utility includes a shell script on UNIX™ and a batch script on Windows™ for customizing backup and restore operations.

When you install ON-Bar with the database server, a default script is included. The name and location of the script depends on the operating system:

UNIX™

The onbar shell script is in the `$(INFORMIXDIR)/bin` directory.

Windows™

The onbar.bat batch script is in the `%INFORMIXDIR%\bin` directory.

When you issue ON-Bar commands from the command line, the arguments are passed to the script, and then to the `onbar_d` utility.

Table 4. ON-Bar utilities

Utility	Description
onbar_d utility	Transfers data between the database server and the storage manager. The onbar command calls the onbar_d utility that starts the onbar-driver. The onbar-driver starts and controls backup and restore activities.
onsmsync utility	Synchronizes the contents of the sysutils database, the emergency boot files, and the storage manager catalogs. Use this utility to purge backups that are no longer needed.
ondblog utility	Changes the database-logging mode. The ondblog utility logs its output in the ON-Bar activity log, <code>bar_act.log</code> .
archecker utility	Verifies backups, and restores table-level data from an archive.

Related reference

[The onsmsync utility on page 107](#)

Related information

[Overview of the archecker utility on page 204](#)

Storage management for ON-Bar

You must use a storage manager to perform backup and restore operations with ON-Bar. The *storage manager* is an application that manages the storage devices and media that contain backups. The storage manager handles all media labeling, mount requests, and storage volumes.

IBM® Informix® Storage Manager (ISM) is included with your database server. HCL Informix® Primary Storage Manager is included with your database server. However, you can choose to use another storage manager that ON-Bar supports and that is compatible with your storage devices.

IBM® Informix® Storage Manager (ISM)

ON-Bar is packaged with IBM® Informix® Storage Manager (ISM).

The ISM server resides on the same computer as ON-Bar and the HCL® Informix® database server; your storage devices are attached to this computer as well. ISM can store data on simple tape drives, optical disk devices, and file systems. ISM also performs the following functions:

- Configures up to four storage devices
- Adds, changes, and deletes administrative users
- Labels and mounts storage volumes on your storage devices
- Manages storage volumes
- Compresses and decompresses data
- Encrypts and decrypts data

Features that ISM supports

IBM® Informix® Storage Manager (ISM) supports the following storage-manager features:

- Back up logical logs and storage spaces to different devices and to specify whether to use encryption or compression for data.
- Write the output of parallel backups to a single device, medium, or volume. Some backup devices can write data faster than the disks used to hold storage spaces can be read.
- Automatic switch from one tape device to another when the volume in the first device fills.
- Migration of data from one backup medium to another. For speed, you can back up logical logs or storage spaces to disk, but you must move them later to tape or other removable media or your disk becomes full.
- Clone copies of backup data for on-site and off-site storage.
- Automatic expiration of data. Once all data on a backup media expires, you can reuse the media.

Features that ISM does not support

IBM® Informix® Storage Manager (ISM) does not support the following features:

- Distributing a single data stream across multiple devices simultaneously, which improves throughput if you have several slow devices
- Using different encryption or compression methods for specified storage spaces or databases
- Scheduling backups
- Support for devices such as tape libraries, jukeboxes, silos, tape autochangers, and stackers
- Remote host operations

You can install some storage managers on a different host from the database server. However, ISM must be installed on the same host as the database server.

Third-party storage managers might support these features.

IBM Spectrum Protect

You can use the IBM® Spectrum Protect product to manage backup or restore operations for ON-Bar. Spectrum Protect is a client/server program that provides storage management solutions to customers in a multivendor computer environment.

Spectrum Protect provides an automated, centrally scheduled, policy-managed backup, archive, and space-management facility for file servers and workstations. HCL® Informix® database servers are Spectrum Protect clients and Spectrum Protect stores data on separate Spectrum Protect servers.

Spectrum Protect features

Spectrum Protect efficiently manages disk, optical, and tape library resources. Spectrum Protect provides the following functions:

- Reduces network complexity with interfaces and functions that span network environments.
- Increases administrator productivity by automating repetitive processes, scheduling unattended processes, and administering Spectrum Protect from anywhere in the network
- Reduces risk of data loss with scheduled routine backups
- Optimizes existing storage resources with automated movement of files from client file systems to Spectrum Protect storage

Consider using IBM® Spectrum Protect if IBM® Informix® Storage Manager (ISM) does not support the features you require. For example, Spectrum Protect supports the following additional features:

- Creating policies to automate storage management and enforce data management goals.
- Automating circulation of media through the storage management process.
- Implementing a progressive backup methodology so that files are backed up incrementally to reduce network traffic, while recovery media is consolidated to provide better performance.
- Using the Network Data Management Protocol to back up and restore file systems stored on a network-attached storage file server.

Spectrum Protect software prerequisites

To back up to an IBM Spectrum Protect server, you must use the HCL® Informix® Interface for Spectrum Protect, which is part of the HCL® Informix® database server installation. The interface contains a shared library that enables ON-Bar to communicate with Spectrum Protect.

Spectrum Protect is distributed separately and supported for Informix® on specific platforms: <https://www.ibm.com/developerworks/wikis/display/idsinterop/IDS+with+IBM+Tivoli+Storage+Manager>. The minimum required version of the Spectrum Protect API for a supported platform is documented in the Informix® machine notes for that platform.

Non-IBM storage managers

ON-Bar supports some third-party storage managers that can manage stackers, robots, and jukeboxes and simple tape and disk devices.

Third-party storage managers might perform these functions, which are not supported by the software manager included with ON-Bar:

- Schedule backups
- Support networked and distributed backups and restores

Contact your Marketing representative for a list of supported storage managers, or go to <https://www.ibm.com/developerworks/wikis/display/idsinterop/IDS+with+Storage+Manager+Vendors>.

Configure the storage manager and ON-Bar

The topics in this section provide the information that you need to plan and to set up ON-Bar with a storage manager.

Related information

[Preparing to back up data on page 44](#)

Configure a storage manager

ON-Bar backup and restore operations require a storage manager that integrates with ON-Bar through an XBSA shared library interface.

You can choose to use the HCL Informix® Primary Storage Manager, the IBM® Spectrum Protect, or a third-party storage manager with ON-Bar. The Informix® Primary Storage Manager is bundled with Informix®. If you are using Spectrum Protect, the XBSA shared library needed for ON-Bar to communicate with Spectrum Protect is bundled with Informix®.

You can choose to use the IBM® Informix® Storage Manager (ISM), the IBM® Spectrum Protect, or a third-party storage manager with ON-Bar. ISM is installed with the database server. If use Spectrum Protect, the XBSA shared library needed for ON-Bar to communicate with Spectrum Protect is bundled with the database server.

The Informix® Primary Storage Manager manages storage for ON-Bar backup and restore operations, including parallel backups, that use file devices (disks) only, not tapes. By default, the HCL Informix® Primary Storage Manager

is automatically configured with the information specified in Informix® Primary Storage Manager and some ON-Bar configuration parameters. This storage manager is also automatically configured when you use the onpsm utility. You can change the configuration. For information, see [HCL Informix Primary Storage Manager on page 181](#) and [Configuring Informix Primary Storage Manager on page 191](#)

Related information

[HCL Informix Primary Storage Manager on page 181](#)

[Setting up Informix Primary Storage Manager on page 190](#)

Determine what storage devices you need

Before backing up, determine what and how many storage devices you need.

Ask the following interrelated questions to determine what storage devices you need. For example, the speed and type of storage devices partly determine the number of storage devices that you need.

- What type of storage devices do you need?

The transaction volume and the size of your database are major factors in determining the type of storage devices that you need.

IBM® Informix® Storage Manager (ISM) supports simple tape devices such as QIC, 4 mm, 8 mm, DLT, optical devices, and disk backups. If ISM cannot manage the storage devices that you need, you need to purchase a different storage manager. For more information, see the *Informix® Storage Manager Administrator's Guide*.

- What is the availability requirement for each device?

Is it important for your storage devices to allow random and sequential access? If so, you cannot use tape storage devices.

- How many storage devices do you need?

ISM supports up to four devices per host.

The number of storage devices that you need depends on the storage devices you have, how much transaction activity occurs on the database server, how fast throughput is, how much time you can allow for backups, and other similar factors.

Storage-manager definitions in the sm_versions file

Most storage managers must have an entry in the `sm_versions` file.

The HCL Informix® Primary Storage Manager and IBM Spectrum Protect do not require an entry in the `sm_versions` file.

The IBM Spectrum Protect does not require an entry in the `sm_versions` file.

The storage-manager definition in the `sm_versions` file uses this format:

```
1 | XBSA_ver|sm_name|sm_ver
```

In the format, `XBSA_ver` is the release version of the XBSA shared library for the storage manager, `sm_name` is the name of the storage manager, and `sm_ver` is the storage-manager version. The maximum field length is 128 characters.

Before ON-Bar starts a backup or restore process with the IBM Spectrum Protect, ISM, and third-party storage managers, ON-Bar calls the currently installed version of the storage-manager-specific XBSA shared library to get its version number. If this version is compatible with the current version of ON-Bar and is defined in the `sm_versions` file, ON-Bar begins the requested operation.

Related reference

[Configuring ISM on page 29](#)

Related information

[Updating the storage-manager definition in the `sm_versions` file for Spectrum Protect on page 35](#)

[Configuring a third-party storage manager on page 36](#)

Configuring ISM

You can use the IBM® Informix® Storage Manager (ISM) with ON-Bar.

The ISM server is installed with the HCL Informix® on UNIX™ or Windows™. Several database server instances can share one ISM instance.



Restriction: Install one copy of ISM on each computer to prevent possible conflicts with the XBSA shared library. Do not run ISM and Legato NetWorker on the same computer because they conflict with each other.

For instructions on how to set up ISM to work with ON-Bar, see the *Informix® Storage Manager Administrator's Guide*.

Before you use ISM for your backups, perform these tasks:

- Set certain environment variables and configuration parameters
- Update the storage-manager definition in the `sm_versions` file.

Related reference

[Storage-manager definitions in the `sm_versions` file on page 28](#)

ISM environment variables and `onconfig` parameters

When you use ISM, you must set the required environment variables in the `onbar` script or in your environment.

If you use ISM, you can specify the volume pool names for storage spaces and logical logs in the `ISM_DATA_POOL` and `ISM_LOG_POOL` configuration parameters in the `onconfig` file. The `ISM_DATA_POOL` configuration parameter specifies the

volume pool that you use for backing up storage spaces. The `ISM_LOG_POOL` configuration parameter specifies the volume pool that you use for backing up logical logs.

If you do not set these configuration parameters, they default to the volume pool names `ISMData` and then `ISMLogs`.

For information, see the *Informix® Storage Manager Administrator's Guide*.

Updating the storage-manager definition in `sm_versions` file for use with ISM

Before ON-Bar starts a backup or restore process, it calls the currently installed version of the storage-manager-specific XBSA shared library to get its version number. If this version is compatible with the current version of ON-Bar and is defined in the `sm_versions` file, ON-Bar begins the requested operation.

To update the storage-manager definition in `sm_versions` file:

1. Copy the `sm_versions.std` template to a new file, `sm_versions` in the `$INFORMIXDIR/etc` directory on UNIX™ or the `%INFORMIXDIR%\etc` directory on Windows™.
2. Issue the `ism_startup -init` command to automatically update `sm_versions` with the correct version number and storage-manager name or manually edit `sm_versions`.



Important: The `ism_startup -init` command erases records of previous backups.

If you are installing an ISM patch, you must manually edit `sm_versions`.

If you manually edit the `sm_versions`, you must put `ism` in the **sm_name** field.

3. Stop any ON-Bar processes (`onbar_d`, `onbar_w`, or `onbar_m`) that are currently running and restart them for the changes to take effect.

Example

The following example shows an ISM definition in the `sm_versions` file:

```
1|1.0.1|ism| ISM.2.20.UC1.114|
```

Configuring Spectrum Protect

To use IBM® Spectrum Protect with HCL® Informix® databases, you must install and configure the IBM Spectrum Protect client on your database server computer and IBM Spectrum Protect on your storage computer.

You must also configure HCL® Informix® Interface for Spectrum Protect and perform other Spectrum Protect configuration tasks on your HCL® Informix® database server computer.

Starting with versions 8.1.2 and 7.1.8, IBM Spectrum Protect requires SSL communication between the client and the server, and therefore additional configuration steps must be followed:

To configure Spectrum Protect:

1. Edit the Spectrum Protect client options files.
2. Assign a Spectrum Protect management class for the server to use for backups.
3. Set the HCL® Informix® Interface for Spectrum Protect environment variables.
4. Register with the Spectrum Protect server.
5. Initialize the HCL® Informix® Interface for Spectrum Protect.
6. Optional. Configure ON-Bar to support optional Spectrum Protect features.

The version of the XBSA shared library for Spectrum Protect is 1.0.3.

For details about Spectrum Protect, read the following manuals:

- *IBM Spectrum Protect Backup-Archive Clients Installation and User's Guide*
- *IBM Spectrum Protect Using the Application Program Interface*
- *IBM Spectrum Protect Administrator's Guide*
- *IBM Spectrum Protect Administrator's Reference*

Editing the Spectrum Protect client options files

The HCL® Informix® Interface for Spectrum Protect communicates with the Spectrum Protect server with the Spectrum Protect API. By default, HCL® Informix® Interface for Spectrum Protect uses the client user options file (`dsm.opt`) and, on UNIX™ systems, the client system options file (`dsm.sys`), both of which are located in the Spectrum Protect API installation directory.

On UNIX™ systems, edit both the `dsm.opt` and the `dsm.sys` files as the **root** user:

- Specify the Spectrum Protect server to use in the client user options file, `dsm.opt`.
- Identify the Spectrum Protect server name, communication method, and server options in the client system options file, `dsm.sys`.

Use the sample `dsm.opt.smp` and `dsm.sys.smp` files distributed with the Spectrum Protect API to help you get started quickly.

On Windows™ systems, specify the Spectrum Protect server name, communication method, and server options in the `dsm.opt` file.

The following example shows a simple `dsm.sys` on UNIX/Linux systems:

```
SErvername  TSMSRV01
TCPPOrt    9027
TCPADMINPort 9028
TCPSeRveraddress  tsmcentaur.myhcl.com
PASSWORdACCESS  GENERATE
PASSWORdDIR  /work/TSM/PAX12/TSMpsswDdir
NODENAME      pax12
ERRORLOGName  /work/TSM/PAX12/dsierror.log
```

The `TCPPOrt` in this file should be same as the `SSLTCPPOrt` used in the Spectrum Protect server.
The `TCPADMINPort` in this file should be same as the `SSLTCPADMINPort` used in the Spectrum Protect server.

See *Spectrum Protect Installing the Clients* and *Spectrum Protect Trace Facility Guide* for information regarding options you can specify in these files.

Editing the Spectrum Protect client user options file

You can edit the IBM® Spectrum Protect client user options file, `dsm.opt`. This file must refer to the correct Spectrum Protect server instance, as listed in the `dsm.sys` file.

Set the following options in the `dsm.opt` file:

SERVERNAME

Identifies which Spectrum Protect server instance, as listed in the `dsm.sys` file, that HCL® Informix® Interface for Spectrum Protect contacts for services.

TRACEFILE

Sends trace output information to a designated file.

TRACEFLAG

Sets specific trace flags

Editing the Spectrum Protect client system options file

You can edit the IBM® Spectrum Protect client systems options file, `dsm.sys`. This file must refer to the correct Spectrum Protect server address and communication method.

The following Spectrum Protect options are the most important to set in the `dsm.sys` file:

SERVERNAME

Specifies the name that you want to use to identify a server when it is referred to in the `dsm.opt` file and to create an instance that contains options for that server.

COMMMETHOD

Identifies the communication method.

TCPSERVERADDRESS

Identifies the Spectrum Protect server.

PASSWORDACCESS

Specifies GENERATE to store the Spectrum Protect password.

PASSWORDDIR

Directory where the Spectrum Protect client will store the SSL encryption related files and keystore.

NODENAME

Specifies the name of the Spectrum Protect client.



Note: Use the INFORMIXSERVER / DBSERVERNAME values in such a way that, it will be easy to identify in the Spectrum Protect catalog.

The SERVERNAME option in the `dsm.opt` and `dsm.sys` files define server instance names only. The TCPSERVERADDRESS option controls which server is contacted.

You can enable deduplication by including the DEDUP=CLIENTORSERVER option in the client system options file. You must also set the **IFX_BAR_USE_DEDUP** environment variable in the database server environment and restart the database server. See the *IBM Spectrum Protect Backup-Archive Client Installation and User's Guide* for information about configuring deduplication.

You can set up multiple server instances in the `dsm.sys` file. See the *IBM Spectrum Protect Backup-Archive Client Installation and User's Guide* for information about multiple server instances.

Related reference

[IFX_BAR_USE_DEDUP environment variable on page 242](#)

Related information

[Configuring ON-Bar for optional Spectrum Protect features on page 36](#)

Assigning a Spectrum Protect management class for a backup

When you back up a database, the default management class for your node is *used*. You can override the default value with a different value that is specified in the INCLUDE option.

The INCLUDE option is placed in the `include-exclude` options file. The file name of the `include-exclude` options file is in the client system options file (`dsm.sys`). For more information, see the *IBM Spectrum Protect Backup-Archive Client Installation and User's Guide*.

Use the following naming conventions for ON-Bar files:

- A database backup:

```
/dbservername/dbservername/dbspacename/level
```

- A log backup:

```
/dbservername/dbservername/server_number/unique_logid
```

For a database backup, an example of the INCLUDE statement is as follows:

```
Include /dbserverA/dbserverA/dbspaceA/* InformixDbMgmt
```

For a logical log backup, an example of the INCLUDE statement is as follows:

```
Include /dbserverA/dbserverA/55/* InformixLogMgmt
```

where the number `55` is the value of the SERVERNUM parameter in the `onconfig` file.

Setting the HCL Informix® Interface for Spectrum Protect environment variables

When you use the HCL Informix® Interface for TSM, you need to set certain environment variables in the environment of the user.

The following table describes these environment variables.

Table 5. HCL Informix® Interface for Spectrum Protect environment variables

Environment variable	Description
DSMI_CONFIG	The fully qualified name for the client user option file (<code>dsm.opt</code>). The default value is <code>dsm.opt</code> in the TSM API installation directory.
DSMI_DIR	<p>On UNIX™, points to the TSM API installed path. This environment variable needs to be defined only if the TSM API is installed in a different path from the default path. The <code>DSMI_DIR</code> environment variable is also used to find the <code>dsm.sys</code> file.</p> <p>On Windows™, specifies the installation location of the TSM Backup-Archive Client. Typically, the TSM Backup-Archive Client is installed in the <code>C:\Tivoli\TSMClient\baclient</code> directory.</p>
DSMI_LOG	<p>Points to the directory that contains the API error log file (<code>dsierror.log</code>).</p> <p>For error log files, create a directory for the error logs to be created in, then set the DSMI_LOG environment variable to that directory. The user informix or the backup operator should have write permission on this directory.</p>

The following example shows how to set up these environment variables for Solaris 32-bit if the TSM API is installed in the `/opt/Tivoli/tsm/client/api` directory:

```
export DSMI_CONFIG=/opt/Tivoli/tsm/client/api/bin/dsm.opt
export DSMI_DIR=/opt/Tivoli/tsm/client/api/bin
export DSMI_LOG=/home/user_a/logdir
```

The following example shows how to set up these environment variables for Windows™ if the TSM API is installed in the `C:\Tivoli\TSMClient\api` directory:

```
set DSMI_CONFIG=C:\Tivoli\TSMClient\api\BIN\dsm.opt
set DSMI_DIR=C:\Tivoli\TSMClient\baclient
set DSMI_LOG=C:\logdir
```

Registering with the Spectrum Protect server

Before backing up to and recovering from an IBM® Spectrum Protect server, you must have a Spectrum Protect registered node name and a password. The process of setting up a node name and password is called *registration*.

After the HCL® Informix® Interface for Spectrum Protect node is registered with a Spectrum Protect server, you can begin using the HCL® Informix® Interface for Spectrum Protect to back up and restore your HCL® Informix® storage spaces and logical logs. If your workstation has a node name assigned to the Spectrum Protect backup-archive client, you should have a

different node name for HCL® Informix® Interface for Spectrum Protect. For information about performing the registration process, see the *IBM Spectrum Protect Backup-Archive Client Installation and User's Guide*.

Initializing the HCL Informix® Interface for Spectrum Protect password

To initialize the password for HCL Informix® Interface for Spectrum Protect, use the **txbsapswd** program. This program sets up a connection with the server instance that you specified in the `dsm.opt` file.

About this task

You must run the **txbsapswd** program as user **root** before using HCL Informix® Interface for TSM.

To initialize the password:

1. Create the directory for the SSL-related files as user "informix" (PASSWORDDIR in the `dsm.sys` file):

```
$ sudo -u informix mkdir /work/TSM/PAX12/TSMpsswdDir
```

2. As user **root**, start the **txbsapswd** program located in `$INFORMIXDIR/bin` directory.
3. Enter the password and press Return. To retain your current password, press Return without a value.
4. Replace ownership and permissions in the `PASSWORDDIR`:

```
$ sudo chown -R informix:informix /work/TSM/PAX12/TSMpsswdDir
```

```
$ sudo chmod -R 750 /work/TSM/PAX12/TSMpsswdDir
```

Updating the storage-manager definition in the `sm_versions` file for Spectrum Protect

You must update the storage-manager definition in `sm_versions` file for ON-Bar to use with IBM® Spectrum Protect.

About this task

Before ON-Bar starts a backup or restore process, it calls the currently installed version of the storage-manager-specific XBSA shared library to get its version number. If this version is compatible with the current version of ON-Bar and is defined in the `sm_versions` file, ON-Bar begins the requested operation.

To update the storage-manager definition in `sm_versions` file:

1. Copy the `sm_versions.std` template to a new file, `sm_versions` in the `$INFORMIXDIR/etc` directory on UNIX™ or the `%INFORMIXDIR%\etc` directory on Windows™.
2. Put `tsm` in the **sm_name** field of the `sm_versions` file. The value `adsm` is also valid but will be deprecated in a future release.
3. Stop any ON-Bar processes (`onbar_d`, `onbar_w`, or `onbar_m`) that are currently running and restart them for the changes to take effect.

Example

The following example shows the IBM Spectrum Protect definition in the `sm_versions` file:

```
1|5.3|tsm|5
```

Related reference

[Storage-manager definitions in the sm_versions file on page 28](#)

Configuring ON-Bar for optional Spectrum Protect features

You can configure ON-Bar to enable or disable optional features in IBM® Spectrum Protect.

Apply all patches and updates to Spectrum Protect before you use the following Spectrum Protect features.

Deduplication

Deduplication eliminates redundant data in backups. To enable the database server to support deduplication, set the **IFX_BAR_USE_DEDUP** environment variable in the Informix® environment and restart the database server. Update the Spectrum Protect client systems option file.

Large transfer buffer size

The default transfer buffer size is 64 KB. Set the **BAR_XFER_BUF_SIZE** configuration parameter to specify a transfer buffer of up to 65 MB.

To limit the transfer buffer size to 64 KB regardless of the value of the **BAR_XFER_BUF_SIZE** configuration parameter, set the **IFX_NO_LONG_BUFFERS** environment variable to 1.

Replicate, import, and export backup objects

Replicating, importing, or exporting backup objects between Spectrum Protect servers requires unique IDs for backup objects. ON-Bar automatically stores IDs in the metadata of backup objects that are unique for all Spectrum Protect server with version 11.70.xC8 or later with version 12.10.xC2 or later.

To disable the ability to restore backup objects that are moved between Spectrum Protect servers, set the **IFX_TSM_OBJINFO_OFF** environment variable to 1.

Related reference

[IFX_BAR_USE_DEDUP environment variable on page 242](#)

[BAR_XFER_BUF_SIZE configuration parameter on page 240](#)

[Editing the Spectrum Protect client system options file on page 32](#)

[IFX_TSM_OBJINFO_OFF environment variable on page 243](#)

[IFX_BAR_NO_LONG_BUFFERS environment variable on page 242](#)

Configuring a third-party storage manager

Storage managers have slightly different installation and configuration requirements. If you use a third-party storage manager, make sure that you follow the manufacturer instructions carefully. If you have difficulty with the storage-manager installation and configuration, contact the manufacturer directly.

About this task

For the list of certified storage managers for your ON-Bar version, consult your marketing representative.



Important: Some storage managers let you specify the data to back up to specific storage devices. Configure the storage manager to back up logical logs to one device and storage spaces to a different device for more efficient backups and restores.

To configure a third-party storage manager:

1. Set ON-Bar configuration parameters and environment variables.
2. Configure the storage manager so that ON-Bar can communicate correctly with it.
For information, see your storage-manager documentation.
3. Configure your storage devices by following the instructions in your storage-manager documentation.
The storage manager must know the device names of the storage devices that it uses.
4. Label your storage volumes.
5. Mount the storage volumes on the storage devices.
6. Create the storage-manager definition in the `sm_versions` file. Use the definition provided by the vendor of the third-party storage manager.
 - a. Copy the `sm_versions.std` template to a new file, `sm_versions` in the `$INFORMIXDIR/etc` directory on UNIX™ or the `%INFORMIXDIR%\etc` directory on Windows™.
 - b. Create your own `sm_versions` file with the correct data for the storage manager by using the format in `sm_versions.std` as a template. To find out which code name to use in `sm_versions` for third-party storage managers, see the storage-manager documentation.
 - c. Stop any ON-Bar processes (`onbar_d`, `onbar_w`, or `onbar_m`) that are currently running and restart them for the changes to take effect.
7. Verify that the `BAR_BSALIB_PATH` configuration parameter points to the correct XBSA shared library for your storage manager.
8. If you enabled deduplication for your storage manager, set the `IFX_BAR_USE_DEDUP` environment variable and restart the database server.
9. ON-Bar uses the value of the `SERVERNUM` configuration parameter as part of the storage path for the logical logs in the storage manager. If the storage manager does not use a wildcard for the server number, set the appropriate server number environment variable for the storage manager.

What to do next

After you configure the storage manager and storage devices and label volumes for your database server and logical-log backups, you are ready to initiate a backup or restore operation with ON-Bar.

Related reference[IFX_BAR_USE_DEDUP environment variable on page 242](#)[Storage-manager definitions in the sm_versions file on page 28](#)

Specifying the location of the XBSA Library

By default, ON-Bar looks for the XBSA shared library in `$INFORMIXDIR/lib/ibsad001.s[ol]` on UNIX™. To specify a different name or location of the XBSA shared library, use the `BAR_BSALIB_PATH` configuration parameter.

UNIX™

You can also make `$INFORMIXDIR/lib/ibsad001.s[ol]` a symbolic link to the correct library.

For example, if you are using ISM, you can do either of the following:

- Link `$INFORMIXDIR/lib/ibsad001.so` to `$INFORMIXDIR/lib/libbsa.so`
- Set `BAR_BSALIB_PATH` to `$INFORMIXDIR/lib/libbsa.so`

For example, if you are using TSM, you can do either of the following:

- Link `$INFORMIXDIR/lib/ibsad001.so` to `$INFORMIXDIR/lib/libtxbsa.so`
- Set `BAR_BSALIB_PATH` to `$INFORMIXDIR/lib/libtxbsa.so`

Windows™

On Windows™, because no default XBSA shared library name exists, you must specify its name and location in the `BAR_BSALIB_PATH` configuration parameter. For example:

- If you are using ISM, set `BAR_BSALIB_PATH` to `%ISMDIR%\bin\libbsa.dll`.
- If you are using TSM, set `BAR_BSALIB_PATH` to `%DIR%\bin\libtxbsa.dll`.

If you are using a third-party storage manager, ON-Bar must use the version of the XBSA library that the storage-manager manufacturer provides. For more information, see [BAR_BSALIB_PATH configuration parameter on page 224](#) and your release notes.



Important: To set the path name of the XBSA library with the `BAR_BSALIB_PATH` configuration parameter in the `onconfig` file, specify the absolute path name. If you specify a relative path name, then the following message is written to the ON-Bar activity log: `BAR_BSALIB_PATH in ONCONFIG is not an absolute path name.`

Validating your storage manager

When you convert or revert the HCL® Informix® database server, the storage manager that you used on the old version might not be validated for the version that you are migrating to. Verify that the storage-manager vendor successfully completed the HCL® Informix® validation process for the database server version and platform.

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

Configuring ON-Bar

Before you begin your first backup, review the default ON-Bar parameters in the `onconfig` file and adjust the values as needed. You can also set an environment variable.

You can configure the behavior of ON-Bar by setting the following configuration parameters and environment variable.

Table 6. ON-Bar configuration parameters and environment variable

Behavior	Configuration parameters
Control the number and size of data buffers and the number of parallel processes.	BAR_NB_XPORT_COUNT configuration parameter on page 235
	BAR_XFER_BUF_SIZE configuration parameter on page 240
	IFX_BAR_NO_LONG_BUFFERS environment variable on page 242
	BAR_MAX_BACKUP configuration parameter on page 232
	BAR_MAX_RESTORE configuration parameter on page 234
Set the debugging level and the location of debug log file.	BAR_DEBUG configuration parameter on page 226
	BAR_DEBUG_LOG configuration parameter on page 230
Change the path of the ON-Bar boot file.	BAR_IXBAR_PATH configuration parameter on page 231
Maintain a history of expired backups.	BAR_HISTORY configuration parameter on page 231
Change the location and contents of ON-Bar activity log.	BAR_IXBAR_PATH configuration parameter on page 231
Maintain a history of expired backups.	BAR_ACT_LOG configuration parameter on page 222
	BAR_PROGRESS_FREQ configuration parameter on page 237

Table 6. ON-Bar configuration parameters and environment variable

(continued)

Behavior	Configuration parameters
Set automatic retrying of failed back ups or restores.	BAR_PERFORMANCE configuration parameter on page 236
Allow a failed restore to be restarted.	BAR_RETRY configuration parameter on page 237
Increase backup size estimate sent to the storage manager.	RESTARTABLE_RESTORE configuration parameter on page 248
Extend the time an RS secondary server waits for a checkpoint during an external backup.	BAR_SIZE_FACTOR configuration parameter on page 239
Configure continuous log backup.	BAR_CKPTSEC_TIMEOUT configuration parameter on page 225
Filter or transform backed up data with an external program.	ALARMPROGRAM configuration parameter on page
Optimize the deduplication capabilities of storage managers.	BACKUP_FILTER configuration parameter on page 221
Disable the ability to replicate, import, or export backup objects among TSM servers.	RESTORE_FILTER configuration parameter on page 249
Force the use of the <code>sm_versions</code> file.	IFX_BAR_USE_DEDUP environment variable on page 242
Do not set the LTAPEDEV configuration parameter to <code>/dev/null</code> or <code>NUL</code> because logical-log backups would be disabled and you can restore only whole-system backups.	IFX_TSM_OBJINFO_OFF environment variable on page 243

ON-Bar security

By default, only the **informix** or **root** users on UNIX™ system or members of the **Informix-Admin** group on Windows™ systems can run ON-Bar commands.

To enable additional users to run ON-Bar commands:

- On UNIX™ systems, create a **bargroup** group and add users to the group. For instructions on how to create a group, see your UNIX™ documentation.
- On Windows™ systems, add the users to the **Informix-Admin** group.



Restriction: For security, it is recommended that ON-Bar commands not be run by the **root** user.

Related reference

[onbar -r syntax: Restoring data on page 66](#)

[onbar -b syntax: Backing up on page 46](#)

[onbar -v syntax: Verifying backups on page 58](#)

[onbar -m syntax: Monitoring recent ON-Bar activity on page 53](#)

[onbar -P syntax: Printing backed-up logical logs on page 55](#)

[onbar -RESTART syntax: Restarting a failed restore on page 87](#)

Verifying the configuration of ON-Bar and your storage manager

Before you begin using ON-Bar and your storage manager, make sure that ON-Bar and your storage manager are set up correctly.

Verify your configuration by checking the items in the following list:

- The storage manager is installed and configured to manage specific storage devices.
- For UNIX™, make sure that the `BAR_BSA LIB_PATH` configuration parameter specifies correctly the XBSA shared library or it is not set and the library is in the default location.
- For Windows™, make sure that the `BAR_BSA LIB_PATH` configuration parameter specifies correctly the XBSA shared library.
- The `sm_versions` file contains a row that identifies the version number of the storage-manager-specific XBSA shared library.

After you verify that ON-Bar and your storage manager are set up correctly, run ON-Bar on your test database to make sure that you can back up and restore data. For more information, follow the instructions in [Back up with ON-Bar on page 43](#).

Files that ON-Bar and storage managers use

ON-Bar, HCL Informix® Primary Storage Manager IBM® Informix® Storage Manager (ISM), and IBM® Spectrum Protect use particular files in your installation.

The following table lists the files that ON-Bar, IBM® Informix® Storage Manager (ISM), and IBM® Spectrum Protect use and the directories where the files are. These names and locations change if you set up the `onconfig` file to values different from the defaults.

Table 7. List of files that ON-Bar, ISM, and Spectrum Protect use

File name	Directory	Purpose
ac_config.std	UNIX™: \$INFORMIXDIR/etc	Template for archecker parameter values.
	Windows™: %INFORMIXDIR%\etc	The ac_config.std file contains the default archecker (archive checking) utility parameters. To use the template, copy it into another file and modify the values.
ac_msg.log	/tmp	The archecker message log.
	%INFORMIXDIR%\etc	When you use archecker with ON-Bar to verify a backup, it writes brief status and error messages to the ON-Bar activity log and writes detailed status and error messages to the archecker message log. Technical Support uses the archecker message log to diagnose problems with backups and restores. Specify the location of the archecker message log with the AC_MSGPATH configuration parameter.
bar_act.log	/tmp	ON-Bar activity log.
	%INFORMIXDIR%	For more information, see bar_act.log file: ON-Bar activity log on page 23 .
bldutil.process_id	/tmp	When the sysutils database is created, error messages appear in this file.
	\tmp	
dsierror.log	\$DSMI_LOG	Spectrum Protect API error log.
dsm.opt	\$DSMI_CONFIG	Spectrum Protect client user option file.
dsm.sys	\$DSMI_DIR	Spectrum Protect client system option file.
Emergency boot files (ixbar* files)	\$INFORMIXDIR/etc	Used in a cold restore. For more information, see ixbar file: ON-Bar emergency boot file on page 23 .
	%INFORMIXDIR%\etc	
ISM catalog	\$INFORMIXDIR/ism	Records information about backup and restore save sets and storage volumes that ISM uses. ISM creates the ISM catalog during the ism_startup initialization. The ISM catalog records are stored in the mm, index, and res files. For more information, see the <i>Informix® Storage Manager Administrator's Guide</i> .
	%ISM%	

Table 7. List of files that ON-Bar, ISM, and Spectrum Protect use (continued)

File name	Directory	Purpose
ISM logs	\$INFORMIXDIR/ism/logs %ISM\logs	Operator alert messages, back-end status, additional ISM information. The ISM log names are <code>daemon.log</code> , <code>messages</code> , and <code>summary</code> .
ISMversion	\$INFORMIXDIR/ism DIR%\bin	Identifies the ISM version. Do not edit this file.
<code>oncfg_servername.servernum</code>	\$INFORMIXDIR/etc %INFORMIXDIR%\etc	Configuration information for ON-Bar restores. The database server creates the <code>oncfg_servername.servernum</code> file when you initialize disk space. The database server updates the file every time that you add or delete a dbspace, a logical-log file, or a chunk. The database server uses the <code>oncfg*</code> file when it salvages logical-log files during a cold restore. The database server uses the <code>oncfg*</code> files, so do not delete them.
<code>save, savegrp, savefs</code>	\$INFORMIXDIR/bin %ISM%\bin	
<code>sm_versions</code>	\$INFORMIXDIR/etc %INFORMIXDIR%\etc	Identifies the version of ISM or a third-party storage manager. To update the storage-manager version, edit the <code>sm_versions</code> file directly. The Informix® Primary Storage Manager does not use the <code>sm_versions.std</code> file.
<code>xbsa.messages</code>	\$INFORMIXDIR/ism/appl ogs %ISM%\applogs	XBSA library call information. ON-Bar and ISM use XBSA to communicate with each other. Technical Support uses the <code>xbsa.messages</code> log to fix problems with ON-Bar and ISM communications.

Back up with ON-Bar

You can use the ON-Bar utility to back up and verify storage spaces (dbspaces, blobspaces, and sbspaces) and logical-log files.

To perform a backup with ON-Bar:

1. Prepare for backup.
2. Back up with ON-Bar
3. Monitor backup progress.
4. Verify backups.
5. Back up storage manager information.

You can customize ON-Bar and storage manager commands in a shell or batch script. You can call ON-Bar from a job-scheduling program.

Related reference

[Customizing ON-Bar and storage-manager commands on page 103](#)

Preparing to back up data

Before you back up storage spaces and logical logs, you must prepare the system and copy critical administrative files.

To prepare to back up data:

1. Configure ON-Bar and your storage manager.
2. Ensure that you have enough logical log space.
ON-Bar checks for available logical-log space at the beginning of a backup. If the logs are nearly full, ON-Bar backs up and frees the logs before attempting to back up the storage spaces. If the logs contain ample space, ON-Bar backs up the storage spaces, then the logical logs.
3. Verify that you have enough temporary disk space.
The database server uses temporary disk space to store the before images of data that are overwritten while backups are occurring and overflow from query processing that occurs in memory. Verify that the **DBSPACETEMP** environment variable and DBSPACETEMP configuration parameter specify dbspaces that have enough space for your needs. If there is not enough room in the specified dbspaces, the backup will fail, root dbspace will be used, or the backup will fail after filling the root dbspace.
4. Back up administrative files to a different location.
5. Run the oncheck -cD command to verify that all database server data is consistent.
You do not need to check for consistency before every level-0 backup. Do not discard a backup that is known to be consistent until the next time that you verify the consistency of your databases.

Related reference

[Configure the storage manager and ON-Bar on page 27](#)

[onbar -b syntax: Backing up on page 46](#)

Related information

[oncheck -cd and oncheck -cD commands: Check pages on page](#)

Administrative files to back up

Although ON-Bar backs up some critical administrative files, you must also include critical files in normal operating-system backups of important configuration files.

Files that ON-Bar backs up

When you back up a storage space, ON-Bar also backs up the following critical files:

- The `onconfig` file
- UNIX™: The `sqlhosts` file
- The ON-Bar emergency boot file: `ixbar.servernum`
- The server boot file: `oncfg_servername.servernum`

You must restore these files if you need to replace disks or if you restore to a second computer system (imported restore).

Look at the `bar_act.log` file to determine whether critical files are successfully backed up. The return code for the `onbar -b` command indicates only whether storage spaces are successfully backed up. The following lines from the `bar_act.log` file show that the ON-Bar emergency boot file, `ixbar.0`, is backed up:

```
Begin backup of critical file '/opt/informix-11.70.fc7/etc/ixbar.0'.
Completed backup of critical file '/opt/informix-11.70.fc7/etc/ixbar.0'
```

Files that you must manually back up

In addition to the critical files, you must also manually back up the following administrative files:

- The `sm_versions` file
- Storage-manager configuration and data files
- Simple-large-object data in blobspaces that are stored on disks
- Simple-large-object data in blobspaces that are stored on disks or optical platters
- Externally stored data such as external tables that a DataBlade® maintains
- The keystore and stash files for encrypting storage spaces, as specified by the `DISK_ENCRYPTION` configuration parameter



Tip: Even though ON-Bar includes the critical files with the files it backs up, it is a good practice to also include the critical files in your system archive. Having the critical files included in both the HCL Informix® and system archives gives you more options if you need them.

Files that ON-Bar re-creates

Although ON-Bar does not back up the following items, ON-Bar automatically re-creates them during a restore. You do not need to make backup copies of these files:

- The `dbspace` pages that are allocated to the database server but that are not yet allocated to a `tblspace` extent
- Mirror chunks, if the corresponding primary chunks are accessible

- Temporary dbspaces

ON-Bar does not back up or restore the data in temporary dbspaces. Upon restore, the database server re-creates empty temporary dbspaces.

onbar -b syntax: Backing up

Use the onbar -b command to back up storage spaces and logical logs.

To run ON-Bar commands, you must be user **root**, user **informix**, a member of the **bargroup** group on UNIX™, or a member of the Informix®-Admin group on Windows™.

- [Usage on page 49](#)
- [Example: Back up a whole system on page 49](#)
- [Example: Back up all online storage spaces and logical logs on page 49](#)
- [Example: Perform an incremental backup on page 50](#)
- [Example: Back up specified storage spaces and all logical logs on page 50](#)
- [Example: Back up a list of storage spaces specified in a file on page 50](#)
- [Example: Back up logical logs on page 50](#)
- [Example: Physical backup on page 50](#)

Syntax for backing up with ON-Bar

```
onbar -b {{{[ -L{0 | 1 | 2}][ -p] | { -filename | dbspace_list | [-w]}} | -l [{ -c | -c | -s}]} [ -o] [ -cf {yes | no | only}] | -F }
```

Table 8. Options for the onbar -b command

Option	Description
-b	Specifies a backup Backs up the storage spaces and logical logs, including the current logical log.
<i>dbspace_list</i>	Specifies the storage spaces to be backed up, separated by blank spaces. If you do not enter <i>dbspace_list</i> or <i>-f filename</i> , ON-Bar backs up all online storage spaces on the database server.
-c	Closes and backs up the current logical log and the other full logical logs.
-C	Starts a continuous log backup. Reserve a dedicated storage device and terminal window because the continuous log backups run indefinitely waiting for logical logs to fill.

Table 8. Options for the onbar -b command (continued)

Option	Description
	To stop a continuous log backup, stop the ON-Bar process with an interrupt command, such as CTRL-C or SIGTERM.
-cf	Specifies whether the critical files are backed up. The critical files are the <code>onconfig</code> file, the <code>sqlhosts</code> file, and the <code>ixbar.servernum</code> file.
	<p>Valid values are:</p> <ul style="list-style-type: none"> • <code>yes</code> = Backs up the critical files. Default when performing a level 0, 1, or 2 backup. • <code>no</code> = Does not back up the critical files. Default when backing up the logical log files. • <code>only</code> = Backs up only the critical files.
-f <i>filename</i>	<p>Backs up the storage spaces that are listed in the text file that is specified by the <i>filename</i> value.</p> <p>Use this option to avoid entering a long list of storage spaces every time that you back up.</p> <p>For more information, see List of storage spaces in a file on page 51.</p>
-F	<p>Performs a fake backup</p> <p>A storage-manager application is not necessary. No backup actually occurs, so no restore is possible from a fake backup. Use fake backups sparingly, if at all. Fake backups might be appropriate in the following situations:</p> <ul style="list-style-type: none"> • Change database logging modes • Change a RAW table to a STANDARD table • Allow the user to use new logs, chunks, or mirrors without performing a backup • In special situations when you, the administrator, judge that a backup is not needed
-L	<p>Specifies the level of backup to perform on storage spaces:</p> <ul style="list-style-type: none"> • 0 = a complete backup (Default) • 1 = changes since the last level-0 backup • 2 = changes since the last level-1 backup

Table 8. Options for the onbar -b command (continued)

Option	Description
	<p>If you request an incremental backup and ON-Bar finds that no previous level backup was performed for a particular storage space, ON-Bar backs up that storage space at the previous level. For example, if you request a level-1 backup, and ON-Bar finds no level-0 backup, it makes a level-0 backup instead.</p>
-l	<p>Performs a backup of full logical-log files.</p> <p>The current logical-log file is not backed up.</p> <p>If you are using ISM, it also backs up the ISM catalog.</p>
-O	<p>Overrides normal backup restrictions.</p> <p>Use this option to back up logical logs when blobspaces are offline.</p> <p>If a log backup occurs when blobspaces are offline, return code 178 displays in the ON-Bar activity log.</p>
-p	<p>Backs up only physical storage spaces without logical logs.</p> <p>A warning message is written to the activity log listing the log unique ID of the latest log file that is required for a restore of the storage spaces. Use this option if logical logs are being continuously backed up. If necessary, a log switch is initiated, so that this log can be backed up. If the current log is already newer than the log with the archive checkpoint of the last storage space, then no log switch is initiated.</p>
-s	<p>Salvages any logical logs that are still on disk after a database server failure. You can run the onbar -l -s command while the server is offline.</p> <p>If possible, use this option before you replace a damaged disk. If you use onbar -r to perform a cold restore on an undamaged disk, ON-Bar automatically salvages the logical logs.</p>
-w	<p>Backs up a whole system, which includes all storage spaces and logical logs based on a single checkpoint.</p> <p>The time of the backup is stored with the backup information. The data in all storage spaces is consistent in a whole-system backup, therefore, you do not need to restore the logical logs to make the data consistent. If you do not save the logical logs, you must use the -w option.</p>

Usage

Before you back up your data, make sure that your data is consistent by running the `oncheck -cD` command.

To run ON-Bar commands, you must be user **root**, user **informix**, or a member of the **bargroup** group on UNIX™, or a member of the **Informix-Admin** group on Windows™. For more information, see [ON-Bar security on page 40](#).

You can back up storage spaces and logical logs when the database server is in online, quiescent, or fast-recovery mode.

The storage-space chunks can be stored on raw disk storage space, in cooked files, or on an NTFS file system (Windows™).

Only online storage spaces are backed up. Use the `onstat -d` command to determine which storage spaces are online. During a backup, if ON-Bar encounters a down dbspace, it skips it and later returns an error. If a storage space is offline, restart the backup when the storage space is back online.

After you begin the backup, monitor its progress in the ON-Bar activity log and database server message log.

You can either back up the logical logs separately or with storage spaces. Back up the logical logs as soon as they fill so that you can reuse them and to protect against data loss if the disks that contain the logs are lost. If the log files fill, the database server pauses until you back up the logical logs. You can either back up the logical logs manually or start a continuous logical-log backup by running the `onbar -b -C` command. Logical-log backups are always level 0. After you close the current logical log, you can back it up.

If you perform whole-system backups and restores, you do not need to restore logical logs. However, back up the logical logs when you use whole-system backups. These log backups allow you to recover your data to a time after the whole-system backup, minimizing data loss.

If you are running continuous logical log backup and then start a whole system backup, the ON-Bar process attempts to save the logical logs. Because the continuous logical log backup is running, an error message is returned indicating that a logical log backup is already running, and the whole system backup returns a non-zero error code. In this case the logical logs are backed up only one time. To avoid the error, create a physical backup with the `onbar -b -w -p` command.

To back up a specific table or set of tables in ON-Bar, store these tables in a separate dbspace and then back up this dbspace. Alternatively, you can perform table level restores with the `archecker` utility.

Example

Example: Back up a whole system

The following command performs a level-0 whole system backup after taking a checkpoint of all online storage spaces and logical logs:

```
onbar -b -w
```

The following command performs a level-1 whole system backup:

```
onbar -b -w -L 1
```

Example

Example: Back up all online storage spaces and logical logs

The following command performs a standard, level-0 backup of all online storage spaces and used logical logs:

```
onbar -b
```

Example**Example: Perform an incremental backup**

The following command performs a standard, level-1 backup:

```
onbar -b -L 1
```

Example**Example: Back up specified storage spaces and all logical logs**

The following command performs a level-0 backup of the dbspaces named **fin_dbspace1** and **fin_dbspace2** and all logical logs:

```
onbar -b fin_dbspace1 fin_dbspace2
```

Example**Example: Back up a list of storage spaces specified in a file**

The following sample file named `listfile3` contains a list of storage spaces to be backed up: **blobsp2.1**, **my_dbspace1**, **blobsp2.2**, **dbsl.1**, **rootdbs.1**, and **dbsl.2**.

```
blobsp2.1
# a comment                ignore this text

        my_dbspace1        # back up this dbspace
; another comment
blobsp2.2                dbsl.1
rootdbs.1        dbsl.2 ; backing up two spaces
```

The following command backs up the storage spaces listed in the `listfile3` file:

```
onbar -b -f listfile3
```

Example**Example: Back up logical logs**

The following command starts a manual logical-log backup:

```
onbar -b -l
```

The following command backs up the current logical-log file:

```
onbar -b -l -c
```

Example

Example: Physical backup

The following command backs up all storage spaces without backing up any logical logs:

```
onbar -b -p -L 0
```

A warning message is written to the ON-Bar activity log file stating that log file backup was not initiated. The message also contains the log unique ID of the latest log file that is required for a restore of the storage spaces. The latest required log file contains the archive checkpoint of the last dbspace backed up.

Example message:

```
2011-12-14 09:30:35 14277 14275 (-43354) WARNING: Logical logs were
not backed up as part of this operation. Logs through log unique ID 9
are needed for restoring this backup. Make sure these logs are backed
up separately.
```

Related reference

[Plan a backup system for a production database server on page 16](#)

[ON-Bar security on page 40](#)

Related information

[Configuring a continuous log restore by using ON-Bar on page 76](#)

[Replacing disks during a restore on page 81](#)

[Preparing to back up data on page 44](#)

List of storage spaces in a file

You can list storage spaces to back up or restore in a file.

The *filename* value can be any valid UNIX™ or Windows™ file name:

- Simple file names, for example: `listfile_1`)
- Relative file names, for example: `../backup_lists/listfile_2` or `..\backup_lists\listfile2`
- absolute file names, for example: `/usr//backup_lists/listfile3` or `c:\backup_lists\listfile3`

The format rules for the file are:

- If you are restoring chunks, list storage space names without paths. Each line can list more than one storage space, separated by spaces or a tab.
- If you are renaming chunks, list the old chunk path name, the old offset, the new chunk path name, and the new offset. Put a blank space or a tab between each item. Put information for each chunk on a separate line.
- Comments begin with a # or a ; symbol and continue to the end of the current line.
- ON-Bar ignores all comment or blank lines in the file.

Use ISM during a backup

Use the `ism_watch` command to monitor backups and restores sent to the IBM® Informix® Storage Manager (ISM) server. During a backup, the ISM server automatically routes storage-space data to volumes in the ISMData volume pool and logical-log files to volumes in the ISMLogs volume pool.

Always keep the volumes from the ISMLogs pool mounted to ensure that a storage device is always available to accept logical-log data. If the volume is not mounted, the backup pauses. For more information about using devices and ISM commands, see the *Informix® Storage Manager Administrator's Guide*.

During the backup operation, ISM creates *save sets* of the backed up data and enters information about the backed up data in the ISM catalog. You can also use the `ism_catalog -create_bootstrap` command to back up the ISM catalog:

If you use the `onbar` script to back up storage spaces and logical logs, it backs up the ISM catalog automatically. If you call `onbar_d` directly, you must use the `ism_catalog -create_bootstrap` command.

Backing up the ISM catalog

The `ism_catalog -create_bootstrap` command creates a full backup of the IBM® Informix® Storage Manager (ISM) catalog when the catalog is first used. Afterwards, the `ism_catalog -create_bootstrap` command performs incremental backups.

Therefore, the media or volume containing the first bootstrap creation is always necessary in order to restore the bootstrap successfully.

To avoid the need for the first media or volume during bootstrap restore, a full bootstrap backup must be performed every time. The level must be specified explicitly by using, the underlying `savegrp` command that `ism_catalog -create_bootstrap` calls.

To perform a full backup of the ISM catalog every time, run the `$INFORMIXDIR/bin/savegrp -O -l full ISMData` command as the user root.

In this command, `-O` is a capital O (not a zero), and `ISMData` is the pool name where the bootstrap data is saved.

You can edit the `$INFORMIXDIR/bin/onbar` shell script to replace the `ism_catalog -create_bootstrap` command with the `savegrp` command.

Redirect the output (both `stdout` and `stderr`) of the `savegrp` command to the ON-Bar activity log file (specified by the `BAR_ACT_LOG` configuration parameter in `onconfig` file).

Backing up blobspaces

You can back up blobspaces in a database that uses transaction logging.

Before you begin

Before you back up a new blobspace, make sure that the log file that recorded the creation of the blobspace is no longer the current log file. You can run the `onstat -l` command to verify the logical-log status.

About this task

When users update or delete simple large objects in blobspaces, the blobpages are not freed for reuse until the log file that contains the delete records is freed. To free the log file, you must back it up.

! **Important:** If you perform a warm restore of a blobspace without backing up the logical logs after updating or deleting data in it, that blobspace might not be restorable.

To back up blobspaces:

1. Verify the logical-log status by running the `onstat -l` or `xctl onstat -l` command.
2. Switch to the next log file by running the `onmode -l` command.
3. Back up the logical logs:

Choose from:

- If the blobspace is online, run the `onbar -b -l -c` command.
 - If the blobspace is offline, run the `onbar -b -l -O` or `onbar -b -O` command. If this backup is successful, ON-Bar returns `178`.
4. Back up the blobspaces by running the `onbar -b` or `onbar -b -w` command.

Related information

`onstat -L` command: [Print the number of free locks on page](#)

onbar -m syntax: Monitoring recent ON-Bar activity

You can monitor recent ON-Bar activity with the `onbar -m` command. Only users who have permission to perform backup and restore operations can use this option.

Monitor recent ON-Bar activity

```
onbar -m { 20 | [lines] } [ -r { 5 | [seconds] } ]
```

Table 9. Options for the onbar -m command

Option	Description
<code>-m</code>	Prints the recent activity of ON-Bar from the activity log file.
<code>lines</code>	Specifies the number of lines to output. Default is 20 lines.
<code>-r</code>	Causes the <code>onbar -m</code> command to repeat.
<code>seconds</code>	Specifies the number of seconds to wait before repeating. Default is 5 seconds.

Related reference[ON-Bar messages and return codes on page 122](#)[Message format in the ON-Bar message log on page 123](#)[ON-Bar security on page 40](#)**Related information**[bar_act.log file: ON-Bar activity log on page 23](#)

Viewing a list of registered backups

You can create a list of the registered ON-Bar backups performed on your system.

To view the list of registered backups:

1. Create a view in the **sysutils** database that contains information from the **bar_action**, **bar_instance**, and **bar_object** catalog tables.

Include the following fields in the view:

- Backup_ID: The internally generated ID for the backup
- Type: Defines whether the backup is a whole system backup, dbspace backup, or logical log backup.
- Object_Name: The name of the object backed up.
- Ifx_Time: Time at which the object was created. For dbspace backups, the checkpoint time that started the backup. For logical logs, the time when the log become full.
- CopyID_HI: The High part of the ID to locate the object in the storage manager.
- CopyID_LO: The Low part of the ID to locate the object in the storage manager.
- Backup_Start: Date and time when the backup started for this object
- Backup_End: Date and time when the backup ended for this object.
- Verify_Date: The time of the last verification made to this object, if any.

2. Run a SELECT statement against the view.

Example**Example**

The following statement creates a view that contains backup information:

```
CREATE VIEW list_backups(Backup_ID, Type, Object_Name, Ifx_Time, CopyID_HI,
    CopyID_LO, Backup_Start, Backup_End, Verify_Date)
AS SELECT * FROM (
SELECT
    act_aid AS backup_id,
    DECODE(act_type, 5, "Whole-System", DECODE(obj_type, "L",
        "Logical log", "Dbspace")) AS Type,
    substr(obj_name,1, 8) AS Object_Name,
    min(DBINFO ('utc_to_datetime', seal_time)) AS Ifx_Time,
    ins_copyid_hi AS CopyID_HI,
    ins_copyid_lo AS CopyID_LO,
    act_start AS Backup_Start,
```

```

    act_end AS Backup_End,
    ins_verify_date AS Verify_Date

FROM
    bar_action A,
    bar_instance I,
    bar_object O

WHERE
    A.act_aid = I.ins_aid AND
    A.act_oid = O.obj_oid AND
    A.act_oid = I.ins_oid AND
    O.obj_type in ("R", "CD", "ND", "L")

GROUP BY 1,2,3,5,6,7,8,9
ORDER BY Ifx_Time, Backup_ID) AS view_list_backups

```

The following query returns all the backups:

```
SELECT * FROM list_backups
```

Related information

[The sysutils Tables on page](#)

onbar -P syntax: Printing backed-up logical logs

You can use the onbar -P command to print logical logs that are backed up using the ON-Bar utility.

To run ON-Bar commands, you must be user **root**, user **informix**, a member of the **bargroup** group on UNIX™, or a member of the Informix®-Admin group on Windows™.

- [Usage on page 57](#)
- [Example: Print a specific transaction on page 57](#)
- [Example: Print multiple logical log files on page 57](#)

Print backed-up logical logs

```
onbar -P -n { unique_id | starting_id-ending_id } [-l] [-q] [-b] [-c] [-uusername] [-tblspace_num] [-
xtransaction_num ]
```

Table 10. Options for the onbar -P command

Option	Purpose
-b	Print logical-log records associated with blobspace blobpages. The database server stores these records on the logical-log backup media as part of blobspace logging.
-c	Use the compression dictionary to expand compressed data.
-l	Print the long listing of the logical-log record.

Table 10. Options for the onbar -P command (continued)

Option	Purpose
	The long listing of a log record includes complex hexadecimal and ASCII dumps of the entire log record.
-n <i>starting_id-ending_id</i>	Print the logical-log records contained in the specified range of log files. The <i>starting_id</i> option is the ID of the first log to print. The <i>ending_id</i> option is ID of the last log to print. The value of the <i>starting_id</i> option must be smaller than the value of the <i>ending_id</i> option. Separate the starting and ending ID values with a hyphen. Do not include blank spaces.
-n <i>unique_id</i>	Print the logical-log records contained in the specified log file. The <i>unique_id</i> option is the unique ID number of the logical log. To determine the unique ID of a specific logical-log file, use the <code>onstat -l</code> command.
-P	Print backed-up logical log information
-q	Do not print the program header
-t <i>tblspace_num</i>	Print the records associated with the <i>tblspace</i> that you specify with the <i>tblspace_num</i> option. Specify the <i>tblspace_num</i> value as either an unsigned integer or hexadecimal value. If you do not use a prefix of <code>0x</code> , the value is interpreted as an integer. The integer must be greater than zero and must exist in the partnum column of the systables system catalog table.
-u <i>username</i>	Print the records for a specific user. The user name must be an existing login name and conform to operating-system-specific rules for login names.
-x <i>transaction_num</i>	Print only the records associated with the transaction that you specify. The <i>transaction_num</i> must be an unsigned integer between zero and <code>TRANSACTIONS -1</code> , inclusive. Additional Information: Use the <code>-x</code> option only in the unlikely situation of an error being generated during a roll-forward. When this situation occurs, the database server sends a message to the message log that includes the transaction ID of the offending transaction. You can

Table 10. Options for the onbar -P command (continued)

Option	Purpose
	use this transaction ID with the -x option to investigate the cause of the error.

Usage

To view the backed-up logical logs, the storage manager must be running.

The output of this command is printed to stdout.

Example

Example: Print a specific transaction

The following command prints information about a single transaction that was performed by the user **informix** against the tblspace 1048722 and is contained in the logical log file 2:

```
onbar -P -n 2 -l -q -b -u "informix" -t 1048722 -x 1
```

The output for this command might be:

```
log uniqid: 2.
1665d0 120 DPT      1          2 0          5
00000078 0002006c 00000010 0000fefe ...x...l .....
00000001 00000000 000077e3 00000000 .....e..w....
00000005 00000005 00002a24 00000001 ..... .*$.
00100004 0a0c21b8 00002a48 00000001 .....!. .*H....
00100006 0a0c2288 00002ea1 00000001 .....". .....
0010001b 0a0c3810 00002bee 00000001 .....8. ..+....
00100015 0a0c3a18 00002a3d 00000001 .....: .*=....
00100005 0a0c57c0 .....W.
166648 60 CKPOINT 1          0 1665d0 1
0000003c 00000042 00000010 0000fefe ...<...B .....
00000001 001665d0 000077e3 00000000 .....e..w....
00010005 00000002 00000002 001665a0 .....e. ....e.
00000007 ffffffff 00084403 ..... ..D.
```

Example

Example: Print multiple logical log files

The following command prints the logical log records for the logical logs files that have IDs of 2, 3, 4, 5, 10, 11, and 12:

```
onbar -P -n 2-5 -n 10-12
```

Related reference

[ON-Bar security on page 40](#)

Related information

[onstat -l command: Print physical and logical log information on page](#)

[onstat -L command: Print the number of free locks on page](#)

[SYSTABLES on page](#)

onbar -v syntax: Verifying backups

Use the `onbar -v` command to verify that backups that were created by the ON-Bar utility are complete and can be restored.

To run ON-Bar commands, you must be user **root**, user **informix**, a member of the **bargroup** group on UNIX™, or a member of the Informix®-Admin group on Windows™.

Sufficient temporary space must be available. For more information, see [Temporary space for backup verification on page 60](#).

- [Usage on page 59](#)
- [Example: Perform a point-in-time verification of a backup on page 59](#)
- [Example: Verify backups of storage spaces listed in a file on page 59](#)
- [Example: ON-Bar activity log verification messages on page 60](#)
- [Example: archecker message log verification messages on page 60](#)

Verify backups

```
onbar -v [-p] [-t "time" ] [{ -f filename | space | [-w] ]
```

Table 11. Options for the onbar -v command

Option	Description
-v	Verifies a backup. The server can be in any mode. If verification is successful, you can restore the storage spaces safely. You can verify a whole-system or physical-only backup. You cannot verify the logical logs.
space	Names of storage spaces to verify. If you enter more than one storage-space name, use a space to separate the names.
-f filename	Verifies the storage spaces that are listed in the text file whose path name <i>filename</i> provides. Use this option to avoid entering a long list of storage spaces every time that you verify them.

Table 11. Options for the onbar -v command (continued)

Option	Description
	<p>You can use any valid UNIX™ or Windows™ path name and file name. For the format of this file, see List of storage spaces in a file on page 51.</p> <p>The file can list multiple storage spaces per line.</p>
-p	Verifies a physical-only backup.
-t " <i>time</i> "	<p>Specifies the date and time to which dbspaces are verified. Must be surrounded by quotation marks.</p> <p>How you enter the time depends on your current GLS locale convention. If the GL_DATETIME environment variable is set, you must specify the date and time according to that variable. If the GLS locale is not set, use ANSI-style date format: <code>YYYY-MM-DD HH:MM:SS</code>.</p>
-w	Verifies a whole-system backup.

Usage

The onbar -v command runs the archecker utility. The archecker utility verifies that all pages required to restore a backup exist on the media in the correct form. After you successfully verify a backup, you can restore it safely.

When you verify a backup, ON-Bar writes summary messages to the `bar_act.log` that report which storage spaces were verified and whether the verification succeeded or failed. The archecker utility writes detailed messages to the `ac_msg.log`. IBM® Software Support uses the `ac_msg.log` to diagnose problems with backups and restores.

The onbar -v command verifies only the smart-large-object extents in an sbpace. For a complete check, use the oncheck -cS command.

The onbar -v command cannot verify the links between data rows and simple large objects in a blobspace. Use the oncheck -cD command instead to verify the links in a blobspace.

Example

Example: Perform a point-in-time verification of a backup

The following command verifies a backup at a point-in-time:

```
onbar -v -t "2011-12-10 10:20:50"
```

Example

Example: Verify backups of storage spaces listed in a file

The following command verifies the backed-up storage spaces that are listed in the file `bkup1`:

```
onbar -v -f /usr/backups/bkup1
```

Example

Example: ON-Bar activity log verification messages

The following examples show messages about verification in the ON-Bar activity log:

The level-0 backup of dbspace **dbs2.2** passed verification, as follows:

```
Begin backup verification of level0 for dbs2.2 (Storage Manager Copy ID:##)
Completed level-0 backup verification successfully.
```

The level-0 backup of **rootdbs** failed verification, as follows:

```
Begin backup verification of level0 for rootdbs (Storage Manager Copy ID:##).
ERROR: Unable to close the physical check: error_message.
```

Example

Example: archecker message log verification messages

More detailed information is available in the archecker message log, as follows:

```
STATUS: Scan PASSED
STATUS: Control page checks PASSED
STATUS: Starting checks of dbspace dbs2.2.
STATUS: Checking dbs2.2:TBLSpace
.
.
STATUS: Tables/Fragments Validated: 1
Archive Validation Passed
```

Related reference

[ON-Bar security on page 40](#)

Related information

[Verifying an expired backup on page 63](#)

Temporary space for backup verification

When you verify backups, 15-25 MB of temporary space must be available.

During backup verification, the archecker utility requires about 15 MB of temporary space for a medium-size system (40-50 GB) and 25 MB for a large system. This temporary space is stored on the file system in the directory that the AC_STORAGE parameter specifies, not in the dbspaces. The temporary files contain bitmap information about the backup and copies of partition pages, free pages in a chunk, reserved pages, and optionally, free pages in a blobspace and debugging information. The archecker utility must have permissions to the temporary directory.

If the backup is verified successfully, these files are deleted. If the backup fails verification, these files remain. Copy them to another location so that IBM® Software Support can review them.

If your database server contains only dbspaces, use the following formula to estimate the amount of temporary space in KB for the archecker temporary files:

```
space = (130 KB * number_of_chunks) + (pagesize * number_of_tables) +
(.05 KB * number_of_logs)
```

For HCL Informix®, if your database server contains blobspaces or sbspaces, use the following formula to estimate the amount of temporary space for the archecker temporary files:

```
space = (130 KB * number_of_chunks) + (pagesize * number_of_tables) +
(.05 KB * number_of_logs) + (pagesize * (num_of_blobpages/252))
```

number_of_chunks

The maximum number of chunks that you estimate for the database server.

pagesize

The system page size in KB.

number_of_tables

The maximum number of tables that you estimate for the database server.

number_of_logs

The number of logical logs on the database server.

num_of_blobpages

The number of blobpages in the blobspaces or the number of sbspaces. (If your database server contains sbspaces, substitute *num_of_blobpages* with the number of sbspaces.)

For example, you would need 12.9 megabytes of temporary disk space on a 50-gigabyte system with a page size of 2 KB. This system does not contain any blobspaces, as the following statement shows:

```
13,252 KB = (130 KB * 25 chunks) + (2 KB * 5000 tables) +
(.05 KB * 50 logs) + (2 KB * 0)
```

To convert KB to MB, divide the result by 1024:

```
12.9 MB = 13,252/1024
```

Verification failures

The verification of a backup can fail for various reasons. If a backup fails verification, do not attempt to restore it.

The causes of a verification failure are unpredictable and range from corruption of the database server to a failed restore because ON-Bar cannot find the backup object on the storage manager. In fact, the restore might appear to be successful but it hides the real problem with the data or media.

Backups with corrupted pages

If the pages are corrupted, the problem is with the databases rather than with the backup or the media.

Run `oncheck -cd` on any tables that produce errors and then redo the backup and verification. To check extents and reserved pages, run `oncheck -ce` and `oncheck -cr`.

Backups with corrupted control information

In this case, all the data is correct, but some of the backup control information is incorrect, which might cause problems with the restore. Ask IBM® Software Support for assistance.

Backups with missing data

When a backup is missing data, it might not be recoverable. After a data loss, try to restore from an older backup. Then restore the current logical logs.

Backups of inconsistent database server data

There are cases where archecker returns "success" to ON-Bar but shows "failure" in the archecker message logs. This situation occurs when archecker verifies that ON-Bar backed up the data correctly, but the database server data was invalid or inconsistent when it was backed up.

Diagnosing why a backup failed verification

If a backup failed verification, you can take steps to diagnose and attempt to fix the problem.

About this task

To diagnose why a backup failed verification:

1. Verify that the **AC_CONFIG** environment variable and the contents of the archecker configuration file are set correctly. If these variables are set incorrectly, the ON-Bar activity log prints a message.
2. Back up the data onto different media.

Do not reuse the original backup media because it might be damaged.

Do not use any backups based on this backup. If the level-0 backup failed verification, do not use the corresponding level-1 and level-2 backups.

3. Verify this new backup.
If verification succeeds, you can restore the storage spaces.
4. Use your storage manager to expire the backup that failed verification and then run the `onsmsync` utility without arguments to remove the bad backup from the **sysutils** and emergency boot files.
5. If verification fails again, call IBM® Software Support and provide them with the following information:
 - Your backup tool name (ON-Bar)
 - The database server `online.log`

- The archecker message log
- The AC_STORAGE directory that contains the bitmap of the backup and copies of important backed-up pages

If only part of the backup is corrupted, IBM® Software Support can help you determine which portion of the backup can be restored in an emergency.

IBM® Software Support might advise you to run oncheck options against a set of tables.

Verifying an expired backup

You can verify an expired backup in case subsequent backups are not valid.

About this task

To verify an expired backup:

1. Check that the status of the backup save set on the storage manager. If the storage manager expired the backup save set, the archecker utility cannot verify it.
2. Use the storage-manager commands for activating the expired backup save set. See your storage-manager documentation.
3. Run the onbar -v command again.

Related reference

[onbar -v syntax: Verifying backups on page 58](#)

Restoring when a backup is missing data

If a backup fails verification because of missing data, you can perform a restore from an older backup.

About this task

To restore when a backup is missing data:

1. Choose the date and time of an older backup than the one that failed. To perform a point-in-time verification, use the `onbar -v -t time space` command.
2. If the older backup passes verification, perform a point-in-time physical restore by using the same *time* value, then perform a log restore, as follows:

```
onbar -r -p -t time space
onbar -r -l
```

3. Expire the corrupted backup at your storage manager.
4. Run the onsmsync command without arguments.

The onsmsync utility removes backups that are no longer held by the storage manager from the emergency boot file and the **sysutils** database, preventing ON-Bar from attempting to use such backups.

Related reference

[The onmsync utility on page 107](#)

[onbar -r syntax: Restoring data on page 66](#)

Restore data with ON-Bar

You can use the ON-Bar utility to restore data that was backed up by the ON-Bar utility.

Before you restore data, use the pre-restore checklist to determine if whether a restore is needed and to prepare for a restore.

To perform a restore with the ON-Bar utility:

1. Make the storage devices that were available during the backup available for the restore.
2. If necessary, add enough temporary space to perform the restore. The logical log restore portion of a warm restore requires temporary space. The minimum amount of temporary space is equal to the smaller of the total amount of allocated logical-log space and the number of log files to be replayed.
3. Run the `onbar -r` command with the appropriate options to restore the data.
4. Monitor the ON-Bar activity log.
5. After the restore is complete, run the `onstat -d` command to verify that all storage spaces are restored. The letter O in the flags column indicates that the chunk is online.

Pre-restore checklist

Use this checklist to determine if a restore is necessary and to prepare for a restore.

To prepare for a restore:

- Determine if you need to restore. If one or more of these problems is true, you perform a restore to fix the problem:
 - Has data been lost or corrupted?
 - Does a committed transaction error need to be undone?
 - Is the database server down or has a disk failed?
 - Is a storage space or chunk down or inconsistent?
- Diagnose the problem by using database server monitoring tools.
- If the root dbspace or the dbspaces that contain the physical log and logical-log files need to be restored, you must perform a cold restore. The database server must be offline during a cold restore. Ask your client users to log off the system.
- Contact the appropriate vendor to resolve the following types of problems before doing a restore:
 - The storage manager
 - The XBSA connection
 - The operating system
 - The storage media

Related information[Database server monitoring on page](#)

Storage space status and required actions

To determine the state of each storage space and its chunks, examine the output of the `onstat -d` command. The storage space status determines the action you need to take to solve the problem. The database server must be online.

The following table describes `onstat -d` command output about chunk status and the actions required to solve the problems. The chunk status information is in the second position of the **flags** column in the first (storage spaces) and second (chunks) sections of the output.

Table 12. Chunk flag descriptions and required actions

chunk flag	Storage space or chunk state	Action required
(No flag)	Storage space no longer exists.	Perform a point-in-time cold restore to a time before the storage space was dropped.
D	Chunk is down or storage space is disabled.	Perform a warm restore of the affected storage space.
I	Chunk is physically restored, but needs a logical restore.	Perform a logical restore.
L	Storage space is being logically restored.	Try the logical restore again.
N	Chunk is renamed and either down or inconsistent.	Perform a warm restore of the chunk when the physical device is available.
O	Chunk is online.	No action required.
P	Storage space is physically restored.	Perform a logical restore, if one is not already in progress.
R	Storage space is being restored.	Perform a physical or logical restore.
X	Storage space or chunk is newly mirrored.	No action required.

Related reference[onbar -r syntax: Restoring data on page 66](#)**Related information**[onstat -d command: Print chunk information on page](#)

Restore save sets with ISM

If you are using IBM® Informix® Storage Manager (ISM), you can restore data from save sets on the storage volume. When the ISM server receives a restore request, the `ism_watch` command prompts you to mount the required storage volume on the storage device. When you mount the volume, the restore resumes.

You can set the retention period for either a save set or volume. Unless all the save sets on the volume are expired, you can use ON-Bar to restore it.

After the retention period for a save set expires, ON-Bar can no longer restore it. To recreate an expired save set, use the `ism_catalog -recreate` from command.

If you set the retention period for a volume, ISM retains the save sets until all the save sets on that volume are expired. To recover an expired volume, use the `ism_catalog -recover` from command.

Storage device availability

Verify that the storage devices and files used in the backup are available for the restore.

If you drop a dbspace or mirror device after a level-0 backup, the dbspace or mirror device must be available to the database server when you begin the restore. If the storage device is not available, the database server cannot write to the chunk and the restore fails.

If you add a chunk after your last backup, the chunk device must be available to the database server when it rolls forward the logical log.

onbar -r syntax: Restoring data

To run a complete restore, use `onbar -r` command.

To run ON-Bar commands, you must be user **root**, user **informix**, a member of the **bargroup** group on UNIX™, or a member of the Informix®-Admin group on Windows™.

- [Usage on page 72](#)
- [Example: Perform a whole-system restore on page 73](#)
- [Example: Restore specific storage spaces on page 73](#)
- [Example: Perform a warm restore in stages on page 73](#)
- [Example: Point-in-time restore on page 73](#)
- [Example: Point-in-time restore with a French locale on page 74](#)
- [Example: Point-in-time restore in stages on page 74](#)
- [Example: Restore a dropped storage space and chunks on page 74](#)
- [Example: Restore critical files on page 74](#)

```
Run a full or physical restore
>>-onbar-- -r-+-----+-----+-----+-----+-----+-----+-----+----->
          '- -p-'  +- -e-----+  '- -w-'
                    +- -encrypt-+
                    '- -decrypt-'
```

```

>+-----+-----+-----+-----+----->
+- -t-"time"+   '- -pw+-----+--'
'- -n--log----'      '--filename--'

>+-----+-----+-----+-----+-----+-->
+- -O-----+-----+-----+-----+-----+
'-+- -rename-- -f--filename-----+--'
| .-----+-----+-----+-----+-----|
| V          |          |          |          |
'|--- -rename-- -p--old_path-- -o--old_offset-- -n--new_path-- -o--new_offset--+-'

>+-----+-----+-----+-----+-----><
+- -f--filename+   '- -cf--+-yes--+-'
| .-----+-----+-----+-----+-----|
| V          |          |          |          |
'|---space+-----+'

```

Run a logical restore

```
onbar -r -l [{ -c | -x | -t "time" | -n log }]
```

Run a tenant restore

```
onbar -r -T tenant_database [{ -encrypt | -decrypt } ] -t "time" [-O]
```

Table 13. Options for the onbar -r command.

Option	Description
-r	<p>Specifies a restore. If the database server is offline, ON-Bar performs a cold restore. If the database server is in online, quiescent, or fast recovery mode, ON-Bar performs a warm restore.</p> <p>In a cold restore, the -r option restores all storage spaces and salvages and restores the logical logs. In a warm restore, the -r option restores all storage spaces that are offline and restores the logical logs.</p> <p>You must specify the -r option first.</p>
space	<p>Specifies which storage spaces to back up as a list of one or more dbspace, blobspace, or sbspace names, separated by blank spaces.</p> <p>ON-Bar restores only the storage spaces listed. If the database server is offline, you must list all the critical dbspaces. You cannot specify temporary spaces.</p>
-C	Continuously restores logical logs from the current logical log tape without sending prompts to mount the tape.

Table 13. Options for the onbar -r command. (continued)

Option	Description
	<p>The server is placed in suspend log restore state, and the command exits after the last applicable log is restored. The server sends a prompt if a log spans tapes. The configuration parameter <code>RESTARTABLE_RESTORE</code> does not affect continuous log restoration.</p>
-cf	<p>Specifies whether the critical files are restored during a cold restore.</p> <p>The critical files are the <code>onconfig</code> file, the <code>sqlhosts</code> file (on UNIX™), the <code>oncfg_servername.servernum</code> file, and the <code>ixbar.servernum</code> file.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> • <code>yes</code> = Restores the critical files. • <code>no</code> = Default. Does not restore the critical files. • <code>only</code> = Restores only the critical files.
-decrypt	<p>Specifies to decrypt any encrypted storage spaces during the physical restore of the spaces.</p>
-encrypt	<p>Specifies to encrypt storage spaces during the physical restore of the spaces. Storage space encryption must be enabled by the <code>DISK_ENCRYPTION</code> configuration parameter. Otherwise, the restore fails.</p> <p>For more information about storage space encryption, see Changing storage space encryption during a restore on page 51.</p>
-e	<p>Specifies an external restore. Run the <code>onbar -r -e</code> command after you externally restore the storage spaces. Marks storage spaces as physically restored, restores the logical logs, and brings the storage spaces online.</p>
-f <i>filename</i>	<p>Specifies the path and file name of a text file that lists the storage spaces to restore or rename.</p> <p>Use this option to avoid entering a long list of storage spaces.</p> <p>For more information, see List of storage spaces in a file on page 51.</p>
-l	<p>Specifies a logical restore only. Restores and rolls forward the logical logs. The logical restore applies only to those storage spaces that are already physically restored.</p>

Table 13. Options for the onbar -r command. (continued)


Option	Description
	<p> Important: To improve performance, replay logical-log transactions in parallel during a warm restore. Use the ON_RECVRY_THREADS configuration parameter to set the number of parallel threads. To replay logical-log transactions in parallel during a cold restore, use the OFF_RECVRY_THREADS configuration parameter. For more information, see your <i>HCL® Informix® Performance Guide</i>.</p>
-n <i>log</i>	<p>Indicates the unique ID of the last logical log to be restored in a cold restore. The database server must be offline.</p> <p>To find the unique ID, use the onstat -l command.</p> <p>A point-in-log restore is a special point-in-time restore. You must restore all storage spaces in a point-in-log restore so that the data is consistent. If any logical logs exist after the specified log, ON-Bar does not restore them and their data is lost. If the specific logical log applies to more than one timeline, ON-Bar uses the latest one.</p> <p>Cannot be combined with the -t option.</p>
-n <i>new_path</i>	<p>Specifies the new path of the chunk. Use with the -rename option.</p>
-O	<p>Overrides internal error checks. Allows the restore of online storage spaces. Forces the recreation of chunk files that no longer exist.</p> <p>Used to override internal error checks to perform the following tasks:</p> <ul style="list-style-type: none"> • Force the restore of online storage spaces. If a storage space in the list of storage spaces to restore is online, ON-Bar takes the storage space offline and then restores it. If this operation succeeds, ON-Bar completes with an exit code of 177. • Force the creation of nonexistent chunk files. If a chunk file for a storage space being restored no longer exists, ON-Bar recreates it. The newly created chunk file is cooked disk space, not raw disk space. If ON-Bar successfully recreates the missing chunk file, ON-Bar completes with an exit code of 179. • Force a cold restore to proceed if a critical storage space is missing. In a cold restore, ON-Bar checks whether every critical space is being restored. This check occasionally causes false warnings. If the warning

Table 13. Options for the onbar -r command. (continued)

Option	Description
	was valid, the restore fails. If the warning was false and ON-Bar successfully restores the server, ON-Bar completes with an exit code of 115.
	Use the -O option with a whole-system restore only to recreate missing chunk files. You cannot use the onbar -r -w -O command when the database server is online because the root dbspace cannot be taken offline during the whole-system restore. Cannot be combined with the -rename option.
<code>-pw [file name]</code>	The -pw option is required only when the storage space encryption feature is enabled and no stash file is in use. Supply an optional path to a file containing the keystore password, otherwise onbar will prompt for a password before performing the restore.
<code>-o new_offset</code>	Specifies the offset of the renamed chunk. Use with the -rename option.
<code>-o old_offset</code>	Specifies the offset of the chunk to be renamed. Use with the -rename option.
<code>-p</code>	<p data-bbox="613 987 966 1014">Specifies a physical restore only.</p> <p data-bbox="613 1050 1412 1239">You must follow a physical restore with a logical restore before data is accessible unless you use a whole-system restore. This option turns off automatic log salvage before a cold restore. If the LTAPEDev configuration parameter is set to <code>/dev/null</code> or <code>NUL</code>, you must use the -p option during a restore.</p>
<code>-p old_path</code>	Specifies the path of the chunk to be renamed. Use with the -rename option.
<code>-rename</code>	<p data-bbox="613 1333 1445 1407">Renames one or more chunks during a cold restore. The database server must be offline.</p> <p data-bbox="613 1438 1461 1596">This option is useful if you need to restore storage spaces to a different disk from the one on which the backup was made. You can rename any type of chunk, including critical chunks and mirror chunks. You can rename chunks that have level-0 backups.</p> <p data-bbox="613 1627 1031 1659">Cannot be combined with the -O option.</p>
<code>-T tenant_database</code>	<p data-bbox="613 1690 1404 1764">Restores a tenant database. The database server must be online. No other warm restores or tenant restores can be in progress.</p> <p data-bbox="613 1795 1445 1871">If you include the -t option, the data is restored to the specified point in time. If you do not include the -t option, the data is restored to the current time.</p>

Table 13. Options for the onbar -r command. (continued)

Option	Description
	<p>Include the -O option unless all of the permanent tenant storage spaces are marked as down. Temporary storage spaces are never backed up or restored.</p> <p>For more information, see Restoring a tenant database to a point in time on page.</p>
-t "time"	<p>Specifies the time of the last transaction to be restored from the logical logs in a cold restore. The database server must be offline. Specifies the time of the last transaction to be restored from the logical logs in a cold restore or a tenant database point-in-time restore. For a cold restore, the database server must be offline. For a tenant database point-in-time restore, the database server must be online.</p> <p>All storage spaces specified are restored to the same point in time. However, for a cold restore, if you perform a physical restore followed by a logical restore, the logical restore can be to a later point in time. For example you might detect that your current backup is corrupted, and that you need to restore the previous backup. In this case, start your physical restore with the timestamp of your previous backup, and subsequently start the logical recovery to a more recent timestamp.</p> <p>A point-in-time restore is typically used to recover from a mistake. For example, if you accidentally dropped a database, you can restore the server to a point in time just before you dropped the database.</p> <p>To determine the appropriate date and time for the point-in-time restore, use the onlog utility. The onlog utility output shows the date and time of the committed transactions in the logical log. All data transactions that occurred after the time you specify in the restore command are lost.</p> <p>Use quotation marks around the date and time. The format for the English locale is <code>yyyy-mm-dd hh:mm:ss</code>. If the GL_DATETIME environment variable is set, you must specify the date and time according to that variable.</p> <p>Cannot be combined with the -n log option.</p>
-w	<p>Performs a whole-system restore of all storage spaces and logical logs from the last whole-system backup. The database server must be offline.</p> <p>After the whole-system restore is complete, the server is in quiescent mode.</p>

Table 13. Options for the onbar -r command. (continued)

Option	Description
	If you specify onbar -r -w without a whole-system backup, return code 147 is returned because ON-Bar cannot find any storage spaces backed up as part of a whole-system backup.
-X	Stops continuous logical log restore. Leaves the server in quiescent mode in a logical restore suspend state without restoring additional logs.

Usage

You can restore storage spaces stored in both raw and cooked files. If your system contains primary and mirror storage spaces, ON-Bar writes to both chunks simultaneously during the restore, except for an external restore. You cannot specify to restore temporary spaces. When you restore the critical dbspaces (for example, the root dbspace), the database server recreates the temporary dbspaces, but they are empty.

ON-Bar restores the storage spaces in parallel if the `BAR_MAX_BACKUP` or `BAR_MAX_RESTORE` configuration parameter is set to a value greater than 1. To speed up restores, you can add additional CPU virtual processors.

You can restore noncritical storage spaces in a warm restore, when the database server is online, in the following circumstances:

- The storage space is online, but one of its chunks is offline, recovering, or inconsistent.
- The storage space is offline or down.

You cannot perform more than one warm restore simultaneously. If you need to restore multiple storage spaces, specify the set of storage spaces to restore to ON-Bar or allow ON-Bar to restore all down storage spaces by not explicitly specifying any spaces.



Tip: For faster performance in a restore, assign separate storage devices for backing up storage spaces and logical logs. If physical and logical backups are mixed together on the storage media, it takes longer to scan the media during a restore.

In certain situations, you might want to perform a restore in stages. If multiple devices are available for the restore, you can restore multiple storage spaces separately or concurrently, and then perform a single logical restore.

By default, ON-Bar restores the latest backup. If you do not want to restore the latest backup, you can restore from an older backup: for example, when backup verification failed or the backup media was lost. You can perform a point-in-time restore or a point-in-log restore. Alternatively, you can expire a bad backup in the storage manager, run the `onsmsync` command, and then restore from the older backup. If you accidentally drop a storage space, you can use a point-in-time restore or a point-in-log restore to recover it.

You can force a restore of online storage spaces (except critical dbspaces) by using the `-O` option. The database server automatically shuts down each storage space before it starts to restore it. Taking the storage space offline ensures that users do not try to update its tables during the restore process.

You can restore critical files during a cold restore by including the `-cf yes` option.

You can rename chunks by specifying new chunks paths and offsets during a cold restore with ON-Bar. This option is useful if you need to restore storage spaces to a different disk from the one on which the backup was made. You can rename any type of chunk, including critical chunks and mirror chunks.

When storage space encryption is enabled, by default storage spaces retain the same encryption state after a restore as during the back up. You can specify to encrypt or decrypt a storage space during a restore with the `-encrypt` or `-decrypt` options.

Example

Example: Perform a whole-system restore

A whole-system restore is a cold restore and must be performed while the server is offline. The following command restores a whole-system backup:

```
onbar -r -w
```

Example

Example: Restore specific storage spaces

The following example restores two specific storage spaces, `fin_dbspace1` and `fin_dbspace2`:

```
onbar -r fin_dbspace1 fin_dbspace2
```

Example

Example: Perform a warm restore in stages

The following commands perform a physical restore, back up logical logs, and perform a logical restore:

```
onbar -r -p
onbar -b -l
onbar -r -l
```

Example

Example: Point-in-time restore

The following command restores database server data to its state at a specific date and time:

```
onbar -r -t "2011-05-10 11:35:57"
```

In this example, the restore replays transactions that committed on or before the specified time, including any transactions with a commit time of 11:35:57. Transactions in progress but not committed by 11:35:57 are rolled back.

Example

Example: Point-in-time restore with a French locale

The default date and time format for the French locale, `fr_fr.8859-1`, uses the format `"%A %d %B %iY %H:%M:%S."`

The following command restores the data to a specific point in time that is formatted for the French locale:

```
onbar -r -t "Lundi 6 Juin 2011 11:20:14"
```

You can set `GL_DATETIME` to a different date and time format that uses the date, month, two-digit year, hours, minutes, and seconds. For example:

```
%.1d %B %iy %H:%M:%S
```

The following command restores to a specific point in time by specifying a two-digit year for the French locale:

```
onbar -r -t "6 Juin 11 11:20:14"
```

Example

Example: Point-in-time restore in stages

The following commands perform a physical restore and a logical restore to the same point in time:

```
onbar -r -p -t "2011-05-10 11:35:57"
onbar -r -l -t "2011-05-10 11:35:57"
```

Example

Example: Restore a dropped storage space and chunks

Suppose that a transaction dropped a storage space named `dbspace1` and deleted chunks at the time 2011-05-10 12:00:00. The following command restores the storage space and recreates the deleted chunks while the server is offline:

```
onbar -r -t "2011-05-10 11:59:59" -O
```

Example

Example: Restore critical files

The following command restores data and the critical files during a cold restore:

```
onbar -r -cf yes
```

Related reference

[Storage space status and required actions on page 65](#)

[External restore commands on page 98](#)

[ON-Bar security on page 40](#)

[onbar -RESTART syntax: Restarting a failed restore on page 87](#)

Related information

[Restoring when a backup is missing data on page 63](#)

[Replacing disks during a restore on page 81](#)

Avoid salvaging logical logs

The `onbar -r` command automatically salvages the logical logs. However, avoid salvaging logical logs in some situations.

Use the `onbar -r -p` and `onbar -r -l` commands to skip log salvage.

If you set the `LTAPEDEV` configuration parameter to `/dev/null` on UNIX™ or to `NUL` on Windows™, the logical logs are not salvaged in any ON-Bar restore (`onbar -r` or `onbar -r -w`, for example).

Avoid salvaging the logical logs in the following situations:

- When you perform an imported restore
 - Salvage the logical logs on the source database server but not on the target database server.
- If you reinitialize the database server (`oninit -i`) before you perform a cold restore
 - Reinitialization creates new logical logs that do not contain the data that you want to restore.
- If you install a new disk for the `dbospace` that contains the logical logs
 - Salvage the logs from the old disk, but not from the new disk.

Related information

[Restoring to a different computer on page 83](#)

Performing a cold restore

If a critical storage space is damaged because of a disk failure or corrupted data, you must perform a cold restore. If a disk fails, you need to replace it before you can perform a cold restore to recover data.

About this task

If you try to perform a cold restore without a backup, data in the storage spaces that were not backed up are lost.

To perform a cold restore:

1. Shut down the server by running the `onmode -ky` command.
2. If the disk that contains the logical-log files must be replaced or repaired, use the `onbar -b -l -s` command to salvage logical-log files on the damaged disk.
 - Otherwise, ON-Bar automatically salvages the logical logs.
3. If necessary, repair or replace the damaged disk.
4. If the files in `INFORMIXDIR` are damaged, copy the back ups of administrative files to their original locations.

Otherwise, you do not need to copy the administrative files.

5. Restore the critical and noncritical storage spaces by running the onbar -r command.
When the restore is complete, the database server is in quiescent mode.
6. Start the server by running the onmode -m command.
7. Synchronize the storage manager by running the onsmsync command.

Related reference

[The onsmsync utility on page 107](#)

Related information

[onmode -k, -m, -s, -u, -j: Change database server mode on page](#)

Configuring a continuous log restore by using ON-Bar

Use continuous log restore to keep a second system (hot backup) available to replace the primary system if the primary system fails.

Before you begin

The version of HCL Informix® must be identical on both the primary and secondary systems.

To configure continuous log restore by using ON-Bar:

1. On the primary system, perform a level-0 backup with the onbar -b -L 0 command.
2. Import the backup objects that were created to the storage manager of the secondary server.
3. On the secondary system, perform a physical restore with the onbar -r -p command.
After the physical restore completes on the secondary system, the database server waits in fast recovery mode to restore logical logs.
4. On the primary system, back up logical logs with the onbar -b -l command.
5. Transfer the backed up logical logs to the secondary system and restore them with the onbar -r -l -C command.
6. Repeat steps [4 on page 76](#) and [5 on page 76](#) for all logical logs that are available to back up and restore.
7. If you are doing continuous log restore on a secondary system as an emergency standby, run the following commands to complete restoring logical logs and quiesce the server:
 - If logical logs are available to restore, run the onbar -r -l command.
 - After all available logical logs are restored, run the onbar -r -l -X command.

Related reference

[onbar -b syntax: Backing up on page 46](#)

Related information

[Continuous log restore on page 10](#)

Restoring data by using a mixed restore

You can use mixed restore to reduce the time until urgent data becomes online and available when you need to restore the server. Urgent data is data that you deem as critical to your business operation.

About this task

In a mixed restore, you first perform a cold restore of the critical dbspaces (the root dbspace and the dbspaces that contain the physical and logical logs) and the dbspaces containing your urgent data. Because you do not restore all dbspaces, you can bring the server online faster. You then restore the remaining storage spaces in one or more warm restores.

To perform a mixed restore:

1. Shut down the database server by running the `onmode -ky` command.
2. Perform a cold restore of the critical and urgent dbspaces by running the `onbar -r` command with the list of critical and urgent dbspace names.
You can specify a point in time to restore from an older backup.
3. Start the server by running the `onmode -m` command.
4. Synchronize the storage manager by running the `onsmsync` command.
5. Perform a warm restore of the remaining storage spaces by running the `onbar -r` command.
You can perform multiple warm restores to prioritize certain storage spaces.

Example

Examples

Example 1: Simple mixed restore

A database server has five dbspaces in addition to the root dbspace: **logdbs**, **dbs_1**, **dbs_2**, **dbs_3**, and **dbs_4**. The logical logs are stored in **logdbs** and the physical log is in the root dbspace. The critical dbspaces that must be restored during the initial cold restore are **rootdbs** and **logdbs**. The dbspace that contains urgent data is **dbs_1**. The following commands shut down the database server, perform a cold restore on the critical and urgent dbspaces, and restart the database server:

```
onmode -ky
onbar -r rootdbs logdbs dbs_1
onmode -m
```

After the database server starts, any data stored in **rootdbs**, **logdbs**, and **dbs_1** dbspaces is accessible.

The following commands synchronize the storage manager and perform a warm restore of the remaining dbspaces, **dbs_2**, **dbs_3**, and **dbs_4**:

```
onsmsync
onbar -r
```

Example 2: Point-in-time mixed restore

The following commands perform a cold restore for a subset of the storage spaces (including all critical dbspaces) in the initial cold restore, perform a warm restore for **dbspace_2** and **dbspace_3**, followed by a warm restore of **dbspace_4** and **dbspace_5**, and finally perform a warm restore of all remaining storage spaces:

```
onbar -r -t "2011-05-10 11:35:57" rootdbs logspace_1 dbspace_1
onmode -m
onsmsync
onbar -r dbspace_2 dbspace_3
onbar -r dbspace_4 dbspace_5
onbar -r
```

Strategies for using a mixed restore

To implement a mixed-restore strategy, carefully select the set of dbspaces in which you place your databases and database objects when you create them.

ON-Bar backs up and restores physical, not logical, entities. Thus, ON-Bar cannot restore a particular database or a particular set of tables. Instead, ON-Bar restores a particular set of storage spaces. It is up to you to track what is stored in those storage spaces.

For example, consider a database with the catalogs in the dbspace **cat_dbs**:

```
create database mydb in cat_dbs with log;
```

A table in this database is fragmented among the dbspaces **tab_dbs_1** and **tab_dbs_2**:

```
create table mytab (i integer, c char(20))
fragment by round robin in tab_dbs_1, tab_dbs_2;
```

An index for the table is stored in the dbspace **idx_dbs**:

```
create index myidx on mytab(i) in idx_dbs;
```

If you need to restore the server, you cannot access all of the data in the example database until you restore the dbspaces containing the database catalogs, table data, and index: in this case, the dbspaces **cat_dbs**, **tab_dbs_1**, **tab_dbs_2**, and **idx_dbs**.

To simplify the management and tracking of your data, divide your set of dbspaces into subsets in which you store data of a particular urgency. When you create your database objects, place them in dbspaces appropriate to their urgency. For example, if you have data with three levels of urgency, you might want to place all the objects (database catalogs, tables, and indexes) associated with your most urgent data in a particular set of dbspaces: for example, **urgent_dbs_1**, **urgent_dbs_2**, ... **urgent_dbs_n**. You would place all the objects associated with less urgent data in a different set of dbspaces: for example, **less_urgent_dbs_1**, **less_urgent_dbs_2**, ... **less_urgent_dbs_k**. Lastly, you would place your remaining data in a different set of dbspaces: for example, **non_urgent_dbs_1**, **non_urgent_dbs_2**, **non_urgent_dbs_r**.

If you need to restore the server, you would first perform a cold restore of all critical dbspaces and dbspaces containing urgent data, **urgent_dbs_1** through **urgent_dbs_n**. For example, assume that logical logs are distributed among two dbspaces, **logdbsp_1** and **logdbsp_2**, and the physical log is in **rootdbs**. The critical dbspaces are therefore **rootdbs**, **logdbsp_1**, and **logdbsp_2**.

You would perform the initial cold restore by issuing the following ON-Bar command:

```
onbar -r rootdbs logdbsp_1 logdbsp_2 urgent_dbs_1 ... urgent_dbs_2
```

You can bring the server online and all business-urgent data is available.

Next, perform a warm restore for the less-urgent data:

```
onmsync
onbar -r less_urgent_dbs_1 less_urgent_dbs_2 ..... less_urgent_dbs_k
```

Finally, you can perform a warm restore for the rest of the server by issuing the following command.

```
onbar -r
```

In a larger system with dozens of dbspaces, you can divide the warm restore portion of the mixed restore into several warm restores, each restoring only a small subset of the dbspaces remaining to be restored in the system.

Recreating chunk files during a restore

If the disk or file system fails, one or more chunk files might be missing from the dbspace. Use the `-O` option to recreate missing chunk files and any necessary directories during a restore.

About this task

The restore fails if insufficient space exists on the file system. The newly created chunk files are cooked files and are owned by group **informix** on UNIX™ or group **Informix-Admin** on Windows™.

Restoring when using cooked chunks

You can recreate missing cooked chunk files during a restore.

About this task



Restriction: ON-Bar does not recreate chunk files during a logical restore if the logical logs contain chunk-creation records.

To restore when using cooked chunks:

1. Install the new disk.
2. Based on your system, perform one of the following tasks:
 - Choose from:**
 - On UNIX™, mount the device as a file system.
 - On Windows™, format the disk.
3. Allocate disk space for the chunk file.
4. Run the `onbar -r -O space` command to recreate the chunk files and restore the dbspace.

Related information

[Allocating cooked file spaces on UNIX on page](#)

Restoring when using raw chunks

You can recreate missing raw chunk files during a restore.

To restore when using raw chunks:

1. Install the new disk.
2. For UNIX™, if you use symbolic links to raw devices, create new links for the down chunks that point to the newly installed disk.
ON-Bar restores the chunk file to where the symbolic link points.
3. Issue the `onbar -r space` command to restore the dbspace.

Related information

[Allocating raw disk space on UNIX on page](#)

[Allocating raw disk space on Windows on page](#)

Reinitializing the database server and restoring data

Reinitializing disk space destroys all existing data managed by the database server. However, you can restore data from a backup that was performed before reinitialization.

Before you begin

You must have a current, level-0 backup of all storage spaces.

About this task

During initialization, ON-Bar saves the emergency boot file elsewhere and starts a new, empty emergency boot file. Therefore, any backups that you performed before reinitializing the database server are not recognized. You must use the copy of the emergency boot file you saved before initialization to restore the previous database server instance.

To reinitialize the database server and restore the old data:

1. Copy the emergency boot file, the `oncfg` file, and the `onconfig` file to a different directory.
2. Set the `FULL_DISK_INIT` configuration parameter to 1 in the `onconfig` file.
3. Shut down the database server.
4. Reinitialize the database server by running the `oninit -i` command.
5. Move the administrative files into the database server directory.
If the administrative files are unavailable, copy them from the last backup into the database server directory.
6. Perform a restore by running the `onbar -r -p -w` command.

Do not salvage the logical logs.

7. Verify that you restored the correct instance of the critical and noncritical storage spaces.

Related information

[The oninit utility on page](#)

Replacing disks during a restore

You can replace disks during a restore by renaming chunks. You rename chunks by specifying new chunks paths and offsets during a cold restore with ON-Bar. This option is useful if you need to restore storage spaces to a different disk from the one on which the backup was made. You can rename any type of chunk, including critical chunks and mirror chunks.

Before you begin

The old chunk must be included in the last level-0 backup.

The following guidelines apply to new chunks:

- The new chunk does not need to exist. You can install the new chunk later and perform a warm restore of a storage space containing it. If you specify a nonexistent chunk, ON-Bar records the rename information in the chunk reserved pages, but does not restore the data. The renamed (but not restored) chunks have a status of offline, designated by an `N` flag in the output of the `onstat -d` command.
- The new chunk must have the correct permissions.
- The new chunk must be included in the last level-0 backup.
- The new chunk path name and offset cannot overlap existing chunks.

About this task



Tip: If you use symbolic links to chunk names, you might not need to rename chunks; you only need to edit the symbolic name definitions.

To rename chunks during a restore:

1. Shut down the database server.
2. Run the `onbar -r` command with the `-rename` option and the chunk information options.
If you are renaming the primary root or mirror root chunk, ON-Bar updates the values of the `ROOTPATH` and `ROOTOFFSET`, or `MIRRORPATH`, and `MIRROROFFSET` configuration parameters. The old version of the `onconfig` file is saved as `$ONCONFIG.localtime`.
3. Perform a level-0 archive so that you can restore the renamed chunks.

Example

Examples

The following table lists example values for two chunks that are used in the examples in this section.

Element	Value for first chunk	Value for second chunk
Old path	/chunk1	/chunk2
Old offset	0	10000
New path	/chunk1N	/chunk2N
New offset	20000	0

Example 1: Rename chunks by supplying chunk information in the command

The following command renames the chunks **chunk1** to **chunk1N** and **chunk2** to **chunk2N**:

```
onbar -r -rename -p /chunk1 -o 0 -n /chunk1N -o 20000
      -rename -p /chunk2 -o 10000 -n /chunk2N -o 0
```

Example 2: Rename chunks by supplying chunk information in a file

Suppose that you have a file named `listfile` that has the following contents:

```
/chunk1 0 /chunk1N 20000
/chunk2 10000 /chunk2N 0
```

The following command renames the chunks **chunk1** to **chunk1N** and **chunk2** to **chunk2N**:

```
onbar -r -rename -f listfile
```

Related reference

[onbar -r syntax: Restoring data on page 66](#)

[onbar -b syntax: Backing up on page 46](#)

Renaming a chunk to a nonexistent device

To rename a chunk to a nonexistent device, specify the new path name, but restore the storage spaces after you install the physical device. This option is useful if you need to rename a chunk and it is convenient to perform a cold restore before you install the new device. When the new chunk device is ready, you can perform a warm restore of a storage space onto it.

About this task

You can combine renaming chunks with existing devices and renaming chunks with nonexistent devices in the same rename operation. This example shows how to rename a single chunk to a nonexistent device name.

The following table lists example values for the chunks used in this example.

Storage space	Old chunk path	Old offset	New chunk path	New offset
sbspace1	/chunk3	0	/chunk3N	0

To rename a chunk to a nonexistent device:

1. Rename the chunk with the following command: `onbar -r -rename -p /chunk3 -o 0 -n /chunk3N -o 0`
2. When you see the following prompt, enter `y` to continue:

```
The chunk /chunk3N does not exist. If you continue, the
restore may fail later for the dbspace which contains this chunk.
Continue without creating this chunk? (y/n)
```

The chunk `/chunk3` is renamed to `/chunk3N`, but the data is not yet restored to `/chunk3N`.

3. Perform a level-0 archive.
4. Add the physical device for `/chunk3N`.
5. Perform a warm restore of **sbspace1** with the `onbar -r sbspace1` command.
6. Perform a level-0 archive.

Restoring to a different computer

You can back up data on one computer and restore the data on a different computer. Importing a restore is useful for disaster recovery or upgrading a database server. After you back up your data and move over the storage-manager objects, you can perform an imported restore. An imported restore involves copying files from the source to the target computer and performing the restore in one of several ways.

Before you begin

Prerequisites:

- Your storage manager must support imported restores.
- A whole-system backup must include all storage spaces; logical logs are optional.

The level-0 backup must include all storage spaces and logical logs.

- Both the source and target computers must be on the same LAN or WAN and must have the following attributes:
 - Identical hardware and operating systems
 - Identical database server versions and editions
 - The same configuration and ROOTPATH information, although the server names and numbers can differ.
 - Identical storage-manager versions
 - Compatible XBSA libraries

About this task



Important: Every chunk (including mirrors) must match exactly in size, location, and offset on the source and target computers for the imported restore to complete.

To perform the imported restore:

1. Install the database server and the storage manager on the target computer.
2. Set up the storage manager on the target database server instance.
 - a. Set the necessary environment variables.
 - b. Define the same type of storage devices as on the source instance.
 - c. Label the storage media with the correct pool names.
 - d. Mount the storage devices.
 - e. Update the `sm_versions` file on the target computer with the storage-manager version.
3. Be sure that the target computer has the devices and links in place for the chunks that match the devices and links on the source computer
4. Perform a level-0 backup (`onbar -b` or `onbar -b -w`) of all storage spaces on the source database server.



Restriction: Do not perform an incremental backup.

5. If you are using IBM® Informix® Storage Manager (ISM), follow these steps:
 - a. Shut down the storage manager on both computers.
 - b. Create a tar file of the storage-manager directories on the source computer.
 - c. Copy this tar file and unpack it on the target computer.

With other storage managers, you might be able to use backup tapes or import the storage-manager directories over the network. For more information, see your storage-manager documentation.

6. Mount the transferred storage volumes.

Choose from:

- If the backup files are on disk, copy them from the source computer to the target computer.
- If the backup is on tapes, mount the transferred volumes on the storage devices that are attached to the target computer. Both the source and target computers must use the same type of storage devices such as 8-mm tape or disk.
- If the backup is on the backup server, retrieve the backup from that backup server.

7. Use storage-manager commands to add the source host name as a client on the target computer.
8. Copy the following files from the source computer to the target computer.

Table 14. Administrative files to copy

File	Action
Emergency boot file	Rename the emergency boot file with the target database server number. For example, rename <code>ixbar.51</code> to <code>ixbar.52</code> . The emergency boot file needs only the entries from the level-0 backup on the source computer. The file name is <code>ixbar.servernum</code> .

File	Action
The <code>oncfg</code> files: <code>oncfg_ser</code> <code>vername.s</code> <code>ervernum</code>	ON-Bar needs the <code>oncfg</code> file to know what dbspaces to retrieve. Rename the <code>oncfg</code> file with the target database server name and number. For example, rename <code>oncfg_bostonserver.51</code> to <code>oncfg_chicagoserver.52</code> . The file name must match the <code>DBSERVERNAME</code> and <code>SERVENUM</code> on the target computer.
The <code>onconfig</code> file	In the <code>onconfig</code> file, update the <code>DBSERVERNAME</code> and <code>SERVENUM</code> parameters with the target database server name and number.
Storage-manager configuration files, if any	The storage-manager configuration files might need updating.

9. Restore the data in one of the following ways:

Table 15. Restore data options

Option	Action
If you did not start the Informix® instance on the target server	Use the <code>onbar -r</code> command to restore the data.
If you are importing a whole-system backup	Use the <code>onbar -r -w -p</code> command to restore the data.
If you started the Informix® instance on the target server.	Restore the data in stages: <ol style="list-style-type: none"> a. Use the <code>onbar -r -p</code> command to restore the physical data. b. Use the <code>onbar -r -l</code> command to restore the logical logs. <p>This process avoids salvaging the logs and any potential corruption of the instance.</p>

10. Before you expire objects on the target computer and the storage manager with the `onmsync` utility, perform one of the following tasks.

Otherwise, `onmsync` expires the incorrect objects.

Choose from:

- Manually edit the emergency boot file `viz ixbar.servernum` in the `$INFORMIXDIR/etc` directory on the target computer. Replace the HCL® Informix® server name that is used on the source computer with the HCL® Informix® server name of the target computer
- Run the `onmsync -b` command as user **informix** on the target computer to regenerate the emergency boot file from the `sysutils` database only. The regenerated emergency boot file reflects the server name of the target computer.

Related reference[Avoid salvaging logical logs on page 75](#)

Example: performing an imported restore with ON-Bar and the Informix® Storage Manager

This example shows how to set up an imported restore of an instance using ON-Bar, the Informix® Storage Manager (ISM), and file devices for the archives.

There is more than one way to perform an imported restore. This example shows the ISM catalog copy method. Another method, the bootstrap recovery method, is described in the HCL Informix® Storage Manager Administrator's Guide.

**Prerequisites:**

- A source machine and a target machine with the same configuration. However, the server name and number can be different.
- A ROOTPATH that is the same on both machines.
- The target machine has the devices and links in place for the chunks and these devices and links match those on the source machine.
- ISM is initialized on both computers.
- The paths to the device directories, volume names, and pool names are the same on both machines.
- User **root** and user **informix** are ISM administrators on both machines.

In this example the directories for dbspace and log backups are:

```
<directory path>/dbspaces1
<directory path>/logfiles1
```

Additional environment parameters set in the source environment are:

```
ISM_server = source computer
export IDS_server
```

Additional environment parameters set on the target machine are:

```
ISM_client = source computer
export IDS_client
SM_server = target computer
export ISM_server
```

1. As user **informix**, perform level-0 backup on the source machine.
2. As user **root**, stop ISM on both computers by running this command: `%ism_shutdown`

3. As user **root**, compress the appropriate ISM directories on the source machine as follows:

```
%cd /nsr
%tar -cvf nsr.tar index mm
```

4. FTP the `nsr.tar` file from the previous step to target machine in binary mode.

5. As user **root**, unpack the `nsr.tar` file on the target machine as follows:

```
%cd /nsr
%tar -xvf nsr.tar
```

6. As user **root** on the source machine, tar the backup directories (devices) as follows:

```
%tar -cvf logfiles1.tar logfiles1
%tar -cvf dbspaces1.tar dbspaces1
```

7. FTP the archived directories from the previous step to the target machine in binary mode.

8. As user **root** on the target machine, overwrite the existing log and archive directories with the directories from the source machine, as follows:

```
%tar -xvf logfiles1.tar
%tar -xvf dbspaces1.tar
```

9. As user **root** on the target machine

a. Start ISM by running this command: `ism_startup`

b. Run the `ism_show -devices` command to show the devices as mounted.

10. On the target machine, create a file (for example, **nsr.txt**) with this content:

```
create type: NSR client; name: source_machine;
remote access: root@target_machine, informix@target_machine;
```

11. As user **informix** on the target machine, run this command:

```
%nsradmin -s target_machine -i nsr.txt
```

The command returns this output: `created resource id <IP information>`

12. As user **informix**, FTP the following files from the `$INFORMIXDIR/etc` directory on the source machine to the `$INFORMIXDIR/etc` directory on the target machine

`ixbar,servernum`

`oncfg_servername,servernum`

13. On the target machine change the file names for the files in the previous step to the corresponding names for the local server.

14. Run the restore command.

onbar -RESTART syntax: Restarting a failed restore

If a failure occurs with the database server, media, storage manager, or ON-Bar during a restore, you can restart the restore from the place that it failed. To restart a failed restore, the `RESTARTABLE_RESTORE` configuration parameter must be set to `ON` in the `onconfig` file when the restore fails.

Restart a restore

`onbar-RESTART`

Table 16. onbar -RESTART command

Option	Description
-RESTART	<p>Restarts a restore after a database server, storage manager, or ON-Bar failure.</p> <p>The <code>RESTARTABLE_RESTORE</code> configuration parameter must be set to ON when the restore failure occurs.</p> <p>You can restart the following types of restores:</p> <ul style="list-style-type: none"> • Whole system • Point in time • Storage spaces • Logical part of a cold restore <p>Do not use the <code>-RESTART</code> option if a failure occurs during a warm logical restore.</p>

Usage

When you enable restartable restore, the logical restore is slower if many logical logs are restored. However, you save time if the restore fails and you restart the restore. Whether a restore is restartable does not affect the speed of the physical restore.

The physical restore restarts at the storage space and level where the failure occurred. If the restore failed while some, but not all, chunks of a storage space were restored, all chunks of that storage space are restored. If storage spaces and incremental backups are restored successfully before the failure, they are not restored again.


If the `BAR_RETRY` configuration parameter is set to 2, ON-Bar automatically tries to restore any failed storage spaces and logical logs again. If the restore is successful, you do not need to restart the restore.

If the `BAR_RETRY` configuration parameter is set to 0 or 1, ON-Bar does not try to restore any failed storage spaces and logical logs again. If the database server is still running, ON-Bar skips the failed storage space and attempts to restore the remaining storage spaces. To complete the restore, run the `onbar -RESTART` command.


The following figure shows how a restartable restore works when the restore failed during a physical restore of **dbspace2**. The level-0, level-1, and level-2 backups of **rootdbs**, and the level-0 and level-1 backups of **dbspace1** and **dbspace2** are successfully restored. The database server fails while restoring the level-1 backup of **dbspace2**. When you restart the restore, ON-Bar restores the level-2 backup of **dbspace 1**, the level-1 and level-2 backups of **dbspace2**, and the logical logs.

Figure 15. Restartable physical restore

Restore failed during a physical restore of dbspace2:



Restart the restore:



If a restore fails during the logical phase and you restart the restore, ON-Bar verifies that the storage spaces are restored, skips the physical restore, and restarts the logical restore. The following figure shows a cold restore that failed while restoring logical log LL-3. When you restart the cold logical restore, log replay starts from the last restored checkpoint. In this example, the last checkpoint is in logical log LL-2.

If a failure occurs during a cold logical restore, ON-Bar restarts the restore at the place of failure.


! **Important:** If a failure occurs during a warm logical restore, restart the restore from the beginning. If the database server is still running, run the `onbar -r -l` command to complete the restore.

Figure 16. Restartable cold logical restore

Cold restore failed during a logical restore of LL-3. The last checkpoint is in LL-2.



Restart the cold restore:



Related reference

[onbar -r syntax: Restoring data on page 66](#)

[BAR_RETRY configuration parameter on page 237](#)

[RESTARTABLE_RESTORE configuration parameter on page 248](#)

[ON-Bar security on page 40](#)

Resolve a failed restore

How you resolve a failed restore depends on the cause of the failure.

You can save some failed restores even if restartable restore is turned off. For example, if the restore fails because of a storage-manager or storage-device error, you can fix the tape drive or storage-manager problem, remount a tape, and then continue the restore.

The following table shows what results to expect when physical restore fails and the value of the BAR_RETRY configuration parameter is > 1.

Table 17. Failed physical restore scenarios

Type of error	RESTARTABLE_RESTORE setting	What to do when the physical restore fails?
Database server, ON-Bar, or storage-manager error (database server is still running)	ON or OFF	<p>ON-Bar tries each failed restore again. If the storage manager failed, fix the storage-manager error.</p> <p>If the tried restore fails, issue <code>onbar -r spaces</code> where <i>spaces</i> is the list of storage spaces not yet restored. Use <code>onstat -d</code> to obtain the list of storage spaces that need to be restored. ON-Bar restores the level-0 backup of each storage space, then the level-1 and level-2 backups, if any.</p>
ON-Bar or storage-manager error (database server is still running)	ON	<p>Issue the <code>onbar -RESTART</code> command.</p> <p>If the storage manager failed, fix the storage-manager error.</p> <p>The restore restarts at the storage space and backup level where the first restore failed. If the level-0 backup of a storage space was successfully restored, the restarted restore skips the level-0 backup and restores the level-1 and level-2 backups, if any.</p>
Database server failure	ON or OFF	<p>Because the database server is down, perform a cold restore. Use <code>onbar -r</code> to restore the critical dbspaces and any noncritical spaces that were not restored the first time.</p>
Database server failure	ON	<p>Issue the <code>onbar -RESTART</code> command.</p> <p>The restore restarts at the storage space and backup level where the first restore failed. If the level-0 backup of a storage space was</p>

Table 17. Failed physical restore scenarios (continued)

Type of error	RESTARTABLE_RESTORE setting	What to do when the physical restore fails?
		successfully restored, the restarted restore skips the level-0 backup and restores the level-1 and level-2 backups, if any.

The following table shows what results to expect when logical restore fails.

Table 18. Failed logical restore scenarios

Type of error	RESTARTABLE_RESTORE setting	What to do when a logical restore fails?
Database server or ON-Bar error in a cold restore (database server is still running)	ON	Issue the onbar -RESTART command. The logical restore restarts at the last checkpoint. If this restore fails, shut down and restart the database server to initiate fast recovery of the logical logs. All logical logs not restored are lost.
Database server or ON-Bar error (database server is still running)	ON or OFF	Issue the onbar -r -l command. The restore restarts at the failed logical log. If onbar -r -l still fails, shut down and restart the database server. The database server completes a fast recovery. All logical logs that were not restored are lost. If fast recovery does not work, you must do a cold restore.
Database server error	ON	If the cold logical restore failed, issue onbar -RESTART. If the warm logical restore failed, issue the onbar -r -l command. If that fails, restart the entire restore from the beginning.
Storage-manager error	ON or OFF	ON-Bar tries each failed logical restore again. If the tried restore fails, the logical restore is suspended. Fix the storage-manager error. Then issue the onbar -r -l command. The restore restarts at the failed logical log.

Related reference

[BAR_RETRY configuration parameter on page 237](#)

[RESTARTABLE_RESTORE configuration parameter on page 248](#)

External backup and restore

These topics discuss recovering data by using external backup and restore.

External backup and restore overview

An external backup and restore eliminates the downtime of systems because the backup and restore operations are performed external to the HCL® Informix® system.

ON-Bar does not move the data during the backup or physical restore. An external backup allows you to copy disks that contain storage-space chunks without using ON-Bar. When disks fail, replace them and use vendor software to restore the data, then use ON-Bar for the logical restore. For more information, see [Data restored in an external restore on page 97](#).

The following are typical scenarios for external backup and restore:

- Availability with disk mirroring

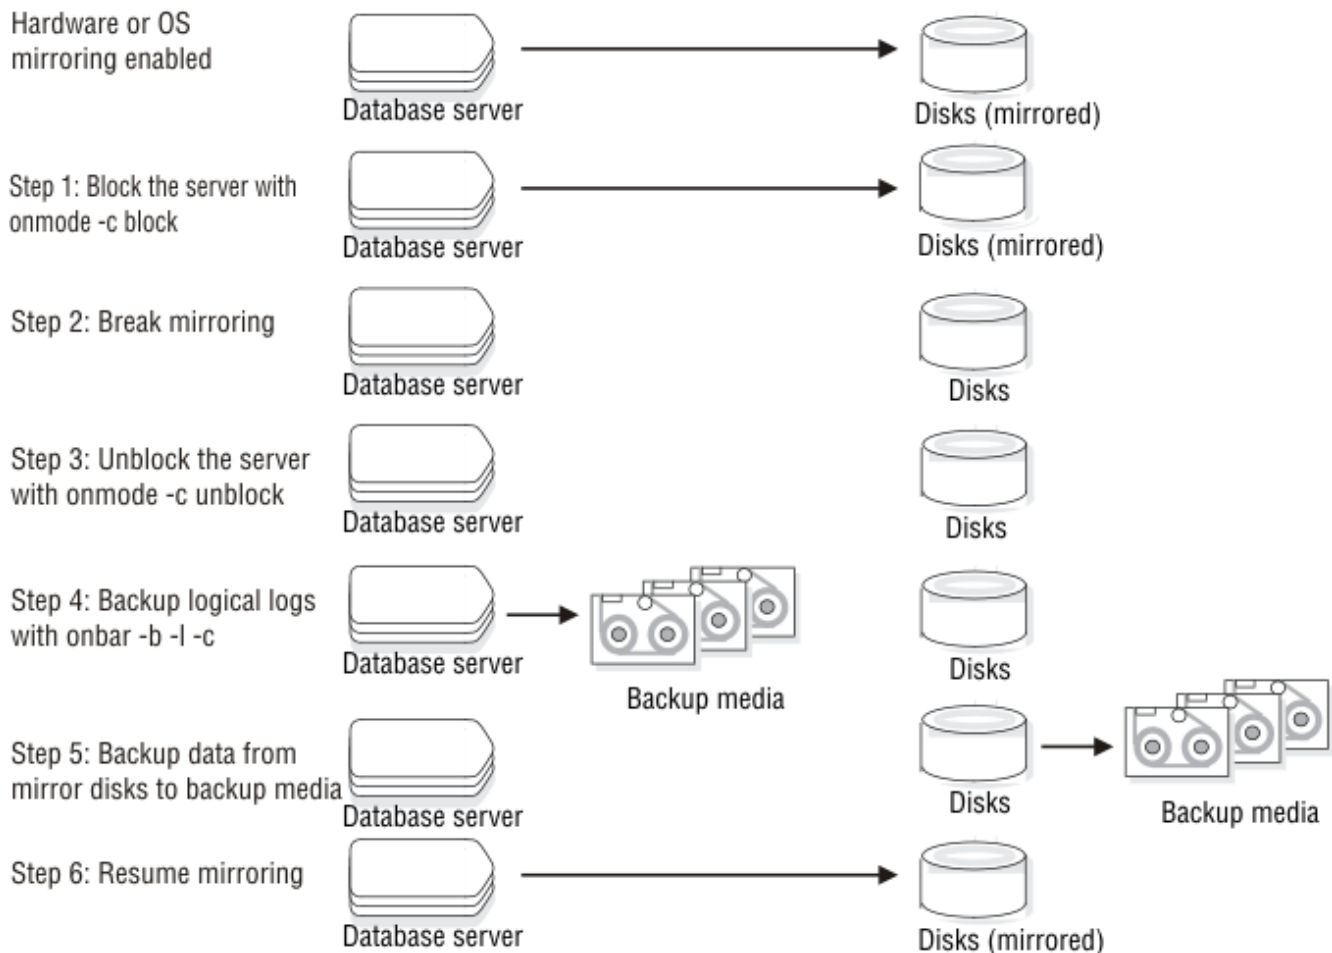
If you use hardware disk mirroring, you can get your system online faster with external backup and restore than with conventional ON-Bar commands.

- Cloning

You can use external backup and restore to clone an existing production system for testing or migration without disturbing the production system.

The following figure shows how to perform a backup with mirroring.

Figure 17. Perform a backup with mirroring



In this configuration, the database server is running continuously, except for the short time when the database server is blocked to break the mirror. The mirrored disks contain a copy of the database server storage spaces. To create a backup, block the database server to stop transactions and disable mirroring. The mirrored disks now contain a copy of the consistent data at a specific point in time. After disabling mirroring, unblock the database server to allow transactions to resume and then backup the logical logs. Copy the data from the offline mirrored disks to back up media with external commands. Now you can resume mirroring.

Block before backing up

Before you begin an external backup, block the database server. Blocking forces a checkpoint, flushes buffers to disk, and blocks user transactions that involve temporary tables.

During the blocking operation, users can access that database server in read-only mode. Then you can physically back up or copy the data to another set of disks or storage media by using operating-system or third-party tools. When you complete the external backup, unblock the database server so that transactions can resume. You should include all the chunk files in each storage space, administrative files, such as `onconfig`, and the emergency boot file, in an external backup.

! **Important:** To make tracking backups easier, you should back up all storage spaces in each external backup.

ON-Bar treats an external backup as equivalent to a level-0 backup. You cannot perform an external backup and then use ON-Bar to perform a level-1 backup, or vice versa because ON-Bar does not have any record of the external backup. For more information, see [Performing an external backup when chunks are not mirrored on page 95](#).

Rules for an external backup

Before you begin an external backup, review the rules for performing an external backup.

The rules that you must follow are:

- The database server must be online or in quiescent mode during an external backup.
- Use ON-Bar to back up all logical logs including the current log so that you can restore the logical logs at the end of the external restore.
- Suspend continuous logical-log backups before you block the database server for an external backup. After the external backup is complete, resume the continuous logical-log backup.

To stop continuous logical-log backup, use the CTRL-C command. To resume continuous logical-log backup, use the `onbar -b -l -C` command.

- Wait until all ON-Bar backup sessions have completed before you block the database server. If any backup sessions are active, the block command displays an error message.
- Any OLTP work or queries are suspended while the database server is blocked. They resume after the database server is unblocked.
- All critical dbspaces of the database server instance must be backed up together simultaneously within the same command bracket of `onmode -c block onmode -c unblock`. Backups of different critical dbspaces done at different times cannot be restored to a consistent system.
- On AIX® operating systems, if the server is running with concurrent I/O because the `DIRECT_IO` configuration parameter is set to enable concurrent I/O, an online external backup program must also use concurrent I/O.

! **Important:** Because the external backup is outside the control of ON-Bar, you must track these backups manually. For more information, see [Track an external backup on page 95](#).

Prepare for an external backup

These topics describe the commands used to prepare for an external backup. For the procedure, see [Performing an external backup when chunks are not mirrored on page 95](#).

Block and unblock database server

This topic shows the syntax of the block and unblock commands on HCL Informix®.

```
onmode -c [{block | unblock}]
```

Element	Purpose	Key considerations
onmode -c	Performs a checkpoint and blocks or unblocks the database server	None.
block	Blocks the database server from any transactions	Sets up the database server for an external backup. While the database server is blocked, users can access it in read-only mode. Sample command: onmode -c block
unblock	Unblocks the database server, allowing data transactions and normal database server operations to resume	Do not unblock until the external backup is finished. Sample command: onmode -c unblock

Track an external backup

The database server and ON-Bar do not track external backups. To track the external backup data, use a third-party storage manager or track the data manually.

The following table shows which items we recommend that you track in an external backup. ON-Bar keeps a limited history of external restores.

Table 19. Items to track when you use external backup and restore

Items to track	Examples
Full path names of each chunk file for each backed up storage space	/work/dbspaces/rootdbs (UNIX™) c:\work\dbspaces\rootdbs (Windows™)
Object type	Critical dbspaces, noncritical storage spaces
ins_copyid_hi and ins_copyid_lo	Copy ID that the storage manager assigns to each backup object
Backup date and time	The times that the database server was blocked and unblocked
Backup media	Tape volume number or disk path name
Database server version	The database server version from which the backup was taken.

Performing an external backup when chunks are not mirrored

About this task

The database server must be online or in quiescent mode during an external backup.

To perform an external backup when chunks are not mirrored:

1. To obtain an external backup, block the database server with the onmode -c block command.
The system takes a checkpoint and suspends all update transactions. Users can access the database server in read-only mode.

2. To back up the storage spaces and administrative files, use a copy command, such as cp, dd, or tar on UNIX™ or copy on Windows™, or a file-backup program.

You must back up all chunks in the storage spaces.

3. To allow normal operations to resume, unblock the database server with the onmode -c unblock command.
4. Back up all the logical logs including the current log so that checkpoint information is available for the external restore.

! **Important:** Because external backup is not done through ON-Bar, you must ensure that you have a backup of the current logical log from the time when you execute the onmode -c block command. Without a backup of this logical-log file, the external backup is not restorable.

5. After you perform an external backup, back up the current log with the onbar -b -l -c command.

Results

If you lose a disk, or the whole system, you are now ready to perform an external restore.

RS secondary server external backup

You can perform an external backup of an RS secondary server. Performing a backup of an RS secondary server blocks that RS secondary server, but does not block the primary server.

You can perform a logical restore from the logs backed up from the primary instance. The backup obtained from the secondary server cannot be restored with level-1 or level-2 backups.

! **Important:** The external backup is not completed if the database instance contains any of the following:

- Nonlogging smart large objects
- Regular blobspaces
- Nonlogging databases
- Raw tables

If an external backup is performed on an instance that contains any of the previously mentioned items, then the backup is incomplete and cannot be used to restore the primary server.

If the backup fails because the checkpoint from the primary has timed out, you can use the BAR_CKPTSEC_TIMEOUT configuration parameter to increase the amount of time, in seconds, that an RS secondary server should wait for a checkpoint to arrive from the primary server while performing an external backup.

Performing an external backup of an RS secondary server

Before you begin

To perform an external backup of an RS secondary server, the STOP_APPLY configuration parameter must not be enabled. If STOP_APPLY is enabled, an error is returned. The server switches to STOP_APPLY mode when a backup is performed on an

RS secondary. After the archive checkpoint is processed, the RS secondary server stops applying logical logs, but continues receiving logs from the primary server.

To perform an external backup of an RS secondary server that has a DELAY_APPLY configuration parameter value greater than 0, it might be necessary to temporarily decrease the parameter value. Performing the backup requires that the RSS process a checkpoint in the logical log, and if no checkpoint is observed within the amount of time that is specified by the onmode -c block *timeout* command in the second step of the following procedure, a backup is not permitted. The DELAY_APPLY configuration parameter can be decreased by the onmode -wf DELAY_APPLY=*setting* command.

About this task

The primary database server must be online or in quiescent mode during an external backup.

To perform an external backup:

1. Ensure that the LOG_STAGING_DIR configuration parameter on the RS secondary server is set to point to a valid staging directory.
2. To obtain an external backup, block the database server with the onmode -c block *timeout* command. The *timeout* parameter indicates the number of seconds that the RS secondary server waits to receive a checkpoint. The *timeout* parameter is valid only when the onmode -c block command is run on an RS secondary server. You must wait for the onmode -c block command to return successfully before you proceed with the external backup.
3. To back up the storage spaces and administrative files, use a copy command, such as cp, dd, or tar on UNIX™ or copy on Windows™, or a file-backup program. You must back up all chunks in the storage spaces.
4. To resume normal operations, unblock the database server by using the onmode -c unblock command.
5. After you perform the external backup, back up the current log and any new logs with the ON-Bar or ontape utilities.



Important: Logical log backup is only possible on the primary server.

If the DELAY_APPLY configuration parameter is set, the logs that are required for the restore process are not necessarily those logs that are currently active on the primary server because some logs could already be archived.

Results

After the backup completes, if the DELAY_APPLY setting on the RS secondary server was decreased, it can be set to its original value by the onmode -wf DELAY_APPLY=*setting* command. After an external backup, you can perform an external restore if a disk or the whole system fails.

Data restored in an external restore

If you lose a disk, or the whole system, you can externally restore data only if it was externally backed up. You must use the same third-party utility for both the external backup and restore. To externally restore the storage spaces, copy the backed-up data to disk. Use the onbar -r -e command to mark the storage spaces as physically restored, replay the logical logs, and

bring the storage spaces back online. If you do not specify an external restore command, the database server thinks that these storage spaces are still down.

You can perform these types of external restores:

- Warm external restore

Mark noncritical storage spaces as physically restored, then perform a logical restore of these storage spaces.

- Cold external restore

Mark storage spaces as physically restored, then perform a logical restore of all storage spaces. Optionally, you can do a point-in-time cold external restore.



Restriction: When you perform a cold external restore, ON-Bar does not first attempt to salvage logical-log files from the database server because the external backup has already copied over the logical-log data.

To salvage logical logs, perform `onbar -l -s` before you copy the external backup and perform the external restore (`onbar -r -e`).

Rename chunks

You can rename chunks in an external cold restore by using the rename options syntax for other restores. Use the following commands to specify new chunk names during restore:

```
onbar -r -e -rename -f filename
```

or

```
onbar -r -e rename -p old_path -o old_offset -n new_path -o new_offset
```

External restore commands

Use the `onbar -r -e` command to perform a warm or cold external restore. This command marks the storage spaces as physically restored and restores the logical logs. The following diagram shows the external restore syntax.

Performing an external restore with ON-Bar

```
onbar -r -e [ <Rename chunks> (explicit id) ] [ { -p | -ttime | -n last_log } ] [ { -o } ] [ { -filename | dbspace_list | -w [ { -ttime | -nlog } ] } ]
```

Element	Purpose	Key considerations
<code>onbar -r</code>	Specifies a restore	In a cold restore, if you do not specify storage space names, all of them are marked as restored.

Element	Purpose	Key considerations
-e	Specifies an external restore	Must be used with the -r option. In a warm external restore, marks the down storage spaces as restored unless the -O option is specified.
<i>dbspace_list</i>	Names one or more storage spaces to be marked as restored in a warm restore	If you do not enter <i>dbspace_list</i> or -f filename and the database server is online or quiescent, ON-Bar marks only the down storage spaces as restored. If you enter more than one storage-space name, use a space to separate the names.
-filename	Restores the storage spaces that are listed in the text file whose path name <i>filename</i> provides	To avoid entering a long list of storage spaces every time, use this option. The <i>filename</i> can be any valid UNIX™ or Windows™ file name.
-n last_log	Indicates the number of the last log to restore	If any logical logs exist after this one, ON-Bar does not restore them and data is lost. The -n option does not work with the -p option.
-O	Restores online storage spaces	None.
-p	Specifies an external physical restore only	After the physical restore completes, you must perform a logical restore.
-t time	Restores the last backup before the specified point in time. If you pick a backup made after the point in time, the restore will fail.	You can use a point-in-time restore in a cold restore only. You must restore all storage spaces. How you enter the time depends on your current GLS locale convention. If the GLS locale is not set, use English-style date format.
-w	Performs a whole-system restore of all storage spaces and logical logs from the last whole-system backup	You must specify the -w option in a cold restore. If you specify <code>onbar -r -w</code> without a whole-system backup, return code 147 appears because ON-Bar cannot find any storage spaces backed up as part of a whole-system backup.

Related reference

[onbar -r syntax: Restoring data on page 66](#)

Rules for an external restore

External restores have specific rules.

The following rules apply to external restores:

- You must externally restore from an external backup. Although the external backup is treated as a level-0 backup, it might actually be an incremental backup not related to HCL Informix®.
- A warm external restore restores only noncritical storage spaces.
- You cannot externally restore temporary dbspaces.
- You cannot externally restore from regular ON-Bar backups.
- You cannot verify that you are restoring from the correct backup and that the storage media is readable with ON-Bar.
- If the external backups are from different times, the external restore uses the beginning logical log from the oldest backup.
- You cannot perform a mixed restore. If critical dbspaces must be restored, you must perform a full cold restore.

The following rules apply to external cold restores:

- Salvage the logical logs (`onbar -b -l -s`) before you switch the disks that contain the critical storage spaces.
- If you are restoring critical dbspaces, the database server must be offline.
- Point-in-time external restores must be cold and restore all storage spaces.
- The external backups of all critical dbspaces of the database server instance must be simultaneous. All critical dbspaces must have to be backed up within the same set of `onmode -c block ... onmode -c unblock` commands.

Performing an external restore

About this task

This section describes procedures for performing cold and warm external restores.

Performing a cold external restore

About this task

If you specify the `onbar -r -e` command in a cold restore, you must restore all storage spaces. Use the `onbar -r -e -p` command to restore all or specific storage spaces.

To perform a cold external restore:

1. Shut down the database server with the `onmode -ky` command.
2. Salvage the logical logs with the `onbar -b -l -s` command.
3. To restore the storage spaces from an external backup, use a copy command, such as `cp`, `dd`, or `tar` on UNIX™ or a file-backup program.

You must restore the storage spaces to the same path as the original data and include all the chunk files.

4. To perform an external restore of all storage spaces and logical logs, use the `onbar -r -e` command.
5. To perform a point-in-time external restore of all storage spaces, use the `onbar -r -e -t datetime` command.

This step brings the database server to fast-recovery mode.

ON-Bar and the database server roll forward the logical logs and bring the storage spaces online.

Performing a warm external restore

About this task

The database server is online during a warm external restore. A warm external restore involves only noncritical storage spaces.

To perform a warm external restore:

1. To restore the storage spaces from an external backup, use a copy command, such as `cp`, `dd`, or `tar` on UNIX™ or a file-backup program.

You must restore the storage spaces to the same path as the original data and include all the chunk files for each restored storage space.

2. Perform a warm external restore of the noncritical storage spaces to bring them online.

Choose from:

- To restore selected storage spaces and all logical logs, use the `onbar -r -e dbspace_list` command.
- To restore the down noncritical storage space named `dbsp1` and logical logs in separate steps, use the following commands:


```
onbar -r -e -p dbsp1
onbar -r -l dbsp1
```
- To restore all the noncritical storage spaces and logical logs, use the `onbar -r -e -O` command.

Examples of external restore commands

The following table contains examples of external restore commands.

External restore command	Action	Comments
<code>onbar -r -e</code>	Complete external restore	In a cold restore, restores everything. In a warm restore, restores all down noncritical storage spaces.
<code>onbar -r -e -p</code> <code>onbar -r -l</code>	Physical external restore and separate logical restore	If the external backups come from different times, you must perform a logical restore. The system restores the logical logs from the oldest external backup.
<code>onbar -r -e <i>dbspace_list</i></code>	External restore of selected storage spaces and logical logs	Use this command in a warm external restore only.
<code>onbar -r -e -p <i>dbspace_list</i></code> <code>onbar -r -l</code>	External restore of selected storage spaces and separate logical restore	Use this command in a warm external restore only.

External restore command	Action	Comments
<code>onbar -r -e -t <i>datetime</i></code>	External point-in-time (cold) restore	Be sure to select a collection of backups from before the specified time.
<code>onbar -r -e rename -p <i>old_path</i> -o <i>old_offset</i> -n <i>new_path</i> -o <i>new_offset</i></code>	External (cold) restore with renamed chunks	Use this command to rename chunks in cold external restore only.
<code>onbar -r -e -w</code> <code>onbar -r -e -p -w</code>	Whole-system external restore	When you use <code>onbar -r -e -w -p</code> , back up all storage spaces in one block and unblock session. That way, all storage spaces have the same checkpoint.

Initializing HDR with an external backup and restore

About this task

You can use external backups to initialize High-Availability Data Replication (HDR).

To initialize HDR with an external backup and restore:

1. Block the source database server with the `onmode -c block` command.
2. Externally back up all chunks on the source database server.
3. When the backup completes, unblock the source database server with the `onmode -c unblock` command.
4. Make the source database server the primary server with the following command: `onmode -d primary secondary_servername`
5. On the target database server, restore the data from the external backup with a copy or file-backup program.
6. On the target database server, restore the external backup of all chunks with the `onbar -r -e -p` command.
On HDR, secondary server can restore only level-0 archives.
7. Make the target database server the secondary server with the following command: `onmode -d secondary primary_servername`
8. If the logical-log records written to the primary database server since step 1 still reside on the primary database server disk, the secondary database server reads these records to perform the logical recovery. Otherwise, perform the logical recovery with the `onbar -r -l` command.

The `database server operational` messages appear in the message log on the primary and secondary servers.

Customize and maintain ON-Bar

These topics discuss the following:

- Customizing ON-Bar and storage-manager commands with the `onbar` script
- Starting `onbar-worker` processes manually
- Expiring and synchronizing the backup history

Customizing ON-Bar and storage-manager commands

You can edit the script that is installed with ON-Bar to customize backup and restore commands, and storage manager commands.

On UNIX™ operating systems, the onbar shell script is in the `$INFORMIXDIR/bin` directory. On Windows™ operating systems, the onbar.bat batch script is in the `%INFORMIXDIR%\bin` directory.

Edit the script and backup a copy of the original file in case you need to revert to it.



Important: Edit the script with caution and test your changes. Do not change the cleanup code at the bottom of the script. Doing so might result in unexpected behavior, for example, leftover temporary files during backup verification.

The script includes commands for IBM® Informix® Storage Manager (ISM) and backs up the ISM catalogs. If you are using a different storage manager, delete the ISM-specific lines, and optionally, add commands for the storage manager that you are using.

The script contains the following sections:

- Add startup processing here

Use this section to initialize the storage manager, if necessary, and set environment variables.

- End startup processing here

This section starts the onbar_d driver and checks the return code. Use this section for onbar_d and storage-manager commands.

- Add cleanup processing here

The code in this section backs up the ISM catalogs to the ISMData volume pool after the backup operation is complete. If you are using a third-party storage manager, delete the ISM-specific information.

If you use a name other than ISMData for the volume pool, change it to the name specified in the ISM_DATA_POOL configuration parameter.

This section removes the archecker temporary files.

- End cleanup processing here

Use this section to return onbar_d error codes.

Related reference

[Back up with ON-Bar on page 43](#)

Updating the ON-Bar script during reinstallation

After you reinstall the database server, you might need to update the script that is installed with ON-Bar. Existing customized scripts were backed up by the installation process so that you can reuse the content.

The installation program installs the default onbar shell script on UNIX™ and the default onbar.bat batch script on Windows™. If the default script differs from the local script, the installation program backs up the local script and issues a message to inform you that the local script was renamed. The naming convention of the renamed file is `onbar.date`, where `date` is the date when the file was renamed. For example, the file `onbar.2012.05.15` was renamed on May 15, 2012.

You can update the default script by adding information to it from the renamed script.

Print the backup boot files

Use the following examples of what to add to the onbar script to print the emergency boot file if the backup is successful. Each time that you issue the `onbar -b` command, the emergency boot file is printed.

The following example is for UNIX™:

```
onbar_d "$@" # receives onbar arguments from command line return_code = $?
# check return code

# if backup (onbar -b) is successful, prints emergency boot file
if [$return_code -eq 0 -a "$1" = "-b"]; then
    servernum='awk '/^SERVERNUM/ {print $2}' $INFORMIXDIR/etc/$ONCONFIG '
    lpr \${INFORMIXDIR}/etc/ixbar.$servernum
fi
exit $return_code
```

The following example is for Windows™:

```
@echo off
%INFORMIXDIR%\bin\onbar_d %*
set onbar_d_return=%errorlevel%

if "%onbar_d_return%" == "0" goto backupcom
goto skip

REM Check if this is a backup command

:backupcom
if "%1" == "-b" goto printboot
goto skip

REM Print the onbar boot file

:printboot
print %INFORMIXDIR%\etc\ixbar.???

REM Set the return code from onbar_d (this must be on the last line of the script)

:skip
%INFORMIXDIR%\bin\set_error %onbar_d_return%
:end
```


Migrate backed-up logical logs to tape

You can set up your storage manager to back up logical logs to disk and then write a script to automatically migrate the logical logs from disk to tape for off-site storage. Edit the onbar script to call this migration script after the onbar_d process completes. The following example shows a script that calls the migration script:

The following example is for UNIX™:

```
onbar_d "$@" # starts the backup or restore
EXIT_CODE=$? # any errors?

PHYS_ONLY=false #if it's physical-only, do nothing
for OPTION in $*; do
    if [ $OPTION = -p ]; then
        PHYS_ONLY = true
    fi
done
if ! PHYS_ONLY; then # if logs were backed up, call another
    migrate_logs # program to move them to tape
fi
```

This example for Windows™ invokes the migration script:

```
%INFORMIXDIR%\bin\onbar_d %*
set onbar_d_return=%errorlevel%

if "%onbar_d_return%" == "0" goto backupcom
goto skip

REM Check if the command is a backup command

:backupcom
if "%1" == "-b" goto m_log
if "%1" == "-l" goto m_log
goto skip

REM Invoke the user-defined program to migrate the logs

:m_log
migrate_log

REM Set the return code from onbar_d (this must be on the last line of the script)

:skip
%INFORMIXDIR%\bin\set_error %onbar_d_return%

:end
```

Expire and synchronize the backup catalogs

ON-Bar maintains a history of backup and restore operations in the **sysutils** database and an extra copy of the backup history in the emergency boot file. ON-Bar uses the **sysutils** database in a warm restore when only a portion of the data is lost.

ON-Bar uses the emergency boot file in a cold restore because the **sysutils** database cannot be accessed. You can use the **onsmsync** utility to regenerate the emergency boot file and expire old backups.

Depending on the command options you supply, the **onsmsync** utility can remove the following items from the **sysutils** database and the emergency boot file:

- Backups that the storage manager has expired
- Old backups based on the age of backup
- Old backups based on the number of times they have been backed up

Use **onsmsync** with the database server online or in quiescent mode to synchronize both the **sysutils** database and the emergency boot file.

To synchronize the **sysutils** database:

1. Bring the database server online or to quiescent mode.
2. Run the **onsmsync** utility without any options.

The **onsmsync** utility synchronizes the **sysutils** database, the storage manager, and the emergency boot file as follows:

- Adds backup history to **sysutils** that is in the emergency boot file but is missing from the **sysutils** database.
- Removes the records of restores, whole-system restores, fake backups, successful and failed backups from the **sysutils** database.
- Expires old logical logs that are no longer needed.
- Regenerates the emergency boot file from the **sysutils** database.

Choose an expiration policy

You can choose from the following three expiration policies:

Retention date (-t)

Deletes all backups before a particular date and time.

Retention interval (-i)

Deletes all backups older than a specified period.

Retention generation (-g)

Keeps a certain number of versions of each backup.

ON-Bar always retains the latest level-0 backup for each storage space. It expires all level-0 backups older than the specified time unless they are required to restore from the oldest retained level-1 backup.

ON-Bar expires all level-1 backups older than the specified time unless they are required to restore from the oldest retained level-2 backup.

ON-Bar retains a whole-system backup that starts before the specified retention time and ends after the specified retention time.

The onsmsync utility

Use the onsmsync utility to synchronize the **sysutils** database and emergency boot file with the storage manager catalog.

If your storage manager is the Informix® Primary Storage Manager, you can also use the export and import options of the onsmsync utility to export Informix® Primary Storage Manager backup objects to an external device and import objects from an external device to a device that is managed by the Informix® Primary Storage Manager. You cannot use the export and import options with other storage managers.

```
onmsync [{ -ggeneration | -ttimestamp | -iinterval }] [-s] [-o] [{ -ffilename | dbspace }][ -cvalue ]| [-b] [{ <Export  
and import options > }]
```

Export and import options (to use only with the Primary Storage Manager)

```
" { -E -pprefix [{ -ggeneration | -ttimestamp | -llog_ID }]| -I -pprefix } "
```

The following table lists all onsmsync command elements, except the elements you use to import and export backup generations. The command elements that you use for importing and exporting are listed in [Table 21: Elements for onsmsync export and import commands on page 108](#).

Table 20. Elements for onsmsync commands

Element	Purpose	Key considerations
-b	Regenerates the emergency boot file (<i>ixbar.servernum</i>) and the sysutils database from each other.	<p>If the <i>ixbar</i> file is empty or does not exist, onsmsync -b re-creates the <i>ixbar</i> file and populates it from the sysutils tables.</p> <p>If the <i>ixbar</i> file is not empty and contains object data, onsmsync -b updates the sysutils database and the <i>ixbar</i> file so that they are in sync.</p> <p>If the <i>ixbar</i> file has entries and the sysutils database was rebuilt, but is empty because it does not contain data, onsmsync -b recreates the sysutils data from the <i>ixbar</i> file.</p> <p>Do not use the -b element with the other onsmsync options.</p> <p>The -b element does not synchronize with the storage manager.</p>
<i>dbspace</i>	Specifies the storage space or storage spaces to expire	If you enter the name of more than one storage space, use a blank space to separate the names.
-f <i>filename</i>	Specifies the path name of a file that contains a list of storage spaces to expire	Use this option to avoid entering a long list of storage spaces. The file name can be any valid UNIX™ or Windows™ file name.

Table 20. Elements for onsmsync commands (continued)

Element	Purpose	Key considerations
<code>-g generation</code>	Specifies the number of versions of each level-0 backup to retain	The latest generation of backups is retained and all earlier ones are expired.
<code>-i interval</code>	Specifies the time interval for retaining backups.	<p>The utility:</p> <ul style="list-style-type: none"> Retains backups that were created after this interval. Expires backups that were created before this interval and removes the backups if expired objects are also removed. <p>Backups older than interval are not expired if they are needed to restore from other backups after that interval. Use the ANSI or GLS format for the <i>interval</i>: YYYY-MM or DD HH:MM:SS</p>
<code>-O</code>	Overrides internal error checks and enforces expiration policy	If used with the <code>-t</code> , <code>-g</code> , or <code>-i</code> option, expires all levels of a backup, even if some of them are needed to restore from a backup that occurred after the expiration date. The <code>-O</code> option does not affect logical-log expiration. See Expire all backups on page 113 .
<code>-s</code>	Skips synchronizing expired backups	The object expiration is based on other arguments if the <code>-s</code> option is provided.
<code>-t timestamp</code>	Expires all backups before a particular date and time	<p>Retains backups that completed after the specified timestamp. Backups that occurred before the timestamp are not expired if they are needed to restore from other backups that occurred after the timestamp.</p> <p>Use the ANSI or GLS_DATETIME format for the timestamp.</p>
<code>-cf</code>	<p>Specifies whether to expire the critical files.</p> <p>When used with the <code>-g</code>, <code>-i</code>, or <code>-t</code>, deletes critical file backups from the Informix® Primary Storage Manager</p>	<p>The critical files are the <code>onconfig</code> file, the <code>sqlhosts</code> file (on UNIX™), the <code>oncfg_servername.servernum</code> file, and the <code>ixbar.servernum</code> file. Valid values are:</p> <ul style="list-style-type: none"> yes = Deletes backups of the critical files. no = Does not delete backups of the critical files. only = Deletes only the backups of critical files.

Table 21. Elements for onsmsync export and import commands

Element	Purpose	Key considerations
<code>-E</code>	Exports a single backup generation to the Informix®	Use this option only if you set up a Informix® Primary Storage Manager external pool.

Table 21. Elements for onsmsync export and import commands (continued)

Element	Purpose	Key considerations
	Primary Storage Manager external pool	If you export a backup generation, you must specify a prefix, which identifies the exported backup. The onsmsync utility creates a subdirectory that includes the prefix in the external pool, and deposits the exported objects in that directory.
<code>-g generation</code>	Specifies the backup generation to export.	The default value is the current backup.
<code>-l</code>	Imports a single backup generation from the external Informix® Primary Storage Manager pool.	If you import a backup generation from an external pool, you must specify the prefix for exported backup. The prefix identifies the backup generation that you want to import.
<code>-l log_ID</code>	Exports the backup generation that includes a logical log ID.	
<code>-p prefix</code>	Specifies the prefix that you are assigning to a backup generation that you are exporting or are using to identify the backup that you want to import.	When the onsmsync utility exports a backup generation, it uses the prefix as the name of the subdirectory in which it places the backup.
<code>-t timestamp</code>	Specifies a backup generation that includes a particular date and time (used only for exporting).	Use the ANSI or GLS_DATETIME format. You might want to include the date and time when you roll out a new version of an application.

Usage

If no options are specified, the onsmsync command synchronizes the **sysutils** database and emergency boot file with the storage-manager catalog. The onsmsync utility compares the backups in the **sysutils** database and emergency boot file with the backups in the storage-manager catalog and then removes all backups that are not in the storage manager catalog from the **sysutils** database and emergency boot file. The onsmsync utility is also capable to expire backups from the storage manager if they are registered in the emergency boot file/sysutils database and they meet the expiration criteria set by the operator.



Tip: To control whether the **sysutils** database maintains a history for expired backups and restores, use the `BAR_HISTORY` configuration parameter. For information, see [BAR_HISTORY configuration parameter on page 231](#).

The order of the commands does not matter except that the storage-space names or file name must come last.

The order of the commands does not matter except:

- The storage-space names or file name must come last.
- When exporting or importing, the -E or -I option must be first. For example, specify `onsmsync -E -g 2`, not `onsmsync -g 2 -E`.

Prerequisites for importing and exporting backup generations on different computers:

- You must have the same version of Informix® on the source and target computers and the computers must use the same operating system.
- You must set up the Informix® Primary Storage Manager and create an external pool on the source and target computers.

When you use the -E or -I options to export or import a backup generation, you must specify the prefix that identifies the subdirectory where the backup generation is placed.

If you use the -E or -I options to export or import a backup generation, you cannot use any `onsmsync` command options that are not related to the export or import operation. For example, you cannot export a backup generation and regenerate the emergency boot file at the same time.

The `onsmsync -I` command renames your current `ixbar` file and creates a new file that contains only the information that is necessary to restore the imported backup.

You can use the `-cf` option with the `-g`, `-i`, or `-t` option to delete critical file backups from the storage manager.

If you apply the `-g` option and the `onsmsync` utility list of objects contains only logical logs and no space backups, the log backups are not expired. In this situation, use the `-t` or `-i` option to expire logical log backups.

Examples

The following example expires backups that started before November 30, 2012:

```
onsmsync -t "2012-11-30 00:00:00"
```

The following command exports the last backup generation to the directory called `gen` in the external pool:

```
onsmsync -E -p gen -g 1
```

The following command exports the fourth most recent backup generation to the directory called `gen` in the external pool:

```
onsmsync -E -p gen -g 4
```

The following command exports the current backup generation to the directory called `gen` in the external pool:

```
onsmsync -E -p gen
```

The following command exports all backup objects in generation 2 to the directory called `gen` in the external pool:

```
onsmsync -E -p gen -g 2
```

The following command exports all backup objects that have a timestamp of `2012-12-31 12:00:00` to the directory called `gen` in the external pool:

```
onsmsync -E -p gen -t "2012-12-31 12:00:00"
```

The following command imports all objects in the subdirectory that is identified by the prefix `gen`:

```
onsmsync -I -p gen
```

The following command imports all backup objects that were exporting using the prefix `gen` and the timestamp of `2012-12-31 12:00:00`. Because the prefix identifies the backup generation, you do not specify a timestamp.

```
onsmsync -I -p gen
```

The following command deletes all but the last two generations of critical file backups:

```
onsmsync -g 2 -cf yes
```

Related reference

[ON-Bar script on page 24](#)

Related information

[Managing storage devices on page 191](#)

[Examples: Manage storage devices with Informix Primary Storage Manager on page 183](#)

[Informix Primary Storage Manager file-naming conventions on page 202](#)

[Restoring when a backup is missing data on page 63](#)

[Performing a cold restore on page 75](#)

Expiring old backups on ISM

IBM® Informix® Storage Manager (ISM) and certain third-party storage managers do not allow the `onsmsync` utility to delete backups from the storage manager.

About this task

You must manually expire or delete the old backups from the storage manager. Then, run `onsmsync` without any parameters.

To expire old backups on ISM:

1. To manually expire the old backups from ISM, use the `ism_config -retention #days` command.

For more information, see the *Informix® Storage Manager Administrator's Guide*.

2. Run `onsmsync` without any options.

Regenerate the emergency boot file

To regenerate the emergency boot file only, use the `onsmsync -b` command.

The `onsmsync -b` command saves the old emergency boot file as `ixbar.server_number.system_time` and regenerates it as `ixbar.server_number`.

Regenerate the sysutils database

If you lose the **sysutils** database, use the `bloutil` utility in `$INFORMIXDIR/etc` on UNIX™ or `%INFORMIXDIR%\etc` on Windows™ to recreate the **sysutils** database with empty tables.

Then use the `onsmsync` utility to recreate the backup and restore data in **sysutils**.



Restriction: If both the **sysutils** database and emergency boot file are missing, you cannot regenerate them with `onsmsync`. Be sure to back up the emergency boot file with your other operating-system files.

Delete a bad backup

The `onsmsync` utility cannot tell which backups failed verification. If the latest backup failed verification but an earlier one was successful, you must manually delete the failed backup records from the storage manager and then run `onsmsync` with no options to synchronize ON-Bar. For more information, see [onbar -v syntax: Verifying backups on page 58](#).

Expire backups based on the retention date

The following example expires backups that started before November 24, 2006, and all fake backups, failed backups, and restores:

```
onsmsync -t "2006-11-24 00:00:00"
```

Expire a generation of backups

The following example retains the latest three sets of level-0 backups and the associated incremental backups, and expires all earlier backups and all restores, fake backups, and failed backups: `onsmsync -g 3`

Expire backups based on the retention interval

The following example expires all backups that are older than three days and all fake backups, failed backups, and restores:

```
onsmsync -i "3 00:00:00"
```

The following example expires all backups older than 18 months (written as 1 year + 6 months):


```
onsmsync -i "1-6"
```

Expire backups with multiple point-in-time restores

If you perform more than one point-in-time restores, multiple timelines for backups exist.

The following figure shows three timelines with their backups.

Figure 19. Multiple timelines for backups



In this example, the second timeline begins with a point-in-time restore to backup 1. The second timeline consists of backups 1, 5, 6, 7, and 8. The third timeline (in bold) consists of backups 1, 5, and 9. The third timeline is considered the current timeline because it contains the latest backup.

When you run the `onsmsync` utility to expire old backups, `onsmsync` removes the old backups from the current timeline, and make sure that the current timeline is restorable from the backup objects that are retained. All other backups that are not in the current timeline are also expired but `onsmsync` does not make sure that the other timelines are restorable from the objects retained.

The `onsmsync` utility applies expiration policies in the following order to make sure that objects from current timeline are expired according to the specified expiration policy and that the current timeline is restorable:

- Apply the expiration policy on all sets of backup objects.
- Unexpire backup objects that belong to the current timeline.
- Apply the expiration policy on the current timeline to ensure that the current timeline is restorable.

At the same time, the expiration policy is applied to backups in other timelines.

For example, if you execute the `onsmsync -g 2` command on the example in the previous figure, backup 1 from the current timeline is expired, and backups 2, 3, 4, 6, and 7 from the first and second timelines are expired. Backups 1, 5, and 9 from the current timeline are retained. Backup 8 is retained from other timelines.

Expire all backups

The `onsmsync` utility retains the latest level-0 backup unless you use the `-O` option. If you use the `-O` and `-t` options, all backups from before the specified time are removed even if they are needed for restore. If you use the `-O` and `-i` options, all backups from before the specified interval are removed even if they are needed for restore.

For example, to expire all backups, specify the following:

```
onsmsync -O -g 0
```



Important: If you use the -O option with the -t, -i, or -g options, you might accidentally delete critical backups, making restores impossible.

Monitor the performance of ON-Bar and the storage managers

You can monitor the performance of ON-Bar and your storage manager. You can specify the level of performance monitoring and have the statistics print to the ON-Bar activity log. The `BAR_PERFORMANCE` configuration parameter specifies whether to gather statistics. The following statistics are gathered:

- Total time spent in XBSA calls.
- Total time spent in Archive API calls.
- Time spent by ON-Bar in transferring data to and from XBSA (storage manager calls).
- Time spent by ON-Bar in transferring data between ON-Bar to HCL Informix®.
- Amount of data transferred to or from the XBSA API.
- Amount of data transferred to or from the Archive API.

Set ON-Bar performance statistics levels

To specify the level of performance statistics that are printed to the ON-Bar activity log, set the `BAR_PERFORMANCE` configuration parameter in the `onconfig` file.

For example, the `BAR_PERFORMANCE 1` setting displays the time spent transferring data between the HCL Informix® instance and the storage manager.

See [BAR_PERFORMANCE configuration parameter on page 236](#) for information about the options for this parameter.

View ON-Bar backup and restore performance statistics

To view ON-Bar performance results, open the ON-Bar activity log file, `bar_act.log`.

The location of the `bar_act.log` file is set by the `BAR_ACT_LOG` configuration parameter.

When the `BAR_PERFORMANCE` configuration parameter is set to 1 or 3, the activity report shows a transfer rate report.

Figure 20. Sample transfer rate performance in the ON-Bar activity log

```

2013-06-03 15:38:02 8597 8595 Begin restore logical log 310 (Storage Manager
copy ID: 28206 0).
2013-06-03 15:38:03 8597 8595 Completed restore logical log 310.
2013-06-03 15:38:08 8597 8595 Completed logical restore.
2013-06-03 15:38:19 8597 8595 PERFORMANCE INFORMATION

                                TRANSFER RATES
-----
| OBJECT          | XBSA API | SERVER API |
| NAME           | xfer-kbytes | xfer-time | xfer-kbytes | xfer-time |
|                |          | RATIO(kb/s) | RATIO(kb/s) | API-TIME |
|-----|-----|-----|-----|-----|
| 309            | 62       | 0.479      | 129       | 1.078 |
| 310            | 62       | 0.407      | 152       | 1.098 |
| rootdbs       | 5828    | 0.618      | 9436     | 1.864 |
| datadbs01     | 62       | 0.488      | 127       | 1.768 |
| datadbs02     | 62       | 0.306      | 203       | 1.568 |
| datadbs03     | 62       | 0.304      | 204       | 1.574 |
| datadbs04     | 62       | 0.306      | 202       | 1.563 |
| datadbs05     | 62       | 0.315      | 197       | 1.585 |
| datadbs06     | 62       | 0.310      | 200       | 1.583 |
| ...           | ...     | ...       | ...     | ...   |
| PID = 8597    | 14722   | 26.758    | 550      | 107.476 |
|                | 14756   | 10.678    | 1382     | 15.829 |
-----
2013-06-03 15:38:19 8597 8595 PERFORMANCE INFORMATION

                                PERFORMANCE CLOCKS
-----
| ITEM DESCRIPTION | TIME SPENT |
|-----|-----|
| Time to Analyze ixbar file | 0.000 |
-----

```

When the `BAR_PERFORMANCE` configuration parameter is set to 2 or 3, the activity report has microsecond timestamps.

Figure 21. Sample processing rates, in microseconds, in the ON-Bar activity log

```

2013-06-03 16:34:04 15272 15270 /usr/informix/bin/onbar_d complete,
returning 0 (0x00)
2013-06-03 16:45:11.608424 17085 17083 /usr/informix/bin/onbar_d -r -w
2013-06-03 16:46:07.926097 17085 17083 Successfully connected to Storage Manager.
2013-06-03 16:46:08.590675 17085 17083 Begin salvage for log 311.
2013-06-03 16:48:07.817487 17085 17083 Completed salvage of logical log 311.
2013-06-03 16:48:08.790782 17085 17083 Begin salvage for log 312.
2013-06-03 16:48:10.129534 17085 17083 Completed salvage of logical log 312.
2013-06-03 17:06:00.836390 17085 17083 Successfully connected to Storage Manager.
...
2013-06-03 17:07:26.357521 17085 17083 Completed cold level 0 restore datadbs07.
2013-06-03 17:07:28.268562 17085 17083 Begin cold level 0 restore datadbs08
(Storage Manager copy ID: 28122 0).
2013-06-03 17:07:29.378405 17085 17083 Completed cold level 0 restore datadbs08.

```

Related reference

[BAR_ACT_LOG configuration parameter on page 222](#)

[BAR_PERFORMANCE configuration parameter on page 236](#)

Related information

[bar_act.log file: ON-Bar activity log on page 23](#)

ON-Bar catalog tables

These topics describe the ON-Bar tables that are stored in the **sysutils** database.

ON-Bar uses these tables for tracking backups and performing restores.

You can query these tables for backup and restore data to evaluate performance or identify object instances for a restore.

The bar_action table

The **bar_action** table lists all backup and restore actions that are attempted against an object, except during certain types of cold restores. Use the information in this table to track backup and restore history.

Table 22. bar_action table columns

Column name	Type	Explanation
act_aid	SERIAL	Action identifier. A unique number within the table. Can be used with act_oid column to join with the bar_instance table.
act_oid	INTEGER	Object identifier. Identifies the backup object against which a backup or restore attempt is made. Can be used with act_aid to join with bar_instance . The act_oid column of the bar_action table equals the obj_oid column of the bar_object table.
act_type	SMALLINT	Identifies the attempted action: 1 for backup, 2 for restore, 3 for a foreign or imported restore, 4 for a fake backup, 5 for a whole-system backup, 6 for a whole-system restore, 7 for expired or deleted objects, 8 for an external restore.
act_status	INTEGER	Identifies the result of the action: 0 if successful, otherwise an ON-Bar-specific error code. For more information, see ON-Bar messages and return codes on page 122 .
act_start	DATETIME YEAR TO SECONDS	The date and time when the action began.
act_end	DATETIME YEAR TO SECONDS	The date and time when the action finished.

The bar_instance table

The **bar_instance** table contains descriptions of each object that is backed up.

ON-Bar writes a record to the **bar_instance** table for each successful backup. ON-Bar might later use the information for a restore operation. For example, if you specify a level-2 backup, ON-Bar uses this table to ensure that a level-1 backup was done previously.

Table 23. bar_instance table columns

Column name	Type	Explanation
ins_aid	INTEGER	Action identifier. Identifies the successful action that created this instance of the backup object. Combined with ins_oid , can be used to join with the bar_action table.
ins_oid	INTEGER	Object identifier. Identifies the affected object. Can be used to join with the bar_object table. Combined with ins_aid , can be used to join with the bar_action table.
ins_time	INTEGER	Timestamp (real clock time). The database server uses this value when it creates the next-level backup. Value represents the number of seconds since midnight, January 1, 1970, Greenwich mean time. The ins_time value is 0.
rsam_time	INTEGER	The backup checkpoint time stamp. Not a clock time. The database server uses this value when it creates the next level backup.
seal_time	INTEGER	The time that the log file was sealed after a backup completed.
prev_seal_time	INTEGER	The time that the previous log file was sealed.
ins_level	SMALLINT	Level of the backup action: 0 for a complete backup, 1 for a backup of any changes to this object since its last level-0 backup, 2 for a backup of any changes since the last level-1 backup. This value is always 0 for logical-log backups.
ins_copyid_hi	INTEGER	The high bits of the instance copy identifier. Combined with ins_copyid_lo , it is a unique value that the storage manager assigns to link the ON-Bar object identifier with the storage-manager object identifier.
ins_copyid_lo	INTEGER	The low bits of the instance copy identifier. Combined with ins_copyid_hi , it is a unique value that the storage manager assigns to link the ON-Bar object identifier with the storage-manager object identifier.
ins_req_aid	INTEGER	Stores the required action ID for a backup object. Used in a restore to determine which level-0 backup goes with the level-1 backup, and which level-1 backup goes with the level-2 backup. For a level-0 backup, the value of ins_req_aid is the same as ins_aid in this table. For example, if this backup is level-1, ins_req_aid holds the action ID of the corresponding level-0 backup of this object.
ins_logstream	INTEGER	Not used.

Table 23. bar_instance table columns (continued)

Column name	Type	Explanation
ins_first_log	INTEGER	In a standard backup, identifies the first logical log required to restore from this backup.
ins_chpt_log	INTEGER	The logical log that contains the archive checkpoint in the dbspace backup.
ins_last_log	INTEGER	In a standard backup, identifies the last logical log required to restore from this backup.
ins_partial	INTEGER	Partial flag from salvage.
ins_sm_id	INTEGER	Not used.
ins_sm_name	CHAR(128)	Not used.
ins_verify	INTEGER	Value is 1 if the backup is verified. Value is 0 if the backup is not verified.
ins_verify_date	DATETIME YEAR TO SECOND	The current date is inserted when a backup is verified. If this backup has not been not verified, a dash represents each date and time.
ins_backup_order	INTEGER	The order in which backups occur.

The bar_ixbar table

The **bar_ixbar** table, which stores a history of all unexpired successful backups in all timelines, is maintained and used by the onsmsync utility only.

The schema of the **bar_ixbar** table is identical to the schema of the **bar_syncdeltab** table, except for its primary key.

Table 24. bar_ixbar table columns

Column name	Type	Explanation
ixb_sm_id	INTEGER	Storage-manager instance ID. Created from BAR_SM in \$ONCONFIG or %ONCONFIG%.
ixb_copyid_hi	INTEGER	The high bits of the instance copy identifier. Combined with ixb_copyid_lo , it is a unique value that the storage manager assigns to link the ON-Bar object identifier with the storage-manager object identifier.
ixb_copyid_lo	INTEGER	The low bits of the instance copy identifier. Combined with ixb_copyid_hi , it is a unique value that the storage manager assigns to link the ON-Bar object identifier with the storage-manager object identifier.

Table 24. bar_ixbar table columns (continued)

Column name	Type	Explanation
ixb_aid	INTEGER	Action identifier, Identifies the successful action that created this instance of the backup object.
ixb_old	INTEGER	Object identifier. Identifies the affected object.
ixb_time	INTEGER	Time stamp (real clock time). The database server uses this value when it creates the next-level backup. Value represents the number of seconds since midnight, January 1, 1970, Greenwich mean time.
ixb_prevtime	INTEGER	Time stamp (real clock time). This value specifies the time stamp of the previous object. Value represents the number of seconds since midnight, January 1, 1970 Greenwich mean time.
ixb_rsam_time	INTEGER	The backup checkpoint time stamp. Not a clock time. The database server uses this value when it creates the next level backup.
ixb_act_start	datetime year to second	The date and time when the action began.
ixb_act_end	datetime year to second	The date and time when the action finished.
ixb_level	SMALLINT	Level of the backup action: 0 for a complete backup, 1 for a backup of any changes to this object since its last level-0 backup, 2 for a backup of any changes since the last level-1 backup. This value is always 0 for logical-log backups.
ixb_req_aid	INTEGER	Stores the required action ID for a backup object. Used in a restore to determine which level-0 backup Goes with the level-1 backup, and which level-1 backup goes with the level-2 backup. For a level-0 backup, the value of ixb_req_aid is the same as ixb_aid in this table. For example, if this backup is level-1, ixb_req_aid holds the action ID of the corresponding level-0 backup of this object.
ixb_first_log	INTEGER	In a standard backup, identifies the first logical log. Required to restore from this backup.
ixb_chpt_log	INTEGER	The ID of the log that contains the rsam_time checkpoint. Used during back up to verify that logs needed for restore are backed up.
ixb_last_log	INTEGER	Log ID of the last log needed during logical restore for this storage space to restore it to the time of the backup.
ixb_lbuflags	INTEGER	Flags describing log backup.

Table 24. bar_ixbar table columns (continued)

Column name	Type	Explanation
ixb_verify	INTEGER	Value is 1 if the backup is verified. Value is 0 if the backup is not verified.
ixb_verify_date	datetime year to second	The current date is inserted when a backup is verified. If this backup has not been verified, a dash represents each date and time.
ixb_sm_name	VARCHAR(128)	Storage-manager instance name. Created from the BAR_SM_NAME parameter in the <code>onconfig</code> file.
ixb_srv_name	VARCHAR(128)	The database server name. Used to ensure that objects are restored to the correct database server. Used when multiple database servers are on the node to ensure that objects are restored in the database server instance to which the object belongs. The database server name can be up to 128 characters.
ixb_obj_name	VARCHAR(128)	The user name for the object. The name can be up to 128 characters.
ixb_obj_type	CHAR(2)	Backup object type: CD critical dbspace L logical log ND noncritical dbspace or sbspace R rootdbs B blobospace

Related reference

[The bar_syncdeltab table on page 122](#)

The bar_object table

The **bar_object** table contains descriptions of each backup object. This table is a list of all storage spaces and logical logs from each database server for which at least one backup attempt was made.

Table 25. `bar_object` table columns

Column name	Type	Explanation
<code>obj_srv_name</code>	VARCHAR(128,0)	The database server name. Used to ensure that objects are restored to the correct database server. Used when multiple database servers are on the node to ensure that objects are restored in the database server instance to which the object belongs. The database server name can be up to 128 characters.
<code>obj_oid</code>	SERIAL	The object identifier. A unique number within the table. Can be used to join with the <code>bar_action</code> and <code>bar_instance</code> tables.
<code>obj_name</code>	VARCHAR(128,0)	The user name for the object. The name can be up to 128 characters.
<code>obj_type</code>	CHAR(2)	Backup object type: CD critical dbspace L logical log ND noncritical dbspace or sbspace R rootdbs B blobospace

The `bar_server` table

The `bar_server` table lists the database servers in an installation. This table is used to ensure that backup objects are returned to their proper places during a restore.

Table 26. `bar_server` table columns

Column name	Type	Explanation
<code>srv_name</code>	VARCHAR(128,0)	DBSERVERNAME value specified in the <code>onconfig</code> file. Database server name can be up to 128 characters.
<code>srv_node</code>	CHAR(256)	Host name of the computer where the database server resides.

Table 26. bar_server table columns (continued)

Column name	Type	Explanation
		The host name can be up to 256 characters.
srv_synctime	INTEGER	The time onsmsync was run.

The bar_syncdeltab table

The **bar_syncdeltab** table is maintained and used by the onsmsync utility only. This table is empty except when onsmsync is running.

The schema of the **bar_syncdeltab** table is identical to the schema of the **bar_ixbar** table, except for its primary key.

Related reference

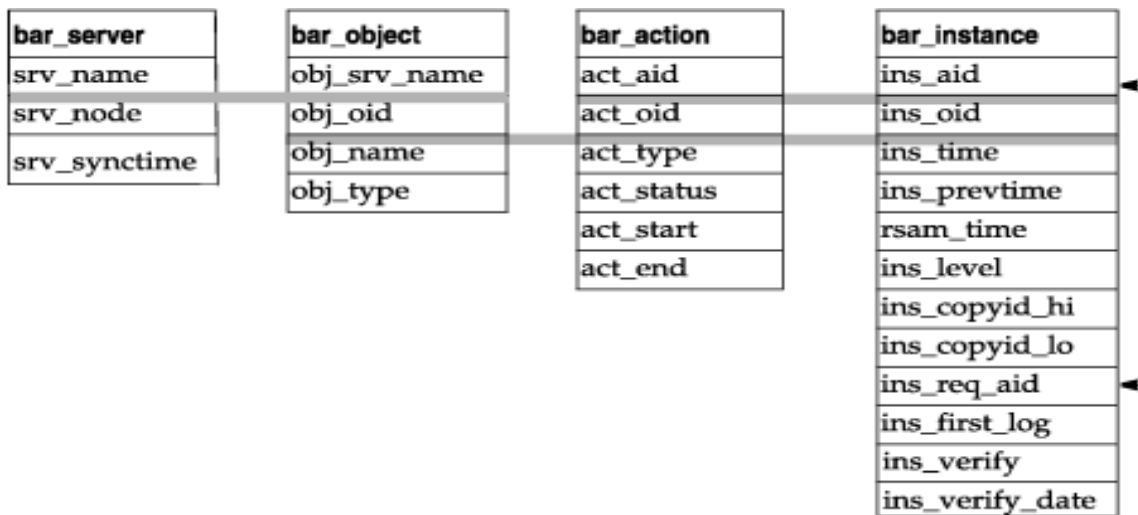
[The bar_ixbar table on page 118](#)

ON-Bar catalog map

This topic contains an example mapping between ON-Bar tables.

The following figure maps the ON-Bar tables on HCL Informix®. In this figure, the gray lines show the relations between tables. The arrows show that the **ins_req_aid** value must be a valid **ins_aid** value.

Figure 22. ON-Bar catalog map on Informix®



ON-Bar messages and return codes

ON-Bar prints informational, progress, warning, and error messages to the ON-Bar activity log file. ON-Bar return codes indicate the status of the command.

Related reference

[onbar -m syntax: Monitoring recent ON-Bar activity on page 53](#)

Related information

[bar_act.log file: ON-Bar activity log on page 23](#)

Message format in the ON-Bar message log

Messages in the ON-Bar activity log file contain timestamps, process IDs, and explanatory text.

A message in the ON-Bar activity log file, `bar_act.log`, has the following format:

```
timestamp process_id parent_process_id message
```

The following table describes each field in the message. No error message numbers appear in the ON-Bar activity log.

Table 27. ON-Bar message format

Message field	Description
<i>timestamp</i>	Date and time when ON-Bar writes the message.
<i>process_id</i>	The number that the operating system uses to identify this instance of ON-Bar.
<i>parent_process_id</i>	The number that the operating system uses to identify the process that executed this instance of ON-Bar.
<i>message</i>	The ON-Bar message text.

The following example illustrates a typical entry in the ON-Bar activity log:

```
1999-08-18 10:09:59 773    772 Completed logical restore.
```



Important: If you receive an XBSA error message, consult the storage-manager logs for more details.

Timestamps when storage managers hang

If a storage manager process hangs, the timestamp for the process in the ON-Bar activity log is inaccurate. An asterisk symbol is appended to the timestamp and the message indicates how long the storage manager process hung. The following example shows that the storage manager process hung for two minutes starting at 10:27. At 10:29, the storage manager completed the backup.

```
2013-02-26 10:27:10* 13410 25695 (-43085) WARNING: BAR_TIMEOUT Storage
                Manager Progress may be stalled for at least 2 minutes.

2013-02-26 10:29:12 13410 25695 Completed level 0 backup dbspace1 (Storage Manager
                copy ID: 1509564809 0).
```

Related reference

[onbar -m syntax: Monitoring recent ON-Bar activity on page 53](#)

Message numbers

The ON-Bar message numbers range from -43000 to -43421.

The following table lists the ON-Bar message groups. Because message numbers do not display in the activity log, the best way to find information about ON-Bar messages is to search for the message text in the error messages file, which is in the subdirectory for your locale under the `$INFORMIXDIR/msg` directory.

ON-Bar message type	Message numbers
ON-Bar usage	-43000 to -43007 and -43357
Options checking	-43010 to -43034
Permission checking	-43035 to -43039
Emergency boot file interface	-43040 to -43059
<code>onconfig</code> file interface	-43060 to -43074
Operating system interface	-43075 to -43099
Database server interface	-43100 to -43229
Back up and restore status	-43230 to -43239
onbar-worker processes	-43240 to -43254
XBSA interface	-43255 to -43301
onsmsync	-43302 to -43319
archecker	-43320 to -43334
ondblog	-43400 to -43424

ON-Bar return codes

You can troubleshoot problems by viewing activity log messages that are applicable for particular return codes.

The following table shows the ON-Bar return codes for all HCL Informix® database servers. These return codes are accompanied by messages in the ON-Bar activity log. For details about an error, review the activity log before you call Technical Support.

Table 28. Common ON-Bar return codes

Decimal value	ON-Bar return code description
2 through 34	These return codes are produced by XBSA. For more information, consult your storage-manager documentation and log files.
100	<p>ON-Bar cannot find something in sysutils, the emergency boot file, or storage-manager catalogs that it needs for processing.</p> <p>Check the ON-Bar activity log for messages that say what could not be found and try to resolve that problem. If the problem recurs, contact Technical Support.</p>
104	<p>Adstar Distributed Storage Manager (ADSM) is in generate-password mode.</p> <p>ON-Bar does not support ADSM running in generate-password mode. For information about changing the ADSM security configuration, refer to your ADSM manual.</p>
115	A critical dbspace is not in the set of dbspaces being cold-restored.
116	The onsmsync utility is already running.
117	The information contained in the sysutils database and the emergency boot file are inconsistent.
118	An error trying to commit a backup object to the storage manager. An error in tenant restore occurred since another tenant restore was in progress.
119	An error in tenant restore occurred since another restore was in progress.
120	The transport buffer size has changed since this object was last backed up. This object cannot be restored. Set the transport-buffer size to its original value and try the restore again. An internal On-Bar error occurred when starting a tenant restore.
121	ON-Bar was unable to determine the list of dbspaces.
122	<p>Deadlock detected.</p> <p>The ON-Bar command is contending with another process. Try the ON-Bar command again.</p>
123	<p>The root dbspace was not in the cold restore.</p> <p>You cannot perform a cold restore without restoring the root dbspace. To resolve the problem, try one of the following procedures:</p> <ul style="list-style-type: none"> • Bring the database server to quiescent or online mode and restore just the storage spaces that need to be restored. • If the database server is offline, issue the onbar -r command to restore all the storage spaces. • Make sure that the root dbspace and other critical dbspaces are listed on the command line or in the -f filename.
124	The buffer had an incomplete page during the backup.

Table 28. Common ON-Bar return codes (continued)

Decimal value	ON-Bar return code description
	For assistance, contact Technical Support.
126	<p>Error processing the emergency boot file.</p> <p>Check the ON-Bar activity log for descriptions of the problem and the emergency boot file for corruption such as non-ASCII characters or lines with varying numbers of columns. If the source of the problem is not obvious, contact Technical Support.</p>
127	<p>Could not write to the emergency boot file.</p> <p>Often, an operating-system error message accompanies this problem. Check the permissions on the following files and directories:</p> <ul style="list-style-type: none"> • \$INFORMIXDIR/etc on UNIX™ or %INFORMIXDIR%\etc on Windows™ • The emergency boot file
128	<p>Data is missing in the object description.</p> <p>For assistance, contact Technical Support.</p>
129	<p>ON-Bar received a different object for restore than it had expected. (The backup object did not match.) The requested backup object might have been deleted or expired from the storage manager.</p> <p>Run onsmsync to synchronize the sysutils database, emergency boot file, and storage-manager catalogs. For assistance, contact Technical Support.</p>
130	<p>Database server is not responding.</p> <p>The database server probably failed during the backup or restore. Run the onstat - command to check the database server status and then:</p> <ul style="list-style-type: none"> • If the operation was a cold restore, restart it. • If the operation was a backup or warm restore, restart the database server and try the backup or warm restore again.
131	<p>A failure occurred in the interface between ON-Bar and the database server.</p> <p>For assistance, contact Technical Support.</p>
132	<p>Function is not in the XBSA shared library.</p> <p>Verify that you are using the correct XBSA for the storage manager. For information, consult your storage-manager manual.</p>
133	Failed to load the XBSA library functions.

Table 28. Common ON-Bar return codes (continued)

Decimal value	ON-Bar return code description
	Verify that you are using the correct XBSA for the storage manager. Ensure that the <code>BAR_BSALIB_PATH</code> value in the <code>onconfig</code> file points to the correct location of the XBSA shared library. For information, consult your storage-manager manual.
134	<p>User wants to restore a logical-log file that is too early.</p> <p>You probably tried a point-in-log restore (<code>onbar -r -l -n</code>) after performing a separate physical restore. The specified logical log is too old to match the backups used in the physical restore. Perform either of the following steps:</p> <ul style="list-style-type: none"> • Rerun the physical restore from an older set of physical backups. • Specify a later logical log in the <code>-n</code> option when you rerun the point-in-log restore. To find the earliest logical log that you can use, look at the emergency boot file. For assistance, contact Technical Support.
136	<p>ON-Bar cannot warm restore the critical dbspaces.</p> <p>Perform either of the following steps:</p> <ul style="list-style-type: none"> • Reissue the warm-restore command without listing any critical dbspaces. • Shut down the database server and perform a cold restore.
137	<p>The <code>MAX_DBSPACE_COUNT</code> was exceeded.</p> <p>For assistance, contact Technical Support.</p>
138	<p>An XBSA error occurred.</p> <p>Verify that you are using the correct XBSA for the storage manager. Also check the <code>bar_act.log</code> for XBSA error messages. For information, consult your storage-manager manual.</p>
139	<p>Either the XBSA version is missing from the <code>sm_versions</code> file or the incorrect XBSA version is in the <code>sm_versions</code> file.</p> <p>Insert the correct XBSA version into the <code>sm_versions</code> file. For more information, consult your storage-manager manual.</p>
140	<p>A fake backup failed.</p> <p>Try the fake backup with the <code>onbar -b -F</code> command again. Only HCL Informix® supports fake backups. If the fake backup fails again, contact Technical Support.</p>
141	<p>ON-Bar received an operating-system signal. Most likely, the user entered the Ctrl-C command to stop an ON-Bar process.</p>

Table 28. Common ON-Bar return codes (continued)

Decimal value	ON-Bar return code description
	Fix the cause of the interruption and then try the ON-Bar command again.
142	<p>ON-Bar was unable to open a file.</p> <p>Verify that the named file exists and that the permissions are correct. Check the ON-Bar activity log for an operating-system error message.</p>
143	<p>ON-Bar was unable to create a child process.</p> <p>If BAR_MAX_BACKUP is not 0, ON-Bar could not create child processes to perform the parallel backup or restore. The operating system probably ran out of resources. Either not enough memory is available to start a new process or no empty slot exists in the process table.</p> <p>Check the operating-system logs, the ON-Bar activity log, or the console.</p>
144	<p>The log backup was stopped because one or more blobspaces were down.</p> <p>Attempt to restore the blobspace. If the restore fails, try the log backup with the onbar -l -O command again. Executing this command might make the blobspace unrestoreable.</p>
145	<p>ON-Bar was unable to acquire more memory space.</p> <p>Wait for system resources to free up and try the ON-Bar command again.</p>
146	<p>ON-Bar was unable to connect to the database server.</p> <p>The network or the database server might be down. For assistance, contact Technical Support.</p>
147	<p>ON-Bar was unable to discover any storage spaces or logical logs to back up or restore.</p> <p>For example, if you specify a point-in-time restore but use a <i>datetime</i> value from before the first standard backup, ON-Bar cannot build a list of storage spaces to restore. This return code also displays if you specify a whole-system restore without having performed a whole-system backup.</p> <p>Verify that the database server and the storage spaces are in the correct state for the backup or restore request. Contact Technical Support.</p>
148	<p>An internal SQL error occurred.</p> <p>Provide Technical Support with the information from the ON-Bar activity log.</p>
149	<p>Either you entered the wrong ON-Bar syntax on the command line or entered an invalid or incorrect <i>datetime</i> value for your GLS environment.</p>

Table 28. Common ON-Bar return codes (continued)

Decimal value	ON-Bar return code description
	Check the command that you tried against the usage message in the ON-Bar activity log. If that does not help, then try the command with quotes around the <i>datetime</i> value again. If your database locale is not English, use the GL_DATE or GL_DATETIME environment variables to set the date and time format.
150	<p>Error collecting data from the <code>onconfig</code> file.</p> <p>Check the permissions, format, and values in the <code>onconfig</code> file. Check that the ONCONFIG environment variable is set correctly.</p>
151	<p>The database server is in an incorrect state for this backup or restore request, or an error occurred while determining the database server state.</p> <p>Either you attempted an operation that is not compatible with the database server mode or ON-Bar is unable to determine the database server state. For example, you cannot do a physical backup with the database server in recovery mode.</p> <p>Check the error message in the ON-Bar activity log. If an ASF error occurred, the following message displays in the ON-Bar activity log:</p> <pre data-bbox="370 978 1458 1010">Fatal error initializing ASF; asfcode = code</pre> <p>To determine the cause of the ASF error, refer to the ASF error code in this message and repeat the backup or restore command. If an ASF error did not occur, change the database server state and repeat the backup or restore command.</p>
152	<p>ON-Bar cannot back up the logical logs.</p> <p>The logical logs are not backed up for either of the following reasons:</p> <ul data-bbox="418 1335 1458 1440" style="list-style-type: none"> • If another log backup is currently running. • If you perform a logical-log backup with the LTAPEDEV parameter set to <code>/dev/null</code> (UNIX™) or <code>NUL</code> (Windows™). <p>You receive this return code when no log backups can be done.</p> <p>To enable log backups, change the LTAPEDEV parameter to a valid value.</p>
153	<p>ON-Bar cannot set the process group ID. If <code>BAR_MAX_BACKUP</code> is set to any value other than 1 and ON-Bar encounters an error setting the process group ID, this value is returned.</p> <p>This message is a warning of a possible operating-system problem.</p>
154	The ON-Bar user does not have the correct permissions.

Table 28. Common ON-Bar return codes (continued)

Decimal value	ON-Bar return code description
	You must be user root or informix or a member of the bargroup group on UNIX™ or a member of the Informix-Admin group on Windows™ to execute ON-Bar commands.
155	The INFORMIXSERVER environment variable is not set. Set the INFORMIXSERVER environment variable to the correct database server name.
156	Backup or restore was not performed because the LTAPEDEV parameter value is not valid. If LTAPEDEV is not set or <code>/dev/null</code> on UNIX™, or if it is <code>NUL</code> on Windows™, the logical logs are not backed up, and ON-Bar returns warning 152.
157	Error attempting to set the INFORMIXSHMBASE environment variable to -1. ON-Bar could not set INFORMIXSHMBASE to -1. For assistance, contact either the system administrator or Technical Support.
158	An internal ON-Bar error occurred. Contact Technical Support.
159	An unexpected error occurred. Contact Technical Support.
160	External restore failed. To determine what went wrong with the external restore, look at the <code>bar_act.log</code> and the <code>online.log</code> files. Ensure that you already performed the manual part of the external restore before you try the <code>onbar-r -e</code> command again to complete the external restore. If that does not work, try the external restore from a different external backup.
161	Restarted restore failed. Verify that <code>RESTARTABLE_RESTORE</code> is set to <code>ON</code> and try the original restore again. For more information, check the ON-Bar activity log and database server message logs.
162	The ON-Bar log file cannot be a symbolic link. Remove the symbolic link or change the <code>onconfig</code> file so that the ON-Bar parameters <code>BAR_DEBUG_LOG</code> or <code>BAR_ACT_LOG</code> point to non-symbolic linked files.
163	The ON-Bar log file must be owned by user informix .

Table 28. Common ON-Bar return codes (continued)

Decimal value	ON-Bar return code description
	Change the ownership of the log file to be owned by user informix or change the BAR_ACT_LOG or BAR_DEBUG_LOG values in the <code>onconfig</code> file to point to different log files.
164	Unable to open file. The file or its directory permissions prevent it from being created or opened. Verify the permissions on the file and its directory.
177	An online dbspace was restored. This return code notifies the user that the -O option overrode the internal checks in ON-Bar. You do not need to take any action.
178	The logical log was backed up while one or more blobspaces were down. This return code notifies the user that the -O option overrode the internal checks in ON-Bar. Examine the data in the blobspace to determine which simple large objects you need to recreate. These blobspaces might not be restorable. For assistance, contact Technical Support.
179	ON-Bar created the chunk needed to restore the dbspace. This return code notifies the user that the -O option overrode the internal checks in ON-Bar. You do not need to take any action.
180	ON-Bar could not create the chunk needed to restore the dbspace. Create the chunk file manually. Try the restore without the -O option again.
181	ON-Bar expired an object that was needed for a backup or restore. The onmsync utility expired an object that might be needed for a restore. You probably specified onmsync with the -O option. If you used the -O option by mistake, contact Technical Support to recover the object from the storage manager.
183	ON-Bar could not obtain the logical-log unique ID from the storage manager. The backup of the specified logical log is missing. Query your storage manager to determine if the backup of the specified logical-log file exists and if it is restorable.
247	On UNIX™, look in <code>/tmp/bar_act.log</code> and the file that the BAR_ACT_LOG parameter points to for clues. (The onbar-merger writes to <code>/tmp/bar_act.log</code> until it has enough information to read the <code>onconfig</code> file.) Resolve the problems that the <code>bar_act.log</code> describes and try the cold restore again. If the cold restore still fails, contact Technical Support.

Table 28. Common ON-Bar return codes (continued)

Decimal value	ON-Bar return code description
252	For assistance, contact Technical Support.

ontape backup and restore system

Related reference

[LTAPEBLK configuration parameter on page 244](#)

[LTAPEDEV configuration parameter on page 246](#)

[LTAPESIZE configuration parameter on page 247](#)

[TAPEBLK configuration parameter on page 250](#)

[TAPEDEV configuration parameter on page 251](#)

[TAPESIZE configuration parameter on page 253](#)

Configure ontape

You configure the ontape utility by setting configuration parameters for backups of storage spaces and logical logs.

You can also set the IFX_BAR_USE_DEDUP environment variable to optimize the backup images for deduplication devices.

Related reference

[IFX_BAR_USE_DEDUP environment variable on page 242](#)

[Comparison of the ON-Bar and ontape utilities on page 11](#)

Set configuration parameters for the ontape utility

The ontape utility uses eight configuration parameters in the `onconfig` file. Two of the configuration parameters specify filter programs for transforming data during backup and restore; the other six are used to create storage-space and logical-log backups.

The `onconfig` file is located in the `$INFORMIXDIR/etc` directory. You specify that file in the **ONCONFIG** environment variable. For a description of the **ONCONFIG** environment variable and instructions on how to set it, see the *HCL® Informix® Guide to SQL: Reference*.

Data transformation filter parameters for ontape

The `BACKUP_FILTER` and `RESTORE_FILTER` configuration parameters specify the names of external programs that you can use to transform data before backup and after a restore.

BACKUP_FILTER

Specifies the location and name of an external filter program used in data transformation. This filter transforms data before backing it up, such as compressing it. The transformed data is then backed up and stored as a single file. The filter path points to the `$(INFORMIXDIR)/bin` directory by default, or an absolute path of the program

RESTORE_FILTER

Specifies the location and name of an external filter program used in data transformation. This filter transforms data back to its original state before the backup, such as extracting it, before returning the data to the server. The filter path points to the `$(INFORMIXDIR)/bin` directory by default, or an absolute path of the program



Prerequisite: The data must have previously been transformed with the `BACKUP_FILTER` parameter.

See [BACKUP_FILTER configuration parameter on page 221](#) and [RESTORE_FILTER configuration parameter on page 249](#) for syntax and usage information, which is the same for ON-Bar and ontape.

Tape and tape device parameters for ontape

The first set of configuration parameters specifies the characteristics of the tape device and tapes for storage-space backups; the second set specifies the characteristics of the tape device and tapes for logical-log backups.

The following list shows backup tape devices and their associated tape parameters.

TAPEDEV

The absolute path name of the tape device or directory file system that is used for storage-space backups. Specify the destination where ontape writes storage space data during an archive and the source from which ontape reads data during a restore.

To configure ontape to use **stdio**, set `TAPEDEV` to `STDIO`.

When backing up to or restoring from a cloud environment, use the following syntax for the `TAPEDEV` configuration parameter:

```
TAPEDEV 'local_path, keep=option, cloud=cloud_vendor, url=url'
```

- **local_path** is the complete path name of the directory where storage spaces backup objects are stored temporarily.
- **option** can be set to *yes* or *no*. If **keep** is set to *yes*, the ontape utility retains the backup objects in the local directory. If **keep** is set to *no*, the backup objects are deleted after they are transferred to or from the cloud storage location.
- **cloud_vendor** is the name of the cloud storage vendor.
- **url** is the cloud storage location where the storage space backup data is stored persistently.

TAPEBLK

The block size of the tapes used for storage-space backups, in kilobytes.

TAPESIZE

The size of the tapes used for storage-space backups, in kilobytes.

The following list shows the logical-log tape devices and their associated tape parameters.

LTAPEDEV

The logical-log tape device or a directory of a file system.

When backing up to or restoring from a cloud environment, use the following syntax for the LTAPEDEV configuration parameter:

```
LTAPEDEV 'local_path, keep=option, cloud=cloud_vendor, url=url'
```

- **local_path** is the complete path name of the directory where log backup objects are stored temporarily.
- **option** can be set to *yes* or *no*. If **keep** is set to *yes*, the ontape utility retains the backup objects in the local directory. If **keep** is set to *no*, the backup objects are deleted after they are transferred to or from the cloud storage location.
- **cloud_vendor** is the name of the cloud storage vendor.
- **url** is the cloud storage location where the log backup data is stored persistently.

LTAPEBLK

The block size of tapes used for logical-log backups, in kilobytes.

LTAPESIZE

The size of tapes used for logical-log backups, in kilobytes.

The following topics contain information about how to set the tape-device, tape-block-size, and tape-size parameters for both storage-space and logical-log backups.

Related information

[Back up to Amazon Simple Storage Service on page 149](#)

Set the tape-device parameters

Specify values for TAPEDEV and LTAPEDEV in the following ways:

- Use separate tape devices, when possible.
- Use symbolic links.
- Specify a directory of a file system.
- For tape devices, specify `/dev/null`.
- Rewind tape devices.
- Configure parameters to perform backup to a cloud.

The following sections explain each of these points.

Related information

[Back up to Amazon Simple Storage Service on page 149](#)

Specify separate devices for storage-space and logical-log backups

When backing up to a tape device, specify different devices for the LTAPEDEV and TAPEDEV parameters in the `onconfig` file. You can schedule these backups independently of each other. You can create a backup on one device at the same time you continuously back up the logical-log files on the other.

If you specify the same device for the LTAPEDEV and TAPEDEV, the logical log can fill, which causes the database server to stop processing during a backup. When this happens, you have two options.

- Stop the backup to free the tape device and back up the logical-log files.
- Leave normal processing suspended until the backup completes.

Precautions to take when you use one tape device

When only one tape device exists and you want to create backups while the database server is online, take the following precautions:

- Configure the database server with a large amount of logical-log space through a combination of many or large logical-log files. (See your *HCL® Informix® Administrator's Guide*.)
- Store all explicitly created temporary tables in a dedicated dbspace and then drop the dbspace before backing up.
- Create the backup when low database activity occurs.
- Free as many logical-log files as possible before you begin the backup.

The logical log can fill up before the backup completes. The backup synchronizes with a checkpoint. A backup might wait for a checkpoint to synchronize activity, but the checkpoint cannot occur until all virtual processors exit critical sections. When database server processing suspends because of a full logical-log file, the virtual processors cannot exit their critical sections and a deadlock results.

Specify tape devices as symbolic links

You can specify the values of LTAPEDEV and TAPEDEV as symbolic links. Using symbolic links enables you to switch to other tape or tape-compatible devices without changing the path name in the `onconfig` file. For example, you can specify the following symbolic link for tape device `/dev/rst0`:

```
ln -s /dev/rst0 /dbfiles/logtape
```

When you set the LTAPEDEV configuration parameter, as the following example shows, you can switch to a different device without changing the LTAPEDEV parameter:

```
LTAPEDEV /dbfiles/logtape
```

You only need to change the symbolic link, as the following example shows:

```
ln -s /usr/backups /dbfiles/logtape
```

A user with one tape device could redirect a logical-log back up to a disk file while using the tape device for a backup.

Specify a file system directory

You can perform a storage-space (level 0, 1, or 2) archive, or a logical-log backup to a directory in the file system by using the `ontape` utility. For each storage-space archive and logical-log backup, `ontape` creates a file in the specified directory.

To specify a file system directory, set the `LTAPEDEV` and `TAPEDEV` configuration parameters to the absolute path name for the directory.

When `ontape` repeats an archive operation, it renames the existing files so that old files are not rewritten. A timestamp is added to the file name, which provides a way for related storage space or logical log files to be organized together.

To learn about the file naming schema, see [Rename existing files on page 145](#).

Specify a remote device

You can perform a storage-space or logical-log backup across your network to a remote device attached to another host computer on UNIX™ and Linux™ platforms.

You should not do a continuous backup to a remote device.

The remote device and the database server computer must have a trusted relationship so that the `rsh` or the `rlogin` utility can connect from the database server computer to the remote device computer without asking for password. You can establish a trusted relationship by configuring the `/etc/hosts.equiv` file, the `~/.rhosts` file, or any equivalent mechanism for your system on the remote device computer. If you want to use a different utility to handle the remote session than the default utility used by your platform, you can set the `DBREMOTECMD` environment variable to the specific utility that you want to use.

To specify a tape device on another host computer, use the following syntax to set the `TAPEDEV` or `LTAPEDEV` configuration parameter:

```
host_machine_name:tape_device_pathname
```

The following example specifies a tape device on the host computer `kyoto`:

```
kyoto:/dev/rmt01
```

For information about the tape size for remote devices, see [Tape size for remote devices on page 138](#).

Related information

[DBREMOTECMD environment variable \(UNIX\) on page](#)

Specify `/dev/null` for a tape device

A best practice is to not use `/dev/null` as the device when backing up. However, if you decide that you do not need to recover transactions from the logical log, you can specify `/dev/null` as a tape device for logical-log backups

When you specify `/dev/null` as a backup tape device, you can avoid the overhead of a level-0 backup that is required after some operations, such as changing the logging status of a database. Obviously, you cannot restore storage spaces from a backup to `/dev/null`.

When you specify the tape device as `/dev/null`, block size and tape size are ignored. If you set the `LTAPEDEV` configuration parameter either to or from `/dev/null`, you must restart the database server for the new setting to take effect.



Important: When you set the `LTAPEDEV` configuration parameter to `/dev/null`, the database server marks the logical-log files as backed up as soon as they become full, effectively discarding logical-log information.

Set `TAPEDEV` to `stdio`

To configure the `ontape` utility to read from standard input or write to standard output, set the `TAPEDEV` configuration parameter to **`stdio`**.

Rewind tape devices before opening and on closing

With `ontape`, you must use *rewindable* tape devices. Before reading from or writing to a tape, the database server performs a series of checks that require the rewind.

Specify the tape-block-size

Use the `TAPEBLK` and `LTAPEBLK` configuration parameters to set the largest block size, in kilobytes, that your tape device permits.

When you set the tape parameter to `/dev/null`, the corresponding block size is ignored.

The `ontape` utility does not check the tape device when you specify the block size. Verify that the tape device can read the block size that you specified. If not, you cannot restore the tape.

Specify the tape size

Use the `TAPESIZE` and `LTAPESIZE` configuration parameter parameters to specify the maximum amount of data that you can write to a tape.

To write or read the tape to the end of the device, set `TAPESIZE` and `LTAPESIZE` to 0. You cannot use this option for remote devices.

When you specify the tape device as `/dev/null`, the corresponding tape size is ignored.

Related reference[TAPESIZE configuration parameter on page 253](#)[LTAPESIZE configuration parameter on page 247](#)

Tape size for remote devices

You can estimate the amount of continuous logical-log backup data that can be stored on tape on remote devices.

When you perform a continuous logical-log backup to a remote device, the amount of data written to the tape is the smaller of LTAPESIZE and the following formula:

```
(sum of space occupied by all logical-log files on disk) -  
(largest logical-log file)
```

The I/O to the remote device completes and the database server frees the logical-log files before a log-full condition occurs.



Restriction: You cannot set tape size to 0 for remote devices.

Changing your ontape configuration

You can change the values of the TAPEDEV, TAPEBLK, and TAPESIZE or LTAPEDEV, LTAPEBLK, and LTAPESIZE configuration parameters after performing a backup and reviewing the current and new parameter values.

Prerequisites: Before you change the parameters for ontape:

- Perform a level-0 backup.
- Examine your configuration file (the file specified in `$INFORMIXDIR/etc/$ONCONFIG`) by executing `onstat -c` while the database server is running.
- If you plan to change TAPEDEV or LTAPEDEV to a different tape device, verify that the tape device can read the block size that you specify with the TAPEBLK or LTAPEBLK configuration parameter. If not, you cannot restore the tape. (The ontape utility does not check the tape device when you specify the block size.)
- Be sure that you are logged in as user **root** or **informix**.

You can change the values of parameters for ontape while the database server is online.

To change the value of TAPEDEV, TAPEBLK, and TAPESIZE or LTAPEDEV, LTAPEBLK, and LTAPESIZE:

1. From the command line, use a text editor to edit your `onconfig` file.
2. Save the file.

The change takes effect immediately. However, if you set either the TAPEDEV parameter or the LTAPEDEV parameter to `dev/null`, you must restart the database server.

Related reference

[LTAPEBLK configuration parameter on page 244](#)

[LTAPEDEV configuration parameter on page 246](#)

[LTAPESIZE configuration parameter on page 247](#)

[TAPEBLK configuration parameter on page 250](#)

[TAPEDEV configuration parameter on page 251](#)

[TAPESIZE configuration parameter on page 253](#)

Back up with ontape

These topics describe how to use the ontape utility to back up storage spaces and logical-log files, and how to change the database logging status. The ontape utility can back up and restore the largest chunk files that your database server supports. The ontape utility cannot back up temporary dbspaces and temporary sbspaces.

Summary of ontape tasks

The ontape utility lets you complete a wide variety of tasks:

- [Change database logging status on page 140](#)
- [Create a backup on page 141](#)
- [Starting a continuous logical-log file backup on page 155](#)
- [ontape utility syntax: Perform a restore on page 159](#)
- [Use external restore commands on page 177](#)

Start ontape

When you need more than one tape during a backup, the ontape utility prompts for each additional tape.

If the database server is in maintenance mode, for example, during a conversion, then the ontape utility can only be started by one of the following users:

- **root**
- **informix**
- The user who started the database server (if not the user **root** or **informix**)



Restriction: Do not start the ontape utility in background mode (that is, with the UNIX™ **&** operator on the command line). You could also need to provide input from the terminal or window. When you execute ontape in background mode, you can miss prompts and delay an operation.



The ontape utility does not include default values for user interaction, nor does it support attempts to retry. When ontape expects a yes-or-no response, it assumes that any response not recognized as a "yes" is "no".

Exit codes for ontape

The ontape utility has the following two exit codes:

0

Indicates a normal exit from ontape.

1

Indicates an exception condition.

Change database logging status

You can use the ontape utility to change the logging status of a database. Most changes in logging mode require a full level-0 backup.

You cannot change the logging mode of an ANSI-compliant database.

You can change an unbuffered logged or buffered logged database to an unlogged database without making a backup.

You can make the following logging modes changes with a level-0 backup:

- An unbuffered logged or buffered logged database to an ANSI database
- An unbuffered logged database to a buffered logged database
- A buffered logged database to an unbuffered logged database

Example

Examples

The following command changes the logging mode of a database named **stores7** to unbuffered logging:

```
ontape -s -L 0 -U stores7
```

The following command changes the logging mode of a database to ANSI-compliant logging:

```
ontape -s -L 0 -A stores7
```

The following command changes the logging mode of a database to unlogged:

```
ontape -N stores7
```

Related reference

[ontape utility syntax: Perform a backup on page 146](#)

Related information[Database-logging status on page](#)

Create a backup

These topics explain how to plan for and create backups of your database server data.

Backup levels that ontape supports

The ontape utility supports level-0, level-1, and level-2 backups.

For information about scheduling backups, see [Plan a recovery strategy on page 14](#).



Tip: Establish a backup schedule that keeps level-1 and level-2 backups small. Schedule frequent level-0 backups to avoid restoring large level-1 and level-2 backups or many logical-log backups.

Level-0 backup

When a fire or flood, for example, completely destroys a computer, you need a level-0 backup to completely restore database server data on the replacement computer. For online backups, the data on the backup tape reflects the contents of the storage spaces at the time the level-0 backup began. (The time the backup started could reflect the last checkpoint before the backup started.)

A level-0 backup can consume lots of time because ontape must write all the pages to tape.

Level-1 backup

A level-1 backup usually takes less time than a level-0 backup because you copy only part of the database server data to the backup tape.

Level-2 backup

A level-2 backup after a level-1 backup usually takes less time than another level-1 backup because only the changes made after the last level-1 backup (instead of the last level-0) get copied to the backup tape.

Back up after changing the physical schema


You must perform a level-0 backup to ensure that you can restore the data after change the physical schema.

Perform a level-0 backup after you make the following administrative changes:

- Changing the TAPEDEV or LTAPEDEV configuration parameter from `/dev/null`
- Adding logging to a database
- Adding a dbspace, blobspace, or sbspace before you can restore it with anything less than a full-system restore
- Starting mirroring for a dbspace that contains logical-log files
- Dropping a logical-log file
- Moving one or more logical-log files
- Changing the size or location of the physical log and after you set up shared memory

- Dropping a chunk before you can reuse the dbspace that contains that chunk
- Renaming a chunk during a cold restore

Consider waiting to make these changes until your next regularly scheduled level-0 backup.

 **Tip:** Although you no longer need to back up immediately after adding a logical-log file, your next backup should be level-0 because the data structures have changed.

Prepare for a backup

When you create a backup, take the following precautions:

- Avoid temp tables during heavy activity.
- Make sure enough logical-log space exists.
- Keep a copy of your configuration file.
- Verify consistency before a level-0 backup.
- Run the database server in the appropriate mode.
- Plan for operator availability.
- Synchronize with other administrative tasks.
- Do not use background mode.
- If necessary, label tapes appropriately.
- If necessary, prepare for writing to standard output.


Avoid temp tables during heavy activity

When you create a temp table during a backup while using the `ontape` utility, that table is placed in `DBSPACETEMP`. When heavy activity occurs during the backup process, the temp table can keep growing and can eventually fill up `DBSPACETEMP`. When this situation occurs, the backup stops and your monitor displays a `NO FREE DISK` error message.

Make sure enough logical-log space exists

When the total available space in the logical log amounts to less than half a single logical-log file, the database server does not create a backup. You must back up the logical-log files and attempt the backup again.

You cannot add mirroring during a backup.

 **Important:** When you use only one available tape device, make sure you back up all your logical-log files before you start your backup to reduce the likelihood of filling the logical log during the backup.

Keep a copy of your configuration file

Keep a copy of the current `onconfig` file when you create a level-0 backup. You need this information to restore database server data from the backup tape.

Verify consistency before a level-0 backup

To ensure the integrity of your backups, periodically verify that all database server data and overhead information is consistent before you create a full-system level-0 backup. You do not check this information before every level-0 backup, but we recommend that you keep the necessary tapes from the most recent backup created immediately after the database server was verified as consistent. For information about consistency checking, see your *HCL® Informix® Administrator's Guide*.

Online and quiescent backups

You can create a backup while the database server is *online* or in *quiescent* mode. The terminal you use to initiate the backup command is dedicated to the backup (displaying messages) until the backup completes. Once you start a backup, the database server must remain in the same mode until the backup finishes; changing the mode terminates the backup activity.

Online backup

You can use an online backup when you want your database server accessible while you create the backup.

Some minor inconveniences can occur during online backups. An online backup can slow checkpoint activity, and that can contribute to a loss in performance. However, this decline in performance is far less costly than the time that you lose when users were denied access to the database server during a backup.

During an online backup, allocation of some disk pages in storage spaces can temporarily freeze. Disk-page allocation is blocked for one chunk at a time until you back up the used pages in the chunk.

Quiescent backup

You create a quiescent backup while the database server is quiescent. Use quiescent backups when you want to eliminate partial transactions in a backup.

Do not use quiescent backups when users need continuous access to the databases.

Back up to tape

When you back up to tape, you must ensure that an operator is available and that you have sufficient media.

Keep an operator available during a backup to mount tapes as prompted. A backup could take several reels of tape. When an operator is not available to mount a new tape when one becomes full, the backup waits. During this wait, when the backup is an online backup, the physical log space could fill up, and that causes the database server to stop the backup. Thus, make sure that an operator is available.

After a tape fills, the *ontape* utility rewinds the tape, displays the tape number for labeling, and prompts the operator to mount the next tape when you need another one. Follow the prompts for labeling and mounting new tapes. A message informs you when the backup is complete.

Label tapes created with ontape

When you label tapes created with the ontape utility, the label must include the following information:

- Backup level
- Date and time
- Tape number that ontape provides

The following example shows what a label can look like:

```
Level 1: Wed Nov 27, 2001 20:45 Tape # 3 of 5
```

Each backup begins with its first tape reel numbered 1. You number each additional tape reel consecutively thereafter. You number a five-tape backup 1 through 5. (Of course, it is possible that you could not know that it is a five-tape backup until it is finished.)

Back up to standard output

A backup to standard output creates an archive in the memory buffer provided by the operating system. If you choose to back up to standard output, you do not need to provide tapes or other storage media.

Backing up to standard output has the following advantages:

- There are no expensive write and read operations to disk or tape.
- You can use operating system utilities to compress or otherwise process the data.
- You can use the archive to create a duplicate of the server by immediately restoring the data onto another database server.

If you back up to standard output, you must also restore from standard input.

When ontape performs a backup to standard output, the data is written to an output file. The directory of the output must have enough disk space to hold the backed-up data. You can use operating system utilities to compress the data. In addition, the user executing the backup command must have write permission to the file to which the backup is diverted or permission to create the file.

When you back up to standard output, ontape does not prompt for user interaction. Error and information messages are written to stderr instead of being directed to standard output.

The TAPESIZE configuration parameter is not used because the capacity of standard output is assumed to be unlimited. The TAPEBLK configuration parameter, however, is valid because it defines the size of the transport buffer between the backend server and the ontape client. You can optimize throughput by setting TAPEBLK to an appropriate value.

You can simultaneously back up and restore a database server to clone it or set up High-Availability Data Replication. For more information, see [Simultaneous backup and restore by using standard I/O on page 173](#).

Related reference

[Restore from standard input on page 171](#)

Back up to a directory

If you choose to back up to a directory, you do not need to provide tapes. Instead, you back up the data to a directory of a local file system or a directory that has been mounted on the local system.

The person who runs the backup must have write permission to the directory. The directory must have enough disk space to hold the backed-up data. You can use operating system utilities to compress the data after it is backed up.

Backing up to a directory has the following advantages:

- Multiple instances can simultaneously back up to the same directory file system.
- You can use operating system utilities to compress or otherwise process the data.
- You can easily configure your system to automatically back up a log file when the file is full.

Set the file directory path

Use the TAPEDEV configuration parameter to specify the absolute path name on a directory of a file system to use for the storage-space archive file. This is the destination where ontape writes storage space data during an archive and the source from which ontape reads data during a restore. You specify the directory where the logical log backup files are written with the LTAPEDEV configuration parameter.



Tip: When you back up to a directory file system, specify the `-d` option to turn off ontape interactive prompts.

Rename existing files

When ontape repeats an archive operation, it renames the existing files so that old files are not rewritten. A timestamp is added to the file name, which provides a way for related storage space or logical log files to be organized together.

Renaming conventions:

- Storage-space archive files

The archive checkpoint time is added, and has the format `servername_YYYYMMDD_hhmmss_archive-level`.

- Logical log backup files

The backup time is added, and has the format `servername_YYYYMMDD_hhmmss`.

For example, the file `My_instance_L0` is renamed to `My_instance_20080913_091527_L0`

When restoring from a file system directory, `ontape` requires that storage-space archive and logical-log backup files be named as specified by the `TAPEDEV` and `LTAPEDEV` parameters. If files have been renamed, including by `ontape` because of repeated archives and backups, files must be manually renamed to their original file names.

Override the default name of the archive files

You can override the default name of the archive files. When `TAPEDEV` or `LTAPEDEV` is a directory path, the default permanent file name consists of `hostname_servernum_Ln` (for levels), and `hostname_servernum_Lognnnnnnnnnn` (for log files). You can override the prefix part of the permanent file name, `hostname_servernum`, by setting the environment variable **IFX_ONTAPE_FILE_PREFIX**.

For example, if you set **IFX_ONTAPE_FILE_PREFIX** to "My_Instance", then during archive, the files are named `My_Instance_L0`, `My_Instance_L1`, `My_Instance_L2`, and, `My_Instance_Log0000000001`, `My_Instance_Log0000000002`, and so on. During restore, `ontape` searches for files in the `TAPEDEV` directory with file names like `My_Instance_L0`, and searches for files in the `LTAPEDEV` directory with file names like `My_Instance_Log0000000001`.

ontape utility syntax: Perform a backup

Use `ontape` utility command options to back up to tape.

Pre-requisites: Before you begin a backup, perform the following steps:

- If necessary, place a write-enabled tape on the tape-drive device that `TAPEDEV` specifies.
- If you set `TAPEDEV` to `STDIO`, ensure that there is enough memory for the backup data.
- If you are using a tape device, put the device online with the appropriate operating-system command.
- Place the database server in online or quiescent mode.
- Log in as the owner of the database server: user **informix** or **root** for a standard installation, or the owner of the non-root installation.

If you are using TAPE devices, do not store more than one backup on the same tape. The tape devices must be of the rewindable type. Begin every backup with a different tape. (Often, a backup spans more than one tape.)

To create a backup, use the `-s` option of the `ontape` command.

Syntax

Create a backup

```
ontape [ < -FILE option > (explicit id) ] [ -v ] { -s { [ -L { 0 | 1 | 2 } ] | -L 0 { -B | -U | -A | -N } database } [ -F ] [ -t { tape_device | STDIO } ] ] | -N database } [ -a ]
```

Element	Purpose	Key considerations
-A	Directs ontape to change the status of the specified database to ANSI-compliant logging.	A database that has ANSI-compliant logging cannot be changed to a different logging mode. A level-0 backup is required to make this logging mode change.
-B	Directs ontape to change the status of the specified database to buffered logging.	A level-0 backup is required to make this logging mode change.
-d	Directs ontape to proceed without interactive prompts.	You can turn off the prompts if you are backing up to or restoring from a directory of a file system. This option does not apply to tape devices, which must pause the backup while you change tapes.
<i>database</i>	The name of the database to change the logging mode.	The database name cannot include a database server name. More than one database name can be specified in the same command.
-F	Directs ontape to perform a fake backup.	<p>A fake backup is only applicable during a backup to standard output.</p> <p>A fake backup is useful for cloning the data in a server. For example, to populate the secondary server in a high-availability cluster.</p> <p>To avoid compromising the normal backup activities, do not keep a record of a fake backup</p> <p>Alternatively, you can use the SQL administration API equivalent: ARCHIVE FAKE. See <i>HCL® Informix® Administrator's Reference</i> for more information.</p>
-L	Directs ontape to create a backup of the level specified.	<p>If you are backing up to tape, use the -L option to specify the backup level as part of the command, you can avoid being prompted for it.</p> <p>If you are backing up to standard output, and do not specify a backup level, ontape performs and level-0 backup.</p>
-N	Directs ontape to end logging for the specified database.	A backup is optional with this logging mode change.
-s	Directs ontape to create a backup.	ontape prompts you to supply the backup level (0, 1, or 2) that you want to create if you do not supply a value using the -L option.

Element	Purpose	Key considerations
-t	Directs ontape to use a different tape device for the current backup or restore.	The -t option overrides the value of the TAPEDEV configuration parameter for the current backup or restore. The -t STDIO option directs ontape to back up to standard output or restore from standard input.
<i>tape_device</i>	The name of the tape device on which to store the backup.	
-U	Directs ontape to change the status of the specified database to unbuffered logging.	A level-0 backup is required to make this logging mode change.
-v	Directs ontape to write informational message to stderr during a backup to standard output.	Verbose mode is useful for monitoring the progress of a backup to standard output.

The ontape utility backs up the storage spaces in the following order: root dbspaces, blobspaces, sbspaces, and dbspaces.

Related reference

[Change database logging status on page 140](#)

Backup examples

Execute the following command to start a backup to tape without specifying a level: `ontape -s`

You can use the **-L** option to specify the level of the backup as part of the command, as the following example shows:

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

Use the **-d** option to avoid interactive prompts when you are backing up to or restoring from a directory: `ontape -s -L 0 -d`

When you do not specify the backup level on the command line, ontape prompts you to enter it. The following figure illustrates a simple ontape backup session.

Figure 24. Example of a simple backup created with ontape

```

ontape -s
Please enter the level of archive to be performed (0, 1, or 2) 0

Please mount tape 1 on /dev/rst0 and press Return to continue ...
16:23:13 Checkpoint Completed: duration was 2 seconds
16:23:13 Level 0 Archive started on rootdbs
16:23:30 Archive on rootdbs Completed.
16:23:31 Checkpoint Completed: duration was 0 seconds

Please label this tape as number 1 in the arc tape sequence.
This tape contains the following logical logs:

3

Program over.

```

The following example shows how to create a level-0 archive of all storage spaces to standard output, which is diverted to a file named `level_0_archive` in the directory `/home`:

```
ontape -s -L 0 >/home/level_0_archive -t STDIO
```

The following example assumes `TAPEDEV STDIO` in `onconfig` and creates a level-1 archive to standard output, which is diverted to a pipe:

```
ontape -v -s -L 1|compress -c >/home/compressed/level_1_archive
```

The **compress** system utility reads from the pipe as input, compresses the data, and writes the data to the file `level_1_archive` in the `/home/compressed` directory. The `ontape` information messages are sent to `stderr`.

Back up raw tables

You can use `ontape` to back up a raw table, however, raw tables are not logged. Therefore, when you restore a raw table, any changes that occurred since the last backup cannot be restored. It is recommended that you use raw tables only for initial loading of data and then alter raw tables to standard tables before performing transactions. For more information, see the *HCL® Informix® Administrator's Guide*.

Back up to Amazon Simple Storage Service

You can use the `ontape` utility to back up and restore data to or from the Amazon Simple Storage Service (S3). You are responsible for terms and any charges associated with your use of the Amazon Simple Storage Service.

Before you begin

Prerequisites:

- You must have an Amazon account to perform cloud storage backups. See the Amazon website for instructions about setting up an account.
- Java™ version 1.5 or later is required.
- Backup objects must be 5 GB or smaller.

About this task

The following steps show how to back up data to the Amazon Simple Storage Service (S3) System and restore from it by using `ontape` backup and restore utility. In this context, cloud storage refers to an online storage service over the Internet. If you choose to back up to cloud storage, you do not need to provide tapes. Instead, you back up the data to a virtual device, most likely located on the Internet.

1. Configure the online storage device.
 - a. Using a web browser, navigate to the Amazon S3 website and log on.
 - b. Obtain an access key ID and a secret access key.
 - c. Store the access credentials in a file. Set the permissions on the file to allow access only to user executing the `ontape` utility.

Result

- On UNIX™ systems, store the values in the file: `$INFORMIXDIR/etc/ifxbkpccloud.credentials`
- On Windows™ systems, store the values in: `%INFORMIXDIR%\etc\ifxbkpccloud.credentials`

The file must have the following format:

```
secretKey=secret_access_key
accessKey=access_key_ID
```

- d. Use the `ifxbkpccloud.jar` utility to create and name a storage device in the region where you intend to store data. Amazon uses the term *bucket* to describe the container for backup data. The storage device name you choose has the same restrictions as those for the bucket name in Amazon S3 and must be unique.

Result

For example, the following command creates a storage device named **mytapedevice** at a **US Standard** region on Amazon S3. Run the command from the `$INFORMIXDIR/bin` directory on UNIX™ systems, or from `%INFORMIXDIR%\bin` on Windows™ systems.

```
java -jar ifxbkpccloud.jar CREATE_DEVICE amazon mytapedevice US_Standard
```

Result

2. Set the `TAPEDEV` and `LTAPEDEV` configuration parameters in the `onconfig` file to point to the cloud storage location. For example:

Result

```
TAPEDEV '/opt/IBM/informix/tapedev_dir, keep = yes, cloud = amazon,
url = https://mytapedevice.s3.amazonaws.com'
LTAPEDEV '/opt/IBM/informix/ltapedev_dir, keep = yes, cloud = amazon,
url = https://mylogdevice.s3.amazonaws.com'
```

3. Back up data to the online storage device by using the `ontape` utility.

Result

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

You can restore data from the cloud storage by using the following command:

```
ontape -r
```

You should use https secure data transmission when transferring data to cloud storage. You should encrypt data before transferring data to a cloud image. To encrypt data, use the `BACKUP_FILTER` and `RESTORE_FILTER` configuration parameters to call an external encryption program. The `archecker` utility does not support table-level restore of data from cloud storage.

Related reference

[Tape and tape device parameters for ontape on page 133](#)

[Set the tape-device parameters on page 134](#)

[Cloud storage file naming conventions on page 151](#)

[The ifxbkpccloud.jar utility on page 151](#)

The ifxbkpccloud.jar utility

Use the `ifxbkpccloud.jar` utility to configure an online storage device for the Amazon Simple Storage Service.

The following options are supported by the `ifxbkpccloud.jar` utility:

- `CREATE_DEVICE provider device [region]`
- `DELETE_DEVICE provider device`
- `LIST_DEVICES provider`
- `DELETE_FILE provider device file`
- `LIST_FILES provider device`

The parameters for the `ifxbkpccloud.jar` commands are defined as follows:

- *provider* is **amazon**.
- *device* is the name of the storage device.
- *region* is one of the following: **US_Standard**, **US_West**, **EU_Ireland** or **AP_Singapore**.
- *file* is the name of backup object (key) stored on Amazon S3.

Error messages from `ifxbkpccloud.jar` are written to `$INFORMIXDIR/ifxbkpccloud.log` on UNIX™ machines and to `%INFORMIXDIR%\ifxbkpccloud.log` on Windows™ machines.

Related information

[Back up to Amazon Simple Storage Service on page 149](#)

Cloud storage file naming conventions

Files associated with cloud storage backups have unique file names.

Data space backup files are saved by using the following format:

```
hostname_servernum_Larchive_level
```

Log backup file names are saved by using the following format:

```
hostname_servernum_lognnnnnnnnnn
```

If the object exists at the cloud storage location, the file is renamed to avoid overwriting old object. Renaming the file adds a timestamp to the object name.

Data space backup files are saved by using the following format:

```
hostname_servernum_YYYYMMDD_hhmmss_Larchive_level
```

Log backup file names are saved by using the following format:

```
hostname_servernum_lognnnnnnnnnn_YYYYMMDD_hhmmss
```

Related information

[Back up to Amazon Simple Storage Service on page 149](#)

When the logical-log files fill during a backup

When the logical log fills during a backup, the console displays a message and the backup suspends normal processing. How you handle the logical-log filling depends on whether you can use one or two tape devices.

When you can use two tape devices

When you can use two tape devices with the database server, log in as the owner of the database server: user **informix** or **root** for a standard installation, or the owner of the non-root installation.

Verify that LTAPEDEV and TAPEDEV specify different path names that correspond to separate tape devices. When they do, back up the logical-log files. See [Create a backup on page 141](#).

When LTAPEDEV and TAPEDEV are identical, assign a different value to the logical-log tape device (LTAPEDEV) and initiate a logical-log-file backup. Otherwise, your options are to either leave normal database server processing suspended until the backup completes or cancel the backup.

When only one tape device is available

When you create a backup with the only available tape device, you cannot back up any logical-log files until you complete the backup. When the logical-log files fill during the backup, normal database server processing halts. You can either stop the backup (by using Ctrl-C only) to free the tape device and back up the logical logs to continue processing, or leave normal processing suspended until the backup completes.

You can take steps to prevent this situation. The section [Start an automatic logical-log backup on page 155](#) describes these steps.

When a backup terminates prematurely

When you cancel or interrupt a backup, sometimes the backup progresses to the point where you can consider it complete. When listed in the monitoring information, as described in [Monitor backup history by using oncheck on page 153](#), you know the backup completed.

Monitor backup history by using oncheck

You can monitor the history of your last full-system backup by using oncheck.

Execute the oncheck -pr command to display reserved-page information for the root dbspace. The last pair of reserved pages contains the following information for the most recent backup:

- Backup level (0, 1, or 2)
- Effective date and time of the backup
- Time stamp describing when the backup began (expressed as a decimal)
- ID number of the logical log that was current when the backup began
- Physical location in the logical log of the checkpoint record (that was written when the backup began)

The effective date and time of the backup equals the date and time of the checkpoint that this backup took as its starting point. This date and time could differ markedly from the time when the backup process was started.

For example, when no one accessed the database server after Tuesday at 7 P.M., and you create a backup Wednesday morning, the effective date and time for that backup is Tuesday night, the time of the last checkpoint. In other words, when there has been no activity after the last checkpoint, the database server does not perform another checkpoint at the start of the backup.

Back up logical-log files with ontape

You must only use ontape to back up logical-log files when you use ontape to make your backup tapes.

In addition to backing up logical-log files, you can use ontape to switch to the next log file, move logical-log files to other dbspaces, or change the size of the logical log. Instructions for those tasks appear in your *HCL® Informix® Administrator's Guide*.

Before you back up the logical-log files

Before you back up the logical-log files, you need to understand the following issues:

- Whether you need to back up the logical-log files
- When you need to back up the logical-log files
- Whether you want to perform an automatic or continuous backup

For more information about these issues, see [Logical-log backup on page 5](#).

Use blobspace TEXT and BYTE data types and logical-log files

You must keep in mind the following two points when you use TEXT and BYTE data types in a database that uses transaction logging:

- To ensure timely reuse of blobpages, back up logical-log files. When users delete TEXT or BYTE values in blobspaces, the blobpages do not become freed for reuse until you free the log file that contains the delete records. To free the log file, you must back it up.
- When you must back up an unavailable blobspace, ontape skips it and makes it impossible to recover the TEXT or BYTE values when it becomes necessary. (However, blobpages from deleted TEXT or BYTE values do become free when the blobspace becomes available even though the TEXT or BYTE values were not backed up.)

In addition, regardless of whether the database uses transaction logging, when you create a blobspace or add a chunk to a blobspace, the blobspace or new chunk is not available for use until the logical-log file that records the event is not the current logical-log file. For information about switching logical-log files, see your *HCL® Informix® Administrator's Guide*.

Use /dev/null when you do not need to recover

When you decide that you do not need to recover transactions or administrative database activities between backups, you can set the database server configuration parameter LTAPEDEV to `/dev/null`.



Important: When you set LTAPEDEV to `/dev/null`, it has the following implications:

- You can only restore the data that your database server manages up to the point of your most recent backup and any previously backed-up logical-log files.
- When you perform a recovery, you must always perform a full-system restore. (See [Full-system restore on page 157](#).) You cannot perform partial restores or restore when the database server is online.

When you set LTAPEDEV to `/dev/null`, the database server marks a logical-log file as backed up (status B) as soon as it becomes full. The database server can then reuse that logical-log file without waiting for you to back it up. As a result, the database server does not preserve any logical-log records.

Fast recovery and rolling back transactions are not impaired when you use `/dev/null` as your log-file backup device. For a description of fast recovery, see your *HCL® Informix® Administrator's Guide*. For information about rolling back transactions, see the ROLLBACK WORK statement in the *HCL® Informix® Guide to SQL: Syntax*.

When to back up logical-log files

You must attempt to back up each logical-log file as soon as it fills. You can tell when you can back up a logical-log file because it has a *used* status. For more information about monitoring the status of logical-log files, see your *HCL® Informix® Administrator's Guide*.

Start an automatic logical-log backup

The database server can operate online when you back up logical-log files. To back up all full logical-log files, use the **-a** option of the `ontape` command.

Request a logical-log backup

```
ontape -a
```

The **-a** option backs up all full logical-log files and prompts you with an option to switch the logical-log files and back up the formerly current log.

When the tape mounted on `LTAPEDEV` becomes full before the end of the logical-log file, `ontape` prompts you to mount a new tape.

When you press the Interrupt key while a backup occurs, the database server finishes the backup and then returns control to you. Any other full logical-log files receive a used status.

To back up all full logical-log files, execute the `ontape -a` command.

Starting a continuous logical-log file backup

When you do not want to monitor the logical-log files and start backups when the logical-log files become full, you can start a continuous backup.

About this task

When you start a continuous backup, the database server automatically backs up each logical-log file as it becomes full. When you perform continuous logical-log file backups, the database server protects you against ever losing more than a partial logical-log file, even in the worst case media failure when a chunk that contains logical-log files fails.

To start a continuous backup of the logical-log files, use the `ontape -c` command. The **-c** option initiates continuous backup of logical-log files. The database server backs up each logical-log file as it becomes full. Continuous backup does not back up the current logical-log file. The database server can operate in online mode when you start continuous backups.

Whether the logical-log files are backed up to tapes or a directory depends on the setting of the `LTAPEDEV` configuration parameter:

- If the `LTAPEDEV` configuration parameter is set to a tape device, someone must always make media available for the backup process. When the specified mounted tape becomes full before the end of the logical-log file, the database server prompts the operator for a new tape. Also, you must dedicate the backup device to the backup process.
- If the `LTAPEDEV` configuration parameter is set to a directory, logical-log files can be backed up unattended. Logical logs are backed up as they fill and a new file is created in the directory for each logical log. Backup is limited by space available for new files.

To back up to a directory, as an alternative to using the `ontape -c` command, you can call the `ontape -a -d` automatic logical log backup command from a script specified by the `ALARMPROGRAM` configuration parameter. You can use either the `alarmprogram` or `script` or the `log_full` script, both of which are found in the `$INFORMIXDIR/etc` directory.

To use the `alarmprogram` script to back up logical logs to a directory:

1. Set the `LTAPEDEV` parameter to an existing directory.

Make sure that this directory is owned by **informix** and group **informix**.

2. Edit the `ALARMPROGRAM` script (`$INFORMIXDIR/etc/alarmprogram.sh` on UNIX™ or Linux™ or `%INFORMIXDIR%\etc\alarmprogram.bat` on Windows™), as follows:
 - a. Set the `BACKUPLOGS` parameter within the file to `y`.
 - b. Change the backup program from `onbar -b -l` to `ontape -a -d`.
3. Restart the database server.

End a continuous logical-log backup

To end continuous logical-log backup, press the Interrupt key (CTRL-C).

When you press the Interrupt key while the database server backs up a logical-log file to a local device, all logs that were backed up before the interrupt are captured on the tape and are marked as backed up by the database server.

When you press the Interrupt key while the database server waits for a logical-log file to fill (and thus is not backing up any logical-log files), all logs that were backed up before the interrupt reside on the tape and are marked as backed up by the database server.

When you press the Interrupt key while the database server performs a continuous backup to a remote device, any logical-log files that were backed up during this operation can or cannot reside on the tape, and are not marked as backed up by the database server (a good reason why you should not do continuous remote backups).

After you stop continuous logging, you must start a new tape for subsequent log backup operations.

You must explicitly request logical-log backups (by using `ontape -a`) until you restart continuous logging.

Devices that logical-log backups must use

The `ontape` utility uses parameters defined in the `onconfig` file to define the tape device for logical-log backups. However, consider the following issues when you choose a logical-log backup device:

- When the logical-log device differs from the backup device, you can plan your backups without considering the competing needs of the backup schedule.
- When you specify `/dev/null` as the logical-log backup device in the configuration parameter `LTAPEDEV`, you avoid having to mount and maintain backup tapes. However, you can only recover data up to the point of your most recent backup tape. You cannot restore work done after the backup. See the warning about setting `LTAPEDEV` to `/dev/null` in [Use /dev/null when you do not need to recover on page 154](#).

If the log backup device on any server node in a high-availability cluster is set to `/dev/null` (on Linux™ or UNIX™) or NUL (on Windows™), then the backup device for all of the other servers within the cluster (including the primary server and any HDR, RSS or SDS secondary servers) must also be set to `/dev/null` (or NUL).

- When your tape device runs slow, the logical log could fill up faster than you can copy it to tape. In this case, you could consider performing the backup to disk and then copying the disk backup to tape.

Restore with ontape

These topics provide instructions for restoring data with the ontape utility for the following procedures:

- A whole-system restore
- A restore of selected dbspaces, blobspaces, and sbspaces

Before you start restoring data, you must understand the concepts in [Restore systems on page 7](#). As explained in that section, a complete recovery of database server data generally consists of a physical restore and a logical restore.

Types of physical restore

If a failure causes the database server to go offline, you must restore all the database server data. This type of restore is a *full-system* restore. You can only restore data to the same version of HCL Informix®. When the failure did not cause the database server to go offline, you can restore only the storage spaces that failed. For illustrations of the restore types, see [Warm, cold, and mixed restores on page 8](#).

Full-system restore

When your database server goes offline because of a disk failure or corrupted data, it means that a critical dbspace was damaged. The following list shows critical dbspaces:

- The root dbspace
- The dbspace that contains the physical log
- A dbspace that contains logical-log files

When you need to restore any critical dbspace, you must perform a full system restore to restore all the data that your database server manages. You must start a full-system restore with a cold restore. See [Cold, warm, or mixed restores on page 158](#).

Restores of dbspaces, blobspaces, and sbspaces

When your database server does not go offline because of a disk failure or corrupted data, the damage occurred to a noncritical dbspace, blobspace, or sbspace.

When you do not need to restore a critical dbspace, you can restore only those storage spaces that contain a damaged chunk or chunks. When a media failure occurs in one chunk of a storage space that spans multiple chunks, all active transactions for that storage space must terminate before the database server can restore it. You can start a restore

operation before the database server finishes the transactions, but the restore becomes delayed until the database server verifies that you finished all transactions that were active at the time of the failure.

Cold, warm, or mixed restores

When you restore the database server data, you must decide whether you can do it while the database server is offline or online. This decision depends in part on the data that you intend to restore.

Cold restores

Perform a *cold restore* while the database server is offline. It consists of both a physical restore and a logical restore. You must perform a cold restore to restore any critical dbspaces.

The database server is offline when you begin a cold restore but it goes into recovery mode after it restores the reserved pages. From that point on it stays in recovery mode until either a logical restore finishes (after which it works in quiescent mode) or you use the onmode utility to shift it to another mode.

You can rename chunks by specifying new chunks paths and offsets during a cold restore. This option is useful if you need to restore storage spaces to a different disk from the one on which the backup was made. You can rename any type of chunk, including critical chunks and mirror chunks. For more information, see [Rename chunks during a restore on page 168](#). You can also rename chunks for an external cold restore; see [Rename chunks on page 178](#) for more information.

A cold restore can be performed after a dbspace has been renamed and a level-0 backup or a backup of the rootdbs and renamed dbspace is performed.

Warm restores

A *warm restore* restores noncritical storage spaces while the database server is in online or quiescent mode. It consists of one or more physical restore operations (when you restore multiple storage spaces concurrently), a logical-log backup, and a logical restore.

During a warm restore, the database server replays backed-up logical-log files for the storage spaces that you restore. To avoid overwriting the current logical log, the database server writes the logical-log files that you designate for replay to temporary space. Therefore, a warm restore requires enough temporary space to hold the logical log or the number of log files being replayed, whichever is smaller. For information about how the database server looks for temporary space, see the discussion of DBSPACETEMP in the *HCL® Informix® Administrator's Guide*.



Important: Make sure that enough temporary space exists for the logical-log portion of the warm restore; the maximum amount of temporary space that the database server needs equals the size of all the logical-log files.

A warm restore can be performed after a dbspace has been renamed and a level-0 archive of the rootdbs and renamed dbspace is taken.

Mixed restores

A *mixed restore* is a cold restore followed by a warm restore. A mixed restore restores some storage spaces during a cold restore (the database server is offline) and some storage spaces during a warm restore (the database server is online). You could do a mixed restore when you perform a full-system restore, but you need to provide access to a particular table or set of tables as soon as possible. In this case, perform a cold restore to restore the critical dbspaces and the dbspaces that contain the important tables.

A cold restore takes less total time to restore all your data than a mixed restore, even though the database server is online during part of a mixed restore because a mixed restore requires two logical restores (one for the cold restore and one for the warm restore). A mixed restore, however, requires the database server to go offline for less time than a cold restore.

The dbspaces not restored during the cold restore do not become available until after the database server restores them during a warm restore, even though a critical dbspace possibly did not damage them.

ontape utility syntax: Perform a restore

Use the `-r` option to perform a full physical and logical restore of the database server data with `ontape`. Use the `-D` option to restore selected storage spaces. Use the `-rename` option to rename chunks during the restore.

You must run the `ontape` command as the owner of the database server: user **informix** or **root** for a standard installation, or the owner of the non-root installation. The owner of the database server for the restore must be the same as the owner of the database server for the backup.

Run a full or physical restore

```
>>-ontape-+-----+-----+-----+-----+-----+-----+-----+----->
      |               (1) | '- -p-+-----+-'
      '-| -FILE option |-----'          '- -e-'


>+-----+-----+-----+-----+-----+-----+-----+----->
+- -encrypt+          '- -pw-+-----+---'
 '- -decrypt-'              '--filename--'

>+-----+-----+-----+-----+-----+-----+-----+----->
|               .------.------.------.------.------.------.------.
|               V
 '- -rename-+----- -p--old_path-- -o--old_offset-- -n--new_path-- -o--new_offset-+---'
           '- -f--filename-----'

>+-----+-----+-----+-----+-----+-----+-----+-----><
|               .------.------.------.------.------.------.------.
|               V               | |               '- -v-' '- -d-'
 '- -D-----dbspace-+-'
```

Run a logical restore

```
ontape {-l {-C | -x} | -s }
```

Element	Purpose	Key considerations
-C	Restores logs from the current logical log tape without sending prompts to mount the tape.	The server is placed in suspend log restore state, and the command exits after the last applicable log is restored. The server sends a prompt if a log spans tapes.
-D	Directs ontape to restore only the storage spaces you specify.	<p>The database server must go into online or quiescent mode to do a warm restore. When you use the -D option, you can restore selected storage spaces.</p> <p>When you do not specify the -D option, ontape performs a full-system restore. The database server must go offline to do a full-system restore. For more information, see Restore selected storage spaces on page 166.</p>
-d	Directs ontape to perform a restore from backup, even when the server version where backup is created, is different than the server version where restore is being performed.	<p> Note: Use when the only available backup was created from a previous version of the server.</p> <p>Backup server version and restore server version must only have minor fix-pack version difference.</p> <p>Compatibility of backups from different versions is not guaranteed.</p>
<i>dbspace</i>	Is the name of a storage space to restore.	You can specify multiple storage spaces, but you must include the root dbspace.
-decrypt	Specifies to decrypt any encrypted storage spaces during the physical restore of the spaces.	For more information about storage space encryption, see Changing storage space encryption during a restore on page .
-e	Directs ontape to perform an external restore	<p>For more information, see Perform an external backup and restore on page 173.</p> <p>This option is compatible with renaming chunks for external cold restores.</p>
-encrypt	Specifies to encrypt storage spaces during the physical restore of the spaces.	<p>Storage space encryption must be enabled by the DISK_ENCRYPTION configuration parameter. Otherwise, storage spaces are not encrypted during the restore.</p> <p>For more information about storage space encryption, see Changing storage space encryption during a restore on page .</p>

Element	Purpose	Key considerations
-f <i>filename</i>	Specifies a file containing the names and offsets of chunks to be renamed and their new locations. Use to rename many chunks at one time.	<p>The file name can be any valid UNIX™ or Windows™ file name, including simple (<i>listfile_1</i>), relative (<i>./backup_lists/listfile_2</i> or <i>..\backup_lists\listfile2</i>), and absolute (<i>/usr/informix/backup_lists/listfile3</i> or <i>c:\informix\backup_lists\listfile3</i>) file names.</p> <p>In the file, list the old chunk path name and offset and the new chunk path name and offset, with a blank space or a tab between each item. Put information for each chunk on a separate line. Blank lines are ignored. Begin comment lines with a # symbol.</p>
-l	Directs ontape to perform a logical restore.	The -l option restores data from the logical-log backup tapes you created after (and including) your last level-0 backup.
-p	Directs ontape to perform a physical data restore.	The -p option restores data from the backup tape you created after (and including) your last level-0 backup. During the restore, the database server is in single-user mode.
-p <i>old_path</i> -o <i>old_offset-n</i> <i>new_path</i> -o <i>new_offset</i>	Specifies the chunk to be renamed and its new location. Use to rename one or more chunks at one time.	<p>The variables for this element are:</p> <p><i>old_path</i> The current path and file name of the chunk.</p> <p><i>old_offset</i> The current offset of the chunk, in kilobytes.</p> <p><i>new_path</i> The new path and file name of the chunk.</p> <p><i>new_offset</i> The new offset of the chunk.</p>
-r	Directs ontape to perform a data restore (both physical and logical).	The -r option restores data from the backup tape and the logical-log backup tapes you created after (and including) your last level-0 backup.
-rename	Directs ontape to rename the specified chunks.	For more information about renaming chunks during a restore, see Rename chunks during a restore on page 168 .
-S	Directs ontape to perform a logical log salvage.	If you want to salvage logical logs, you must use the -S option before performing a restore from standard input. The LTAPEDEV configuration parameter must be set to the logical log tape device.

Element	Purpose	Key considerations
-t STDIO	Directs ontape to restore from standard input.	The -t option overrides the value of the TAPEDEV configuration parameter for the current restore.
-v	Directs ontape to write informational message to stderr during a restore from standard input.	Verbose mode is useful for monitoring the progress of a restore from standard input.
-X	Quiesces a server in logical restore suspend state without restoring additional logs.	Include this option with -r -l to end continuous log restore of logical logs.



Note: The -pw option is required only when the [Storage space encryption](#) feature is enabled and no stash file is in use. Supply an optional path to a file containing the keystore password, otherwise ontape will prompt for a password before performing the restore.

Restore the whole system

This section outlines the prerequisites and steps you need to complete to restore your entire database server with ontape.

The following list summarizes the main steps in a full-system restore:

1. Gather the appropriate backup and logical log tapes.
2. Decide on a complete cold or a mixed restore.
3. Verify your database server configuration.
4. Perform a cold restore.

Familiarize yourself with these instructions before you attempt a full-system restore.

Gather backup and logical-log tapes before restoring

Before you restore an entire database system, you must gather backup and logical log tapes. If you changed the names of backup and logical log files, you must also manually rename the files to their original file names.

Backup tapes

Gather all the tapes from your latest level-0 backup that contain the storage spaces you are restoring and any subsequent level-1 or level-2 backups.

Identify the tape that has the latest level-0 backup of the root dbspace on it; you must use this tape first.

Logical-log tapes

If at the time of the archive checkpoint, an open transaction started, gather all logical-log tapes before you perform the level-0 backup.

Gather all the logical-log tapes from the backup after the latest level-0 backup of the storage spaces you are restoring.

When using `ontape` to create an archive backup of the system, a snapshot of the logical logs is included with the archive. At the end of the archive, the system displays a message that indicates what logical logs are included in the archive. The snapshot is included in the archive so that if any open transactions exist at the time of the backup, those transactions can be reconciled when the archive is restored. Then:

- If you decide not to replay any logical logs, the system can be brought to a consistent state.
- If you decide to replay logical logs, the logs contained within the archive backup are discarded and you must replay transactions from the logical log backups.

The starting log file is the oldest log file containing an open transaction at the time of the last restored archive. You can identify that log file from the message that was displayed when the last archive was restored.

Example:

- The `ontape -s -L 0` command performs a level-0 backup of the system and displays a message that states that the archive contains logs 2-4.
- The `ontape -s -L 1` command performs a level 1 incremental backup of the system and displays a message that states that the archive contains logs 8-9.

If you restore only the level-0 archive and want to replay logs, you need log backups starting with log 2. If you restore the level-0 and level-1 archives and want to replay logs, you need log backups starting with log 8.

The restore of the logical log files uses an archive format, not a log file format. However, the logs contained within the restored archive are in log file format, not an archive format.

File names when restoring from directory

When restoring from a file system directory, `ontape` requires that storage-space archive and logical-log backup files be named as specified by the `TAPEDEV` and `LTAPEDEV` configuration parameters. If files were renamed, including by `ontape` because of repeated archives and backups, you must manually rename the files to their original file names. To learn about the naming conventions for storage-space archive files and logical-log backup files, see [Back up to a directory on page 145](#) for the naming conventions of these files.

Decide on a complete cold or a mixed restore

As mentioned in [Cold, warm, or mixed restores on page 158](#), when you restore your entire database server, you can restore the critical dbspaces (and any other storage spaces you want to come online quickly) during a cold restore, and then restore the remaining storage spaces during a warm restore. Decide before you start the restore if you want a cold restore or a mixed restore.

Verify your database server configuration

During a cold restore, you cannot set up shared memory, add chunks, or change tape devices. Thus, when you begin the restore, the current database server configuration must remain compatible with, and accommodate, all parameter values assigned after the time of the most recent backup.

For guidance, use the copies of the configuration file that you create at the time of each backup. However, do not set all current parameters to the same values as were recorded at the last backup. Pay attention to the following three groups of parameters:

- Shared-memory parameters
- Mirroring parameters
- Device parameters

Set shared-memory parameters to maximum assigned value

Make sure that you set your current shared-memory parameters to the maximum value assigned after the level-0 backup. For example, if you decrease the value of `USERTHREADS` from 45 to 30 sometime after the level-0 backup, you must begin the restore with `USERTHREADS` set at 45, and not at 30, even though the configuration file copy for the last backup could register the value of `USERTHREADS` set at 30. (When you do not possess a record of the maximum value of `USERTHREADS` after the level-0 backup, set the value as high as you think necessary. You could reassign values to `BUFFERPOOL`, `LOCKS`, and `TBLSPACES` as well because the minimum values for these three parameters are based on the value of `USERTHREADS`.)

Set mirroring configuration to level-0 backup state

Verify that your current mirroring configuration matches the configuration that was in effect at the time of the last level-0 backup. Because it is recommended that you create a level-0 backup after each change in your mirroring configuration, this creates no problems. The most critical parameters are the mirroring parameters that appear in the configuration file, `MIRRORPATH` and `MIRROROFFSET`.

Verify that the raw devices or files are available

Verify that the raw devices or files that you used for storage (of the storage spaces being restored) after the level-0 backup are available.

For example, if you drop a dbspace or mirroring for a dbspace after your level-0 backup, you must make the dbspace or mirror chunk device available to the database server when you begin the restore. When the database server attempts to write to the chunk and cannot find it, the restore does not complete. Similarly, if you add a chunk after your last backup, you must make the chunk device available to the database server when it begins to roll forward the logical logs.

Perform a cold restore

To perform a cold restore with ontape, the database server must be offline.

You must run the `ontape` command as the owner of the database server: user **informix** or **root** for a standard installation, or the owner of the non-root installation. The owner of the database server for the restore must be the same as the owner of the database server for the backup.

Run the following `ontape` command to restore all the storage spaces: `ontape -r`

When you perform a mixed restore, you restore only some of the storage spaces during the cold restore. You must restore at least all the critical dbspaces, as the following example shows:

```
ontape -r -D rootdbs llogdbs plogdbs
```

Salvage logical-log files

Before the cold restore starts, the console prompts you to salvage the logical-log files on disk. To salvage the logical-log files, use a new tape. It saves log records that you did not back up and enables you to recover your database server data up to the point of the failure.

The following example shows a log salvage:

```
...
Continue restore? (y/n) y
Do you want to back up the logs? (y/n) y

Please mount tape 1 on /dev/ltapedev and press Return to continue.
Would you like to back up any of logs 31 - 32? (y/n) y
Logical logs 31 - 32 may be backed up.
Enter the id of the oldest log that you would like to backup? 31

Please label this tape as number 1 in the log tape sequence.

This tape contains the following logical logs:
    31-32
Log salvage is complete, continuing restore of archive.
Restore a level 1 archive (y/N) y
Ready for level 1 tape
...
```

Mount tapes during the restore

During the cold restore, `ontape` prompts you to mount tapes with the appropriate backup files.

When restoring from a directory, the prompt specifies the absolute path name of the directory. Before responding to the prompt, you can copy or rename the file in the directory.

You can avoid the prompt by using the `ontape -d` option. When using this option, ensure that storage-space archive and logical-log backup files exist in the directory, as specified by the `TAPEDEV` and `LTAPEDEV` parameters. The `ontape` utility scans the directories for the files and uses them for the restore. After restoring the newest applicable logical-log backup file, `ontape` automatically commits the restore and brings the HCL Informix® instance into quiescent mode.

Restore logical log files

When you perform a mixed restore, you must restore all the logical-log files backed up after the last level-0 backup.

When you perform a full restore, you can choose not to restore logical-log files. When you do not back up your logical-log files or choose not to restore them, you can restore your data only up to the state it was in at the time of your last backup. For more information, see [Back up logical-log files with ontape on page 153](#).

To restore the logical logs, use the `ontape -l` command.

Bring the database server online when the restore is over

At the end of the cold restore, the database server is in quiescent mode. You can bring the database server online and continue processing as usual.

When you restore only some of your storage spaces during the cold restore, you can start a warm restore of the remaining storage spaces after you bring the database server online.

Restore selected storage spaces

These topics outline the steps that you must perform during a restore of selected storage spaces with `ontape` while the database server is in online or quiescent mode (a warm restore). During a warm restore, you do not need to worry about shared-memory parameters as you do for cold restores.

Before you attempt a restore, familiarize yourself with these instructions.

The following list describes the main steps in a warm restore:

Gather the appropriate tapes

Gather the appropriate backup and logical-log tapes.

Backup tapes

Before you start your restore, gather together all the tapes from your latest level-0 backup that contain the storage spaces you are restoring and any subsequent level-1 or level-2 backups.

Logical-log tapes

Gather together all the logical-log tapes from the logical-log backup after the latest level-0 backup of the storage spaces you are restoring.

Ensure that needed device are available

Verify that storage devices and files are available before you begin a restore. For example, when you drop a `dbspace` or mirroring for a `dbspace` after your level-0 backup, you must ensure that the `dbspace` or mirror chunk device is available to the database server when you begin the restore. If the storage device is not available, the database server cannot write to the chunk and the restore fails.

When you add a chunk after your last backup, you must ensure that the chunk device is available to the database server when it rolls forward the logical logs.

Back up logical-log files

Before you start a warm restore (even when you perform the warm restore as part of a mixed restore), you must back up your logical-log files. See [Back up logical-log files with ontape on page 153](#).

After the warm restore, you must roll forward your logical-log files to bring the dbspaces that you are restoring to a state of consistency with the other dbspaces in the system. Failure to roll forward the logical log after restoring a selected dbspace results in the following message from ontape:

```
Partial system restore is incomplete.
```

Perform a warm restore

To perform a warm restore with ontape, the database server must operate in online or quiescent mode.

You must run the ontape command as the owner of the database server: user **informix** or **root** for a standard installation, or the owner of the non-root installation. The owner of the database server for the restore must be the same as the owner of the database server for the backup.

To restore selected storage spaces, run the ontape command, with the options that the following example shows:

```
ontape -r -D dbspace1 dbspace2
```

You cannot restore critical dbspaces during a warm restore; you must restore them as part of a cold restore, described in [Restore the whole system on page 162](#).

During the restore, ontape prompts you to mount tapes with the appropriate backup files.

At the end of the warm restore, the storage spaces that were down go online.

Restore raw tables

When you use ontape to restore a raw table, it contains only data that existed on disk at the time of the backup. Because raw tables are not logged, any changes that occurred since the last backup cannot be restored. For more information, see [Back up raw tables on page 149](#) and the *HCL® Informix® Administrator's Guide*.

Configuring continuous log restore with ontape

Before you begin

Ensure that the version of HCL Informix® is identical on both the primary and secondary systems.

About this task

Use continuous log restore to restart a log restore with newly available logs after all currently available logs have been restored. For more information, see [Continuous log restore on page 10](#).

To configure continuous log restore with ontape:

1. On the primary system, perform a level-0 archive with the `ontape -s -L 0` command.
2. On the secondary system, copy the files or mount the tape (as assigned by LTAPEDEV) and perform a physical restore with the `ontape -p` command.
3. Respond to the following prompts:

```
Continue restore? Y
Do you want to back up the logs? N
Restore a level 1 archive? N
```

After the physical restore completes, the database instance waits in fast recovery mode to restore logical logs.

4. On the primary system, back up logical logs with the `ontape -a` command.
5. On the secondary system, copy the files or mount the tape that contains the backed up logical logs from the primary system. Perform a logical log restore with the `ontape -l -C` command.
6. Repeat steps [4 on page 168](#) and [5 on page 168](#) for all logical logs that are available to back up and restore.
7. If you are doing continuous log restore on a secondary system as an emergency standby, run the following commands to complete restoring logical logs and quiesce the server
 - If logical logs are available to restore, use the `ontape -l` command.
 - After all available logical logs are restored, use the `ontape -l -X` command.

Related information

[Continuous log restore on page 10](#)

Rename chunks during a restore

You can rename chunks during a cold restore with ontape. This option is useful if you need to restore storage spaces to a different disk from the one on which the backup was made. You can rename any type of chunk, including critical chunks and mirror chunks.

The ontape rename chunk restore only works for cold restores.

The critical dbspaces (for example, the rootdbs) must be restored during a cold restore. If you do not specify the list of dbspaces to be restored, then the server restores the critical dbspaces and all the other dbspaces. But if you specify the list of dbspaces to be restored, then the critical dbspaces must be included in the list.

For the syntax of renaming chunks with ontape, see [ontape utility syntax: Perform a restore on page 159](#).



Tip: If you use symbolic links to chunk names, you might not need to rename chunks; you need only edit the symbolic name definitions. For more information, see the *HCL® Informix® Administrator's Guide*.



You can rename chunks during an external cold restore. See [Rename chunks on page 178](#) for more information.

Validation sequence for renaming chunks

During a cold restore, ontape performs the following validations to rename chunks:

- It validates that the old chunk path names and offsets exist in the archive reserved pages.
- It validates that the new chunk path names and offsets do not overlap each other or existing chunks.
- If renaming the primary root or mirror root chunk, it updates the `onconfig` file parameters `ROOTPATH` and `ROOTOFFSET`, or `MIRRORPATH`, and `MIRROROFFSET`. The old version of the `onconfig` file is saved as `$ONCONFIG.localtime`.
- It restores the data from the old chunks to the new chunks (if the new chunks exist).
- It writes the rename information for each chunk to the online log.

If either of the validation steps fails, the renaming process stops and ontape writes an error message to the ontape activity log.



Important:

- Perform a level-0 archive after you rename chunks; otherwise your next restore fails.
- If you add a chunk after performing a level-0 archive, that chunk cannot be renamed during a restore. Also, you cannot safely specify that chunk as a new path in the mapping list.
- Renaming chunks for database servers participating in HDR involves a significant amount of offline time for both database servers. For more information, see the *HCL® Informix® Administrator's Guide*.

New chunk requirements

To rename a chunk, follow these guidelines for new chunks:

- The new chunk does not need to exist

You can install the new chunk later and perform a warm restore of a storage space containing it. If you specify a nonexistent chunk, ontape records the rename information in the chunk reserved pages, but does not restore the data. The renamed (but not restored) chunks have a status of offline, designated by `D`, in the `onstat -d` chunk status command output.

- New chunks must have the proper permissions.

Rename operations fail unless the chunks have the proper permissions. For more information, see the *HCL® Informix® Administrator's Guide*.

Rename chunks with command-line options

To rename the chunks by supplying information about the command line, use this command:

```
ontape -r -rename -p /chunk1 -o 0 -n /chunk1N -o 20000
        -rename -p /chunk2 -o 10000 -n /chunk2N -o 0
```

Perform a level-0 archive after the rename and restore operation is complete.

Rename chunks with a file

To rename the chunks by supplying a file named `listfile`, use the following command: `ontape -r -rename -f listfile`

The contents of the `listfile` file are:

```
/chunk1 0 /chunk1N 20000
/chunk2 10000 /chunk2N 0
```

Perform a level-0 archive after the rename and restore operation is complete.

Rename chunks while specifying other options

To rename the chunks with command-line options while performing a restore of **dbspace1** and **dbspace2** where the **rootdbs** is the rootdbs, use the following command:

```
ontape -r -rename -p /chunk1 -o 0 -n /chunk1N -o 20000
        -rename -p /chunk2 -o 10000 -n /chunk2N -o 0
        -D rootdbs dbspace1 dbspace2
```

Alternatively, to rename the chunks by using file while performing a restore of **dbspace1** and **dbspace2**, use the following command:

```
ontape -r -rename -f listfile -D rootdbs dbspace1 dbspace2
```

Perform a level-0 archive after the rename and restore operation is complete.

Rename a chunk to a nonexistent device

To rename a chunk to a device that does not yet exist, you specify the new path name, but you do not restore its storage spaces until after you install the physical device. This option is useful if you need to rename a chunk and it is convenient to perform a cold restore before you install the new device. When the new chunk device is ready, you can perform a warm restore of a storage space onto it.

You can combine renaming chunks with existing devices and renaming chunks with nonexistent devices in the same rename operation. This example shows how to rename a single chunk to a nonexistent device name.

The following table lists example values for the chunks used in this example.

Storage space	Old chunk path	Old offset	New chunk path	New offset
sbspace1	/chunk3	0	/chunk3N	0

Renaming a chunk to a nonexistent device

About this task

To rename a chunk to a nonexistent device:

1. Rename the chunk: using the following command: `ontape -r -rename -p /chunk3 -o 0 -n /chunk3N -o 0`
2. When the following prompt appears, enter `y` to continue:

```
The chunk /chunk3N does not exist. If you continue, the restore
may fail later for the dbspace which contains this chunk.
Continue without creating this chunk? (y/n)
```

The chunk `/chunk3` is renamed to `/chunk3N`, but the data has not yet been restored to `/chunk3N`.

3. Perform a level-0 archive.
4. Add the physical device for `/chunk3N`.
5. Perform a warm restore of **sbspace1** with the `ontape -r -D sbspace1` command.
6. Perform a level-0 archive.

Restore from standard input

You can perform a restore from standard input, you must first have performed a backup to standard output.

When you perform a restore from standard input, `ontape` does not prompt you for options or information. If `ontape` cannot perform the operation with the information you provided in the restore command, `ontape` exits with an appropriate error.

Restoring from standard input differs from restoring from tapes in the following ways:

- No logical restore or logical log salvage occurs.
 - To perform a logical restore, use the `ontape -l` command after the physical restore.
 - To salvage logical logs, use the `ontape -S` command before the physical restore.
- You are not prompted to confirm the restore. Informational messages about the archive are sent to `stderr`.

If you detect a problem, you can interrupt the restore during the 10 second delay between the completion of the archive information and starting the database server.

Examples

In the following example, `ontape` performs a physical restore from the file `level_0_archive`, which contains the archive previously performed to standard output:

```
cat /home/level_0_archive | ontape -p
```

In the following example, `ontape` performs a restore of a level-0 archive, followed by a restore of a level-1 archive:

```
cat /home/level_0_archive /home/level_1_archive | ontape -r
```

In the following example, `ontape` performs a restore of `sbspace1`:

```
cat/home/level_0_archive | ontape -r -D spspace1 -t STDIO
```

When these restores are completed, the database server is left in single-user mode.

Related reference[Back up to standard output on page 144](#)

Restore data to a remote server

You can restore data to a remote server with the `ontape` utility.

The remote server must have the following characteristics:

- Identical hardware and operating systems
- Identical database server versions and editions
- The same configuration and `ROOTPATH` information, although the server names and numbers can differ.

You can restore data to a remote server with the following command:

```
ontape -s -L 0 -F | rsh remote_server "ontape -p"
```

However, the process might hang after completing successfully. You have three primary options:

- Terminate the remote shell process
- Execute the remote shell from the remote server with the following command:

```
rsh local_server "ontape -s -L 0 -F" | ontape -p
```

- Redirect the standard output (stdout) and standard error (stderr) on the remote server with the following command from the `sh` or `bash` shell:

```
ontape -p >/dev/null 2>&1
```

- You can simplify this redirection by placing it in a shell script, `ontape.sh`, on the remote server. You can issue the following command from the local server:

```
ontape -s -L 0 -F | rsh remote_server /my/path/ontape.sh
```

- The shell script `ontape.sh` contains the following text:

```
#!/bin/sh
#define some environment variables, such as

INFORMIXDIR=/... ; export INFORMIXDIR
INFORMIXSQLHOSTS=/...; export
INFORMIXSQLHOSTS ONCONFIG=/...; export ONCONFIG
INFORMIXSERVER=/...; export INFORMIXSERVER
PATH=/...; export PATH
# invoke ontape with stdout/stderr redirection

ontape -p >/dev/null 2>&1
```

Simultaneous backup and restore by using standard I/O

To clone a database server or quickly set up High-Availability Data Replication (HDR), you can perform a simultaneous backup to standard output and restore from standard input. If you perform the backup and restore solely to duplicate a database server, use the `-F` option to prevent the archive from being saved.

On HDR, the secondary server can restore only level-0 archives.

To use standard I/O to perform the backup and restore, set the **TAPEDEV** configuration parameter to **STDIO**, or you can specify `-t STDIO` from the command line.

For example, if the **TAPEDEV** configuration parameter is set to **STDIO**, the following command loads data into the secondary server on an HDR pair (named **secondary_host**).

```
ontape -s -L 0 -F | rsh secondary_host "ontape -p"
```

In the next example, assume that the **TAPEDEV** configuration parameter is not set. The following command loads data into the secondary server of an HDR pair (named **secondary_host**):

```
ontape -s -L 0 -F -t STDIO | rsh secondary_host "ontape -t STDIO -p"
```

The examples perform a fake level-0 archive of the database server on the local computer, pipe the data to the remote computer by using the `rsh` system utility, and perform a physical restore on the remote computer by reading the data directly from the pipe.



Important: The previous examples require that the **INFORMIXDIR**, **INFORMIXSERVER**, **INFORMIXSQLHOSTS**, and **ONCONFIG** environment variables be set in the default environment for the user on the remote computer on which the command is executed. The user must be **informix** or **root**.

Perform an external backup and restore

These topics discuss performing an external backup and recovering data by restoring it with the `ontape` utility.

Recover data by using an external backup and restore

You can perform an *external backup and restore*, which eliminates the downtime of systems because the backup and restore operations are performed external to the HCL® Informix® system.

The `ontape` utility does not move the data during the backup or physical restore. An external backup allows you to copy disks that contain storage-space chunks without using `ontape`. When disks fail, replace them and use vendor software to restore the data, then use `ontape` for the logical restore. For more information, see [Data that is restored in an external restore on page 176](#).

The following are typical scenarios for external backup and restore:

- Availability with disk mirroring

If you use hardware disk mirroring, you can get your system online faster with external backup and restore than with conventional ontape commands.

- Cloning

You can use external backup and restore to clone an existing production system for testing or migration without disturbing the production system.

Data that is backed up in an external backup

Before you begin an external backup, block the database server. Blocking forces a checkpoint, flushes buffers to disk, and blocks user transactions that involve temporary tables. During the blocking operation, users can access that database server in read-only mode. Then you can physically back up or copy the data to another set of disks or storage media by using operating-system or third-party tools. When you complete the external backup, unblock the database server so that transactions can resume. You should include all the chunk files in each storage space and administrative files, such as `onconfig`, in an external backup.



Important: To make tracking backups easier, it is recommended that you back up all storage spaces in each external backup.

The ontape utility treats an external backup as equivalent to a level-0 backup. You cannot perform an external backup and then use ontape to perform a level-1 backup, or vice versa because ontape does not have any record of the external backup. For more information, see [Performing a cold external restore on page 178](#).

Rules for an external backup

Before you begin an external backup, keep in mind the following rules:

- The database server must be online or quiescent during an external backup.
- Use ontape to back up all logical logs including the current log so that you can restore the logical logs at the end of the external restore.
- Suspend continuous logical-log backups before you block the database server for an external backup. After the external backup is complete, resume the continuous logical-log backup.
- Wait until all ontape backup sessions have completed before you block the database server. If any backup sessions are active, the block command displays an error message.
- Any OLTP work or queries are suspended while the database server is blocked. They resume after the database server is unblocked.
- All critical dbspaces of the database server instance must be backed up together simultaneously within the same command bracket of `onmode -c block ... onmode -c unblock`. Backups of different critical dbspaces done at different times cannot be restored to a consistent system.

! **Important:** Because the external backup is outside the control of ontape, you must track these backups manually. For more information, see [Track an external backup on page 95](#).

Performing an external backup

About this task

The database server must be online or in quiescent mode during an external backup.

To perform an external backup without disk mirroring:

1. To obtain an external backup, block the database server with the `onmode -c block` command.
The system takes a checkpoint and suspends all update transactions. Users can access the database server in read-only mode.
2. To back up the storage spaces and administrative files, use a copy command, such as `cp`, `dd`, or `tar` on UNIX™ or `copy` on Windows™, or a file-backup program.

You must back up all chunks in the storage spaces.

3. To allow normal operations to resume, unblock the database server with the `onmode -c unblock` command.
4. Back up all the logical logs including the current log so that checkpoint information is available for the external restore.

! **Important:** Because external backup is not done through ontape, you must ensure that you have a backup of the current logical log from the time when you execute the `onmode -c block` command. Without a backup of this logical-log file, the external backup is not restorable.

5. After you perform an external backup, back up the current log. using the `ontape -a` command.

Results

If you lose a disk or the whole system, you are now ready to perform an external restore.

Prepare for an external backup

These topics describe the commands used to prepare for an external backup. For the procedure, see [Performing an external backup on page 175](#).

Block and unblock the database server

This section shows the syntax of the `block` and `unblock` commands.

```
onmode -c [{block | unblock}]
```

Element	Purpose	Key considerations
-c	Performs a checkpoint and blocks or unblocks the database server	None.
block	Blocks the database server from any transactions	Sets up the database server for an external backup. While the database server is blocked, users can access it in read-only mode. Sample command: <code>onmode -c block</code>
unblock	Unblocks the database server, allowing data transactions and normal database server operations to resume	Do not unblock until the external backup is finished. Sample command: <code>onmode -c unblock</code>

Track an external backup

The database server and ontape do not track external backups. To track the external backup data, use a third-party storage manager or track the data manually. The following table shows the items we recommend that you track in an external backup.

Table 29. Items to track when you use external backup and restore

Items to track	Examples
Full path names of each chunk file for each backed up storage space	UNIX™: <code>/work/dbspaces/rootdbs</code> Windows™: <code>c:\work\dbspaces\rootdbs</code>
Object type	Critical dbspaces, noncritical storage spaces
<code>ins_copyid_hi</code> and <code>ins_copyid_lo</code>	Copy ID that the storage manager assigns to each backup object
Backup date and time	The times that the database server was blocked and unblocked
Backup media	Tape volume number or disk path name
Database server version	Version 14.10

Data that is restored in an external restore

If you lose a disk or the whole system, you can externally restore data only if it was externally backed up. You must use the same third-party utility for both the external backup and restore. To externally restore the storage spaces, copy the backed up data to disk. Use the `ontape -p -e` command to mark the storage spaces as physically restored, replay the logical logs with the `ontape -l` command, and bring the storage spaces back online. If you do not specify an external restore command, the database server cannot update the status of these storage spaces to online.

You can only perform a cold external restore with ontape. A cold external restore marks storage spaces as physically restored, then performs a logical restore of all storage spaces.

When you perform a cold external restore, ontape does not first attempt to salvage logical-log files from the database server because the external backup has already copied over the logical-log data.

To salvage logical logs, perform `ontape -S` before you copy the external backup and perform the external restore (`ontape -p -e`).

Use external restore commands

Use the `ontape -p -e` command to perform a cold external restore. This command marks the storage spaces as physically restored. The following diagram shows the external physical restore syntax.

Perform an external physical restore

`-p-e`

Element	Purpose	Key considerations
<code>-e</code>	Specifies an external restore	Must be used with the <code>-p</code> option.
<code>-p</code>	Specifies a physical restore	In a cold restore, if you do not specify storage space names, all of them are marked as restored. After the physical restore completes, you must perform a logical restore.

Use the `ontape -l` command to perform a logical restore. For more information, see [ontape utility syntax: Perform a restore on page 159](#).

Rules for an external restore

Before you begin an external restore, know what you can and cannot restore from an external backup and be aware of the rules for an external restore.

These requirements and rules are:

- You must externally restore from an external backup. Although the external backup is treated as a level-0 backup, it might actually be an incremental backup that was created by another source than Informix®.
- You cannot externally restore temporary dbspaces.
- You cannot externally restore from regular `ontape` backups.
- You cannot verify that you are restoring from the correct backup and that the storage media is readable with `ontape`.
- If the external backups are from different times, the external restore uses the beginning logical log from the oldest backup.
- Salvage the logical logs (`ontape -l`) before you switch the disks that contain the critical storage spaces.
- If you are restoring critical dbspaces, the database server must be offline.
- If you are restoring the rootdbs, disable mirroring during the restore.
- The external backups of all critical dbspaces of the database server instance must have been simultaneous. All critical dbspaces must have been backed up within the same `onmode -c block onmode -c unblock` command bracket.

Rename chunks

You can rename chunks in an external cold restore by using the rename options syntax for other restores.

Use the following commands to rename chunks during an external cold restore:

```
ontape -p -e -rename -f filename
```

or

```
ontape -p -e -rename -p old_path -o old_offset-n new_path-o new_offset
```

Performing a cold external restore

About this task

If you specify the `ontape -p -e` command in a cold restore, you must restore all storage spaces. Use the `ontape -p -e` command to restore all storage spaces.

To perform a cold external restore:

1. Shut down the database server with the `onmode -ky` command.
2. To restore the storage spaces from an external backup, use a copy command, such as `cp`, `dd`, or `tar` on UNIX™ or a file-backup program.

You must restore the storage spaces to the same path as the original data.

3. To perform an external restore of all storage spaces followed by a logical restore, use the following commands:
 - `ontape -p -e`
 - `ontape -l`

Examples of external restore commands

The following table contains an example of external restore commands.

External restore command	Action	Comments
<code>ontape -p -e</code> <code>ontape -l</code>	Physical external restore and logical restore	The system restores the logical logs from the oldest external backup.
<code>ontape -p -e -rename -f</code>	External cold restore with renamed chunks	

Initializing HDR with an external backup and restore

About this task

You can use external backups to initialize High-Availability Data Replication (HDR).

To initialize HDR with an external backup and restore:

1. Block the source database server with the `onmode -c block` command.
2. Externally back up all chunks on the source database server.
3. When the backup completes, unblock the source database server with the `onmode -c unblock` command.
4. Make the source database server the primary server with the following command: `onmode -d primary secondary_servename`
5. On the target database server, restore the data from the external backup with a copy or file-backup program.
6. On the target database server, restore the external backup of all chunks with the `ontape -p -e` command.
7. Make the target database server the secondary server with the following command: `onmode -d secondary primary_servename`
8. If the logical-log records written to the primary database server since step 1 still reside on the primary database server disk, the secondary database server reads these records to perform the logical recovery. Otherwise, perform the logical recovery with the `ontape -l` command.

The `database server operational` messages appear in the message log on the primary and secondary servers.

Backup and restore a Remote Secondary Server(RSS)

It is possible to archive a Remote Secondary Server and to back up logical logs on that node, using [onbar on page 20](#) or [ontape on page 132](#). This archive may then be used to rebuild the RSS if necessary, saving time over using an archive that is taken on the primary and copied to the secondary machine.

Although an archive and log backups taken on an RSS node may be used to restore a primary node, they are recommended for this purpose only when no other archive exists. The logical log position on the secondary is often several logs behind that of the primary, which means the last logical logs backed up on the RSS may not be the last log completed on the primary. This is not a problem when the archive is used to recreate the RSS, because once reconnected to the primary the logs will be resynchronized and no transactions will be lost.

Prerequisites for taking archives and log backups on an RSS node are as follows:

1. Enable the feature by setting the [BAR_SEC_ALLOW_BACKUP on page 238](#) configuration parameter to 1 and restarting the RSS. This parameter may not be tuned dynamically.
2. Ensure that the RSS node has one or more active [temporary dbspaces on page](#) , and that they are listed in the [DBSPACETEMP on page](#) configuration parameter. Once an archive begins, all pages physically logged for a particular space will need to be stored until that space has been archived. The temporary dbspaces listed in DBSPACETEMP are used for this before-image storage. The total amount of temporary space required will vary according to the update load on the RSS during the archive, which will fail to complete if it runs out of temporary space.
3. Set the [TAPEDEV, LTAPEDEV, and BAR-related configuration parameters on page](#) to appropriate values depending on your preferred backup utility.
4. The primary node must not contain any of the following non-logged objects:
 - 1) BLOB spaces
 - 2) Non-logged smartblobs

3) No-log databases

4) Raw tables

If any of these unlogged objects are present in the instance, an archive taken on the RSS node will fail because the archive would be incomplete and therefore unusable for restoration on a primary node.

5. In order to back up logical logs on an RSS node the [LTAPEDEV on page 246](#) configuration parameter must *not* be set to the Null device. If logical logs will not be backed up on the RSS node, LTAPEDEV *must* be set to the Null device. Note that these two rules apply on an RSS node only when BAR_SEC_ALLOW_BACKUP is set to 1. When BAR_SEC_ALLOW_BACKUP is set to 0 the setting of LTAPEDEV on the RSS node must be in sync with that of the primary, and no logical log backups will be allowed or required on the RSS node regardless of the LTAPEDEV setting.

If it becomes necessary to recreate the RSS node from an archive and log backups taken on that node, the same procedures used for restoring archives taken on the primary may be used with the local archive. No new configuration changes or steps are required.

Even when taking archives and log backups on an RSS node, it is recommended that archives and log backups be taken on the primary as well, because restoration of the primary is likely to be faster with a local backup and there is a greater chance that all committed transactions will be salvaged and restored. Again, the risk of lost transactions is not an issue when recreating the RSS from a local backup as long as the primary node can forward all missing logs to the RSS during synchronization.

In an emergency it is possible to perform a cold restore of a primary node using an archive taken on an RSS and log backups taken either on the RSS or the primary. This process is made simpler with the ontape utility but is also feasible with onbar.

- If using ontape, simply restore the archive and roll forward the log backups as usual, regardless of the origin of these backup elements.
- If using onbar the storage manager containing all backup objects must be available on the primary, and you must replace the [ixbar file on page](#) on the primary with the ixbar file from the RSS node before performing the restore.

Using an archive taken on an RSS node for a warm restore on the primary node is not recommended.

Related reference

[BAR_SEC_ALLOW_BACKUP configuration parameter on page 238](#)

Informix® Primary Storage Manager

The HCL Informix® Primary Storage Manager manages storage for ON-Bar backup and restore operations, including parallel backups, that use file devices (disks).

HCL Informix® Primary Storage Manager

HCL Informix® Primary Storage Manager is an application that manages storage devices used for backup and restore requests that are issued by ON-Bar. This storage manager supports both serial and parallel processing for backup and restore requests.

Informix® Primary Storage Manager consists of the following components:

onpsm utility

A command-line utility that you can use to perform the following tasks:

- Create, modify, and delete storage devices
- Define and modify the maximum sizes for devices
- Move backup information from one device to another within a device pool
- Determine whether volumes, storage objects, and devices are locked or busy
- Release locked volumes, storage objects, and devices
- Verify volume names and labels

XBSA shared library

A unique version of the X/Open Backup Services API (XBSA) shared library that ON-Bar and the Informix® Primary Storage Manager use to communicate with each other. When ON-Bar stores or retrieves data that is stored on storage devices, the storage manager coordinates the request through the XBSA interface at the device level. You specify the location of the XBSA shared library with the `BAR_BSALIB_PATH` configuration parameter.

Storage catalog tables

A set of flat files that track information about all storage objects, devices, and device pools. These files are required to restore backup objects that are created by Informix® Primary Storage Manager. By default, these files are stored in the `$INFORMIXDIR/etc/psm` directory. You can use the `PSM_CATALOG_PATH` configuration parameter to specify another location for the storage catalog tables.



Important:

- Back up the storage catalog tables with your operating system tools as part of a disaster recovery strategy. The storage catalog tables are not backed up with the database instance and they are not associated with system catalog tables.
- To prevent the storage catalog tables from getting too large, delete old generations of backups regularly. Use the `onmsync` utility to manage expiration policies.

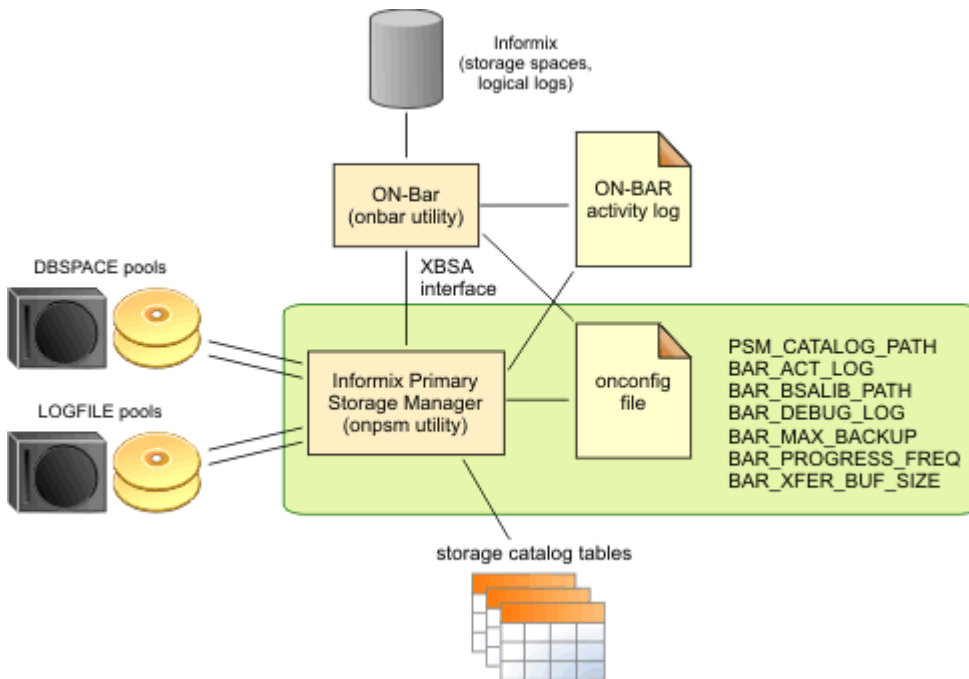
The configuration parameters that you use to configure the Informix® Primary Storage Manager are in the `onconfig` file.

You define and maintain storage devices with the `onpsm` command-line utility. You can configure one device at a time or generate a device-configuration file to configure multiple devices. During backups, Informix® Primary Storage Manager

selects a device from a pool of available devices. If the device becomes full or fails, the storage manager automatically moves to another device in the same pool.

Informix® Primary Storage Manager writes informational, warning and error messages to the storage manager activity log. You can use the PSM_ACT_LOG configuration parameter to specify the location of the activity log. If the PSM_ACT_LOG configuration parameter does not contain information, the storage manager puts activity information in the directory specified with the BAR_ACT_LOG configuration parameter.

Figure 28. Components of Informix® Primary Storage Manager



Storage manager feature	Explanation
Storage devices to use with the storage manager	File devices only The storage manager automatically creates a default device when a catalog is created. The default device is \$INFORMIXDIR/backups. You can remove the default device.
Buffer transfer size	Unlimited
Encryption and compression	Achieved with BACKUP_FILTER, RESTORE_FILTER FILTERS in ON-Bar (The storage manager does not provide encryption or compression.)
Expiration policies of the storage manager	No expiration policies. (You manually expire backup objects from the storage manager with the onmsync utility. The onmsync object expiration commands remove objects from the storage manager.)

You can perform an imported restore with ON-Bar and the Informix® Primary Storage Manager. In an imported restore, you back up the Informix® instance on one machine and restore the instance on a different machine. Use the `onmsmsync` export and import options to export the backup objects from the storage manager on the backup machine and import the backup objects into the storage manager on the restore machine.

Backups to Cloud and STDIO devices

Using STDIO devices for backup and restore:

- PSM will write/read a stream of data to an external utility (i.e. `sftp` or `curl`)
- Operator to provide parameters to invoke the utility for read/write/drop operations
- Transfer of data to the third party program will happen using STDIO, PSM will write to the standard input of the utility and will read from its standard output
- The Operator do not have direct access to the stream of data
- In order to use this feature the operator must create a PSM device of type "STDIO"
- A device of type STDIO will require you provide the path to the program to be executed as the device name (i.e. `/usr/bin/curl`)
- Also you must provide the arguments to call the program during backup, restore and drop
- You can optionally provide the maximum size of a file, if the backup is bigger than this size, it will be broken up in pieces of this size

The new command line is:

```
onpsm -D -add /usr/bin/sftp.sh -t STDIO --stdio_warg "BACKUP @obj_name1@.@obj_id@.@obj_part@"
--stdio_rarg "RESTORE @obj_name1@.@obj_id@.@obj_part@"
--stdio_darg "DELETE @obj_name1@.@obj_id@.@obj_part@" --max_part_size <size in KB>
```

Related reference

[Configure a storage manager on page 27](#)

[The `onpsm` utility for storage management on page 192](#)

[Informix Primary Storage Manager configuration parameters on page 260](#)

Related information

[Backup Services API \(XBSA\) on page 22](#)

[Device pools on page 201](#)

[Examples: Manage storage devices with Informix Primary Storage Manager on page 183](#)

Examples: Manage storage devices with Informix® Primary Storage Manager

Learn how to set up and use Informix® Primary Storage Manager to manage storage devices that the `onbar` utility uses for backing up and restoring instances. Each example shows how you can use the storage manager for a specific backup strategy.

Before you begin

Prerequisites:

- Informix® 14.10 is installed with the ON-Bar utility.
- Environment variable INFORMIXDIR is set to the path where the database server is installed.
- Environment variable ONCONFIG is set to the file in `$INFORMIXDIR/etc` that contains the configuration parameters for your database. The name of the file must be unique for each database server instance.
- User **informix** or root privileges.

About this task

[Example 1: Storing backups for an instance on page 184](#)

[Example 2: Storing backups for two instances on page 186](#)

[Example 3: Exporting backups to and restoring them from another directory on page 188](#)

[Example 4: Exporting a backup from one server and importing it into another server on page 189](#)

In these examples, *storage manager* refers to Informix® Primary Storage Manager.

Related reference

[Informix Primary Storage Manager configuration parameters on page 260](#)

[The onpsm utility for storage management on page 192](#)

[The onmsync utility on page 107](#)

Related information

[HCL Informix Primary Storage Manager on page 181](#)

Example 1: Storing backups for an instance

This example shows how to set up and use Informix® Primary Storage Manager to back up the data and logical logs for a single database server instance to a directory: `$INFORMIXDIR/backups`.

About this task

In this example, you update the configuration file so that the Informix® Primary Storage Manager can communicate with ON-Bar and you specify the directory where you want backups stored. Then you use the **onbar** utility to perform a standard, level-0 backup of all online storage spaces and used logical logs. You validate the backup by checking the messages that were logged and by using the **onpsm** utility to confirm that storage objects were created.

1. Set the BAR_BSALIB_PATH configuration parameter to the full path and name of the shared library for the storage manager.

Example

For example, on Linux™, Solaris:

```
BAR_BSALIB_PATH $INFORMIXDIR/lib/libbsapsm.so
```


You must use the version of the XBSA shared library that is provided for Informix® Primary Storage Manager. If you do not specify the path with the `BAR_BSALIB_PATH` configuration parameter, you must ensure that the XBSA library is in the default location on your operating system.

2. If needed, create the directory in which to store the backup objects.

By default, the storage manger includes the default pools LOGPOOL and DBSPPOOL, with the default directory `$(INFORMIXDIR)/backups` in each pool.

- If you want to use the default backup directory, verify that the `$(INFORMIXDIR)/backups` directory exists.
- If you want to use a different backup directory, use the `onpsm -D add` command to add a new backup directory for LOGPOOL and DBSPPOOL. For example, run the following commands to add different backup directories for the LOGPOOL and DBSPPOOL pools:

```
onpsm -D add /backups/infx/logs -g LOGPOOL -p HIGHEST -t FILE
onpsm -D add /backups/infx/spaces -g DBSPPOOL -p HIGHEST -t FILE
```

Use the HIGHEST priority for the device that should be filled first. Only one device in a pool can have the priority setting of HIGHEST.

3. Run the **onbar** utility to perform a standard, level-0 backup of all online storage spaces and used logical logs.

Example

```
onbar -b -L 0
```

Result

If the storage catalog tables do not exist, they are created in the `$(INFORMIXDIR)/etc/psm` directory.

4. Validate that the storage manager is set up and that the backup objects are created.
 - a. Look in the ON-Bar activity log to confirm that the storage manager is ready and that ON-Bar recognizes the storage manager.

Example

For example, the first message is from the storage manager and the second message is from the backup utility:

```
2012-01-03 15:51:23 11193 2569 Informix PSM is ready.
2012-01-03 15:51:23 11193 2569 Using Informix PSM version 14.10.FC1
as the Storage Manager. XBSA API version is 1.0.3.
```

By default, the storage manager posts messages to the ON-Bar activity log. The location of the activity log is set by the `BAR_ACT_LOG` configuration parameter. If you want the storage manager messages to be logged separately, you must set the `PSM_ACT_LOG` configuration parameter.

- b. Run the `onpsm -O list` command to list the storage objects that were created:

Result

The list, as shown in the following example, includes the storage object IDs, the date the storage objects were created, the size of the storage objects, and where the storage objects are in the storage device. The object IDs are also stored in the `ixbar` file and are used by ON-Bar to locate the objects.

```

=====
Object List Report

Obj ID   Date Created      Size (MB)   Logical Path
-----
         2012-08-06 12:02:10   12.5   /serv1/rootdbs/0/serv1.1
         2012-08-06 12:02:12     0.1   /serv1/logdbs/0/serv1.1
         2012-08-06 12:02:12     0.1   /serv1/dbs2/0/serv1.1
         2012-08-06 12:02:12     0.1   /serv1/dbs1/0/serv1.1
         2012-08-06 12:02:13     0.1   /serv1/physdbs/0/serv1.1
         2012-08-06 12:02:14     0.3   /serv1/10/9/serv1.1
         2012-08-06 12:02:14     0.0   /serv1/crit_files/ixbar/serv1.1
         2012-08-06 12:02:14     0.0   /serv1/crit_files/oncfg/serv1.1
         2012-08-06 12:02:14     0.1   /serv1/crit_files/onconfig/serv1.1
         2012-08-06 12:02:14     0.0   /serv1/crit_files/sqlhosts/serv1.1
=====

```

c. Run the `onpsm -D list` command to display a list that shows that the device was added to the DBSPool and LOGPOOL pools. The following example shows output of the command:

Result

Type	Prio	Block/Size (MB)	Pool Name	Device Name---
FILE	HIGHEST	--/--	DBSPool	/backups/infx/logs
FILE	HIGHEST	--/--	LOGPOOL	/backups/infx/spaces

Results

With a few simple steps, you configured the storage manager and performed a full backup of an instance to a file device. Very little configuration was required because the storage manager uses the default settings for various ON-Bar configuration parameters.

What to do next

Storage catalog tables are not included in a backup. Be sure to back up the storage catalog tables with your operating system tools as part of a disaster recovery strategy. If the storage catalog tables are lost, the **onbar** utility cannot restore the backup objects that Informix® Primary Storage Manager created. The location of the storage catalog tables is set by the `PSM_CATALOG_PATH` configuration parameter (default = `$INFORMIXDIR/etc/psm`).

To restore the instance from the backup objects, use the **onbar** utility. The storage manager tracks the backup objects and storage devices for you.

Example 2: Storing backups for two instances

This example shows how to configure one instance of Informix® Primary Storage Manager to manage the storage devices for two database server instances in a multiple residency environment.

About this task

In this example, you set up two independent database server environments on the same computer. Each database server is installed in a separate directory: (`/usr/informix/ids1210fc1` and `/usr/informix/ids1210fc1b`) and has a database server instance. Storage for backup operations on both database server instances is managed by one instance of Informix® Primary Storage Manager. Pools of storage devices for physical and logical data are configured for each instance.

1. For *each* instance, edit the **onconfig** file to configure storage management for ON-Bar.

Table 30. Configuration parameters and their associated values

Configuration parameter	Value
BAR_BSA LIB_PATH Specify the full path and name of the shared library for the storage manager.	<code>/usr/informix/ids1210fc1b/lib/libbsapsm.so</code>
PSM_CATALOG_PATH Specify the path of the storage catalog tables.	<code>/usr/informix/ids1210fc1b/etc/psm</code>
PSM_DBS_POOL Specify the name for a group of devices for storing online data (dbspace) backups.	FC1: <code>DBSPool_FC1</code> FC1B: <code>DBSPool_FC1B</code>
PSM_LOG_POOL Specify the name for a group of devices for storing online logical log backups.	FC1: <code>LOGPool_FC1</code> FC1B: <code>LOGPool_FC1B</code>

2. For *each* instance, create a directory in which to store the backup objects.

Example

```
mkdir $INFORMIXDIR/backups/dev_for_1201fc1
mkdir $INFORMIXDIR/backups/dev_for_1201fc1b
```

3. Run the **onpsm** utility to create device pools for each instance. For example, specify:

Example

```
onpsm -P add DBSPool_FC1
onpsm -P add LOGPool_FC1
onpsm -P add DBSPool_FC1B
onpsm -P add LOGPool_FC1B
```

4. Run the **onpsm** utility to add the storage devices.

Example

```
onpsm -D add $INFORMIXDIR/backups/dev_for_1201fc1 -t FILE -g DBSPool_FC1
onpsm -D add $INFORMIXDIR/backups/dev_for_1201fc1 -t FILE -g LOGPool_FC1
onpsm -D add $INFORMIXDIR/backups/dev_for_1201fc1b -t FILE -g DBSPool_FC1B
onpsm -D add $INFORMIXDIR/backups/dev_for_1201fc1b -t FILE -g LOGPool_FC1B
```

5. For *each* instance, run the **onbar** utility to perform a standard, level-0 backup of all online storage spaces and used logical logs.

Example

```
onbar -b -L 0
```

Result

6. Validate that the storage manager is set up and that the backup objects are created.

- a. For *each* instance, look in the ON-Bar activity log to confirm that the storage manager is ready and that ON-Bar recognizes the storage manager. For example, look for this information:

Example

```
2012-01-03 15:51:23 11193 2569 Informix PSM is ready.
2012-01-03 15:51:23 11193 2569 Using Informix PSM version 14.10.FC1
as the Storage Manager. XBSA API version is 1.0.3.
```

- b. Use the **onpsm** utility to list the storage objects that were created:

Example

```
onpsm -O list
```

Result

The report includes the storage object IDs, the date the storage objects were created, the size of the storage objects, and the location of the storage objects in the storage device.

Example 3: Exporting backups to and restoring them from another directory

This example shows how to export backups to a new directory and import the backup objects from that directory.

About this task

Suppose that you keep five generations of backups. As an added precaution, you also keep copies of the most recent backups in a separate directory. In this example, you use the **onmsync** utility to export your most recent backup to and import it from the Informix® Primary Storage Manager external pool in a separate directory.

The storage manager tracks devices in the external device pool (EXTPOOL) so it can copy objects to and from external devices. (Although the storage manager tracks devices, it does not track files and objects that are inside the EXTPOOL pool in the storage manager catalogs.)

1. Store backups for an instance, following the steps in [Example 1: Storing backups for an instance on page 184](#).
2. Run the **onpsm -D list** command to check that there is a device in the EXTPOOL pool.

- a. If there is no device in the EXTPOOL pool, add one using the **onpsm -D add** command.

Example

The following example shows how to add a device with the path `/export/informix/psm_exportdir` to the EXTPOOL pool.

```
$ onpsm -D add /export/informix/psm_exportdir -g EXTPOOL -t FILE
```

3. Run the `onsmsync` command to export all backup objects in the generation 1 level-0 backup, using prefix `pw_sept5`, which becomes the name of the subdirectory in which the utility places the backup:

Example

```
onsmsync -E -p pw_sept5 -g 1
```

After you run the `onsmsync -E` command to export the backup objects, you will see a subdirectory in the `EXTPOOL` directory that includes a directory holding the backup objects and a file called `export.bom`.

What to do next

Suppose that something happens to the backup generation stored in your primary backup directory and you want to import the `pw_sept5` backup generation from the second directory. To import the backup generation:

1. Run the `onsmsync` command to import all backup objects in the `pw_sept5` subdirectory:

```
onsmsync -I -p pw_sept5
```

Use your own file-transfer methods to move the exported backups, as needed, to other machines.

Example 4: Exporting a backup from one server and importing it into another server

This example shows how to use the `onsmsync` utility to export a backup from a database server that has the name `informix_serv1`. Then the example shows how to use the `onsmsync` utility to import the data into a server that has the name `informix_serv2`.

1. Set up and export files on database server `informix_serv1`:
 - a. Set the `INFORMIXDIR`, `INFORMIXSERVER`, `ONCONFIG`, `PATH`, `INFORMIXSQLHOSTS` environment variables for `informix_serv1`.
 - b. Run the `onpsm -D list` command to check that there is a device in the `EXTPOOL` pool. If there is no device in the `EXTPOOL` pool, add one using the `onpsm -D add` command.
 - c. Run the `onsmsync` command to export all backup objects in the generation 1 level-0 backup, using prefix `serv1_20120810`, which becomes the name of the subdirectory in which the utility places the backup:

Example

```
$ onsmsync -E -p serv1_20120810 -g 1
```

2. Prepare to import files on the second database server, `informix_serv2`, as follows:
 - a. Set the `INFORMIXDIR`, `INFORMIXSERVER`, `ONCONFIG`, `PATH`, `INFORMIXSQLHOSTS` environment variables for `informix_serv2`.
 - b. Run the `onpsm -D list` command to determine if the `EXTPOOL` has the same device that you viewed or added in step 1b. (This could occur for shared devices). If there is no device in the `EXTPOOL` pool, add one using the `onpsm -D add` command.

- c. Copy the previously exported backup objects (for example, subdirectory `serv1_20120810`) into the EXTPOOL device from which you will import the backup objects.
- d. Run the following command to import backup objects from EXTPOOL device:

Example

```
$ onmsync -I -p serv1_20120810
```

Result

After you run the `onmsync -I` command to import the backup objects, the objects are stored in the new LOGPOOL and DBSPool pools.

- e. Run the `onpsm -O list` command to view the imported objects.

Result

Notice that the import command also creates a new `ixbar` file in `$INFORMIXDIR/etc/` directory.

```
$ ls -l $INFORMIXDIR/etc/*ixbar*
-rw-rw-- 1 informix informix    0 Aug 10 19:44
  /usr/informix/etc/ixbar.12.20120810.194441
-rw-rw-- 1 informix informix 2704 Aug 10 19:44
  /usr/informix/etc/ixbar.12
```

The new `ixbar` file lists the imported backup objects so that you can perform an ON-BAR cold restore to restore the `informix_serv1` instance from the first database server to the `informix_serv2` instance on the second database server.

Setting up Informix® Primary Storage Manager

Setting up involves gathering and specifying information about your storage devices and, if necessary, changing the default configuration of the storage manager.

Related reference

[Configure a storage manager on page 27](#)

[The onpsm utility for storage management on page 192](#)

[Informix Primary Storage Manager configuration parameters on page 260](#)

Collecting information about file directories and devices

You must gather information about and configure at least one file directory or device for each of the DBSPool and LOGPOOL pools before ON-Bar can use the HCL Informix® Primary Storage Manager.

About this task

Before defining directories or devices, gather the following information:

- The full path names and types of the devices that you plan to use for your backup storage.
- The amount of space that you want to commit to ON-Bar backups.

Configuring Informix® Primary Storage Manager

By default, the HCL Informix® Primary Storage Manager is automatically configured with the information specified in the storage manager and with some ON-Bar configuration parameters. It is also automatically configured when you use the `onpsm` utility. You can change the configuration.

About this task

The Informix® Primary Storage Manager uses file devices (disks) only, not tapes. You cannot configure the storage manager to use tapes.

To manually configure the Informix® Primary Storage Manager:

1. Update the `BAR_BSALIB_PATH` configuration parameter to point to the storage manager library.

For example, on Linux™ or Solaris, specify:

```
BAR_BSALIB_PATH $INFORMIXDIR/lib/libbsapsm.so
```

2. Specify the destination and source devices for backup and restore operations by using the **onpsm** utility.
3. Change the default configuration for the storage manager if necessary for your environment:
 - a. To override the default values for the location of storage manager log files and catalogs, debugging activity, and pool names, specify new values in the Informix® Primary Storage Manager configuration parameters.
 - b. To specify a larger transfer buffer with ON-Bar and the Informix® Primary Storage Manager, increase the size in the `BAR_XFER_BUF_SIZE` configuration parameter.
 - c. To change the frequency of the progress messages in the ON-Bar activity log, update the value specified in the `BAR_PROGRESS_FREQ` configuration parameter.
 - d. To change the number of processes that ON-Bar runs concurrently, update the value specified in the `BAR_MAX_BACKUP` configuration parameter.

Related reference

[Informix Primary Storage Manager configuration parameters on page 260](#)

[ON-Bar and ontape configuration parameters and environment variable on page 221](#)

[The onpsm utility for storage management on page 192](#)

Managing storage devices

Use the `onpsm` utility to add, monitor, and remove storage devices and to manage HCL Informix® Primary Storage Manager catalogs, locks, and objects. Use the `onsmsync` utility to export ON-Bar backups to and import them from external pools and to expire backups.

Related reference

[The onpsm utility for storage management on page 192](#)

[The onsmsync utility on page 107](#)

The onpsm utility for storage management

Use the onpsm utility to manage the HCL Informix® Primary Storage Manager catalogs, devices, locks, and objects.

Pre-requisite: To run the onpsm utility, you must be user **root** or **informix**.

Syntax

```
onpsm { { <Catalog options> | <Device options> | <Object options> | <Pool options> } | [ { -h } ] | [ { -v | -version | -version all } ] }
```

Catalog options

```
" -C { check [ { -l | -n } ] | [ detail ] | [ export ] | [ { import | -y } ] | [ { init | init -d | -y } ] | [ { unlock } ] }
```

Device options

```
" -D { add path -p priority -g pool_name -t type [ -s size ] | del path { -g pool_name [ { -y } ] | -d [ { -y } ] } | list { -u | -l dev_def_file.txt } | purge path | scan path | update path [ -p priority ] [ -s size ] }
```

Object options

```
" -O { del -o object_id [ { -y } ] | detail -o object_id | dump -o object_id | list }
```

Pool options

```
" P { add pool_name | del pool_name [ { -y } ] | list }
```

Table 31. onpsm utility catalog options

Element	Purpose	Key Considerations
-C check	Checks storage manager catalog tables, which store metadata about the pools	This command identifies files that have problems.

Table 31. onpsm utility catalog options (continued)

Element	Purpose	Key Considerations
	and devices that the storage manager manages	
-C check -l	Displays index keys while checking the catalog tables	
-C check -n	Indicates that the storage manager does not fix errors that are found	
-C detail	Shows details about the storage manager catalog tables	
-C export	Exports the Informix® Primary Storage Manager catalog tables to a directory called <code>psm_catalog.exp</code>	
-C import	Replaces the current Informix® Primary Storage Manager catalog with a catalog that is recreated from files that are in the <code>psm_catalog.exp</code> directory	Only import the catalog if you have a system problem, lost your current catalog, and need to revert to the exported catalog. If you need to import a catalog, run the <code>onpsm -C init</code> command before you run the <code>onpsm -C import</code> command.
-C init	Deletes storage manager catalog tables	
-C init -d	Deletes storage manager catalog table and the backup objects in file devices	
-C unlock	Unlocks the storage manager catalog	If the storage manager exits abnormally from a backup or restore session because a failure occurred, storage manager catalog tables might remain locked. If catalog tables are locked, you can release the locks.
-y	Specifies to not to ask for confirmation before deleting catalog tables	

Table 32. onpsm utility device options

Element	Purpose	Key Considerations
-D add	Adds a device to the pool specified with the <code>-g</code> option	Before adding devices, gather information about the device. See Collecting information about file directories and devices on page 190 .
-D del	Removes a device:	If you delete the device using <code>-g</code> option, you can restore the objects if necessary.

Table 32. onpsm utility device options (continued)

Element	Purpose	Key Considerations
	<ul style="list-style-type: none"> • If you use the -g option, removes a device from the pool specified with the -g option, while retaining the device objects in the Informix® Primary Storage Manager catalog. • If you use the -d option, removes the device from all pools and removes all backup objects from the file system in that device 	<p>If you remove a device, the storage manager cannot add new objects to the device.</p>
-D list	Displays a list of all devices in the system	
-D purge	Removes missing storage manager objects from the Informix® Primary Storage Manager catalog	
-D scan	<p>Scans objects in the device to verify that the objects exist in the Informix® Primary Storage Manager catalog so the objects can be restored if necessary</p> <p>If an object is not in the catalog, this command adds the object to the catalog.</p>	<p>If the command cannot add an object to the catalog, the command ignores the missing file.</p> <p>Before a missing object can be added to a catalog, the following conditions must occur:</p> <ul style="list-style-type: none"> • The object ID must not be assigned to any other object in the storage manager. • Files must not be renamed or relocated to different directories inside the device • The object version must not be assigned to any other object in the storage manager.
-D update	Modifies information about a device	<p>If you want to modify information about more than one device, run a separate command for each device.</p>
<i>path</i>	Full name and path to the device (for TAPE devices) or to a directory (for FILE devices)	<p>The path must be in the format appropriate to the operating system to which the device is attached .</p> <p>The name of the device must be unique within a pool.</p> <p>You can include the same device in multiple pools.</p> <p>If you are deleting, listing, purging, scanning, or updating information, the path must be to an existing device.</p>

Table 32. onpsm utility device options (continued)

Element	Purpose	Key Considerations
-d	Deletes the pool from all pools and deletes backup objects	
-g <i>pool_name</i>	The pool in which to add the device, either DBSPool, LOG POOL, or EXTPool	<p>Information about pools is stored in the Informix® Primary Storage Manager catalog.</p> <p>If you do not provide a pool name, the command fails.</p> <p>Specify:</p> <ul style="list-style-type: none"> • DBSPool for backups of dbspaces, blobspaces, and sbspaces • LOGPOOL for backups of logical logs • EXTPool that serves as a staging area from which you can move specific backups or backup generations to permanent storage or onto a different computer.
-l <i>dev_def_file.txt</i>	Loads information about the device from a device-definition file	
-p <i>priority</i>	Priority of the device, either HIGHEST, HIGH, LOW, or READ-ONLY	<p>The storage manager fills high-priority devices in a pool before placing data into low-priority devices in that pool. If the high-priority devices are busy at the moment when the storage manager is ready to fill a pool, the storage manager uses low priority devices.</p> <p>Only one device in a pool can have the priority of HIGHEST.</p> <p>If multiple devices have the same priority in the same pool, the storage manager determines which device to use first.</p> <p>When a device becomes full, the storage manager changes the priority to READ-ONLY. You can change the priority after you add more space to the device.</p>
-s <i>size</i>	For tape devices only, the maximum storage capacity of the device in kilobytes	The size is optional for tape devices. If a size is not specified, or if you specify 0, the storage manager interprets the size as unlimited. When the size is unlimited, the device is not considered full until it returns an error that specifies that the device is full.

Table 32. onpsm utility device options (continued)

Element	Purpose	Key Considerations
		To specify the size enter the numeric value of the size followed by the suffix B, K, M, G, T, or P (for bytes, kilobytes, megabytes, gigabytes, terabytes, or petabytes). The suffix can be upper or lowercase.
-t <i>type</i>	Type of device, either FILE or TAPE	Information about devices is stored in the Informix® Primary Storage Manager catalog,
-u	Unloads information about the device to a device-definition file	The device definition file is a text file with a specific format. The storage manager uses the file to recreate the information when you run an onpsm command with the load option.
-y	Specifies not to ask for confirmation to complete the requested action	

Table 33. onpsm object options

Element	Purpose	Key Considerations
-O del	Deletes physical objects from a pool	
-O detail	Displays details about the specified object. Details include the location of the object.	
-O dump	Extracts the object data to a file in the current directory	
-o <i>object_id</i>	Identifies the particular object	You can delete or dump one or more objects with a single command, as shown in Usage on page 197 .
-O list	Displays all objects in a pool	For each object, the list includes the date and time the object was created, the size of the object, and the path name of the object.
-y	Specifies not to ask for confirmation to complete the requested action.	

Table 34. onpsm pool options

Element	Purpose	Key Considerations
-P add <i>pool_name</i>	Adds a new pool	
-P del <i>pool_name</i>	Deletes the specified pool	

Table 34. onpsm pool options (continued)

Element	Purpose	Key Considerations
-P list	Lists all pools in the system	
-y	Specifies not to ask for confirmation to complete the requested action.	

Table 35. onpsm utility general options

Element	Purpose	Key Considerations
-h	Displays help information	
-V	Displays the software version number and the serial number	For more details about the standard -V and -version options, see Obtaining utility version information on page .
-version	Displays the software version number, serial number, and additional information such as the host, operating system, build date, and the Global Language Support (GLS) version	For more details about the standard -V and -version options, see Obtaining utility version information on page .
-version all	Displays onpsm version information and information about the PSM shared library	

Usage

When you run an onpsm command to define a device, the storage manager automatically creates storage manager catalogs if they do not exist.

The default device for the storage manager is `$(INFORMIXDIR)/backups`. The device, which is low priority, is automatically created when the catalog is created. You can remove the default device.

When you create a device, the storage manager automatically creates the directory for the device if the directory does not exist. The storage manager uses the directory path that you specify in the `onpsm -D add` command.

You can delete one or more objects with a single command, for example, by running a command that has this format:

```
onpsm _0 del -o obj_1 -o obj_2
```

You can also dump one or more objects with a single command, for example, by running a command that has this format:

```
onpsm _0 dump -o obj_1 -o obj_2
```

If the data is not needed, run the `onmsync` utility to delete backup objects from the Informix® Primary Storage Manager.

Some third-party storage managers do not allow the `onmsync` utility to delete backup objects from the storage manager. If you have a third-party storage manager, you might need to manually delete backup objects that you no longer need.

Example

Examples

The following command adds a file device with the path name `$INFORMIXDIR/backups` in the DBSPool pool:

```
onpsm -D add $INFORMIXDIR/backups -g DBSPool -t FILE -p HIGH
```

The following command checks Informix® Primary Storage Manager catalog tables and indicates that the storage manager does not fix any errors found during the check:

```
onpsm -C check -n
```

The following command lists objects in pools, including the date and time the object was created, the size of the object, and the path name of the object.

```
onpsm -O list
```

Related information

[HCL Informix Primary Storage Manager on page 181](#)

[Setting up Informix Primary Storage Manager on page 190](#)

[Configuring Informix Primary Storage Manager on page 191](#)

[Managing storage devices on page 191](#)

[Examples: Manage storage devices with Informix Primary Storage Manager on page 183](#)

[Device pools on page 201](#)

onpsm -C detail output

Use the `onpsm -C detail` command to view details about the storage manager catalog tables.

Sample onpsm -C detail command output

```
D:\IFMXDATA\gacpsm>onpsm -C detail
```

```
Informix® Primary Storage Manager State:
```

```
PSM Unique ID      : 1358735848
Catalog Location   : D:\database\1210\etc\psm\
Catalog State      : Locked
Catalog Owner      : 2
Catalog Lock Mode  : Regular
```

```
Sessions:
```

Session ID	Process ID
2	1576
3	4556
4	5555

```
Informix® PSM Locked Objects
```

Session	Object Id	Date	Time	Server	Object Name
---------	-----------	------	------	--------	-------------

```

3      116 2012-12-09 20:54:34 /gacpsm_tcp /gacpsm_tcp/168/242
4      117 2012-12-09 21:07:44 /gacpsm_tcp /gacpsm_tcp/168/242

```

The output contains the following sections:

Informix® Primary Storage Manager State

Shows general information about all of the sessions that are active in the system and whether the catalog is locked.

PSM Unique ID

ID for the catalog

Catalog Location

Path to the catalog

Catalog State

Indicates whether the catalog is locked

Catalog Owner

Informix® Primary Storage Manager session ID

Catalog Lock Mode

Lock category

For example, `Regular` means that the lock is a user lock.

Sessions

Lists all of the sessions that are active in the system and the process IDs that match the sessions.

Session ID

ID of the session. This is the same ID that appears in the `Catalog Owner` field.

Process ID

Internal ID for the ON-Bar, archecker, or storage manager process that is locking the catalog.

Informix® PSM Locked Objects

Shows the locks in devices or objects that are held by storage manager sessions. For each lock, the output shows the session number, the object ID, the date and time that the lock was placed, the server, and the object name.

onpsm -D list output

The `onpsm -D list` command displays information about all of the devices in each HCL Informix® Primary Storage Manager pool. You can use this list to determine whether you need to change information about your devices.

Sample onpsm -D list command output

Type	Prio	Block/Size (MB)	Pool Name	Device Name
FILE	LOW	--/--	DBSPOOL	/informix/backups
FILE	LOW	--/--	LOGPOOL	/informix/backups

Type

Type of device, either FILE or TAPE (Currently only the FILE type is supported.)

Prio

Priority of the device, either HIGH, HIGHEST, LOW, or READ-ONLY

HIGH is the default priority if a priority is not specified. Only one device in a pool can have a priority of HIGHEST.

Block Size

Size of the device (only applies to devices of the TAPE type)

Pool Name

The name of the pool (DBSPOOL, LOGPOOL, or EXTPOOL)

In the example output above, there are no EXTPOOL devices.

Device Name

Complete path name for the device

onpsm -O list output

The onpsm -O list command displays all of the objects stored in a pool.

Sample onpsm -O list command output

Object List Report				Logical Path
Obj ID	Date Created	Size (MB)	Name.Version (omits piece #)	
1	2012-07-06 14:39:47	12.0	/gacpsm_tcp/rootdbs/0/gacpsm_tcp.1	
2	2012-07-06 14:41:18	12.0	/gacpsm_tcp/rootdbs/0/gacpsm_tcp.2	
3	2012-07-11 13:42:10	3.9	/gacpsm_tcp/160/14/gacpsm_tcp.1	
4	2012-07-11 13:42:13	3.9	/gacpsm_tcp/160/15/gacpsm_tcp.1	
5	2012-07-11 13:42:15	3.9	/gacpsm_tcp/160/16/gacpsm_tcp.1	
6	2012-07-11 13:42:15	3.9	/gacpsm_tcp/160/17/gacpsm_tcp.1	
7	2012-07-11 13:42:15	3.9	/gacpsm_tcp/160/18/gacpsm_tcp.1	
8	2012-07-11 13:42:15	3.9	/gacpsm_tcp/160/19/gacpsm_tcp.1	
9	2012-07-11 13:42:16	3.9	/gacpsm_tcp/160/20/gacpsm_tcp.1	
10	2012-07-11 13:42:16	3.9	/gacpsm_tcp/160/21/gacpsm_tcp.1	

Obj ID

The ID of the stored object

Date Created

The date and time when the object was created

Size

The size of the object in megabytes

Name.Version

The path name of the object followed by `.` and the version of the object (for example, `.2` to indicate version 2)

Device pools

The HCL Informix® Primary Storage Manager pool is a named group of disk devices that you use as a repository for backups.

When storing backup objects, the Informix® Primary Storage Manager selects particular devices from the pool and moves automatically from one device to another when devices are full or when they fail. You maintain pools by using the `onpsm` utility to add, modify, view, and drop devices in the pool.

Three pools are available:

DBSPOOL

Holds online backups of dbspaces, blobspaces, and sbspaces

LOGPOOL

Holds online backups of logical logs

EXTPOOL

Serves as a staging area for exporting a backup set of objects to a single large, external logical object or for importing the backed up objects. You can move specific backups or backup generations in this pool to permanent storage or onto a different computer. Files in the EXTERNAL pool are offline. The files are not visible to ON-Bar, and the Informix® Primary Storage Manager does not track them.

Related reference

[PSM_DBS_POOL configuration parameter on page 263](#)

[PSM_LOG_POOL configuration parameter on page 266](#)

[The onpsm utility for storage management on page 192](#)

Related information

[HCL Informix Primary Storage Manager on page 181](#)

[Informix Primary Storage Manager file-naming conventions on page 202](#)

Device-configuration file for the Informix® Primary Storage Manager

The `onpsm` utility can generate a device-configuration file, which is a text file that contains information about a storage device. The utility uses this information to re-create the devices.

The configuration file contains the following information about each device:

DEVICE

The complete path to the device

TYPE

Type of device

FILE = file directory on a disk device

POOL

The pool that contains the device, either DBSPool, LOGPOOL, or EXTPool

BLOCKSIZE

Not applicable for disk devices.

SIZE

Not applicable for disk devices.

PRIORITY

The priority of the device

For example, the device-configuration file might contain the following information

```
DEVICE=/vobs/tristarm/sqlldist/psm_backup/dbspace
TYPE=FILE
POOL=DBSPool
BLOCKSIZE=0
SIZE=0
PRIORITY=HIGH
```

Informix® Primary Storage Manager file-naming conventions

When creating files that store your backup data, the HCL Informix® Primary Storage Manager uses specific file-naming conventions.

For the DBSPool and LOGPOOL device pools, the path name of the storage file consists of:

1. Category information:
 - For space backups, the category consists of the device name, the database server name, and the space name
 - For log backups, the category consists of the device name, the database server name, server number, and a log file number
2. For space backups, the backup level, either: 0, 1, or 2
3. A version number, which is an integer that starts at 1 for an object of that category and is incremented for each subsequent backup of an object of that category.
4. An ID that identifies a piece of the backed up object

For space backups, file names have this format:

```
/ device/DBSERVERNAME/ dbspace/ backup_level/ DBSERVERNAME.version.piece
```

For example, the name of the file for the third piece of a second level 0 backup of the dbspace `rootdbs` for the server named `SERVER1` is:

```
/my_device/SERVER1/rootdbs/0/SERVER1.2.3
```

For log backups, file names have this format:

```
/device/DBSERVERNAME/SERVERNUM/LOG_UNIQUE_ID/DBSERVERNAME.version.piece
```

If you use the `onsmsync` utility with the `-E` option to export a backup generation in the Informix® Primary Storage Manager, the `onsmsync` utility creates and places the backup files in a subdirectory of the storage manager `EXTPOOL` device. You must provide a *prefix* (the name of a subdirectory) when you use the `onsmsync` utility with the `-E` or `-I` options. ON-Bar uses the specified path as a key to communication with the storage manager to store and retrieve objects.

Related reference

[The `onsmsync` utility on page 107](#)

Related information

[Device pools on page 201](#)

Message logs for Informix® Primary Storage Manager

The Informix® Primary Storage Manager writes messages to a storage manager activity log and debug log.

The message logs are stored in the directories specified in the `BAR_DEBUG_LOG` or `BAR_ACT_LOG` configuration parameters. You can use the `PSM_ACT_LOG` and `PSM_DEBUG_LOG` configuration parameters to specify another directory for each of these logs.

The `PSM_DEBUG` configuration parameter specifies the level of debugging activity that is captured in the debug log.

Related reference

[PSM_ACT_LOG configuration parameter on page 261](#)

[PSM_DEBUG_LOG configuration parameter on page 265](#)

[PSM_DEBUG configuration parameter on page 263](#)

archecker table level restore utility

archecker table level restore utility

You can use the `archecker` utility to perform point-in-time table-level restores that extract tables or portion of tables from archives and logical logs.

The `archecker` utility restores tables by specifying the source table to be extracted, the destination table where the data is placed, and an `INSERT` statement that links the two tables.

For information about using the archecker utility to verify backups, see [onbar -v syntax: Verifying backups on page 58](#).

Related reference

[The archecker utility configuration parameters and environment variable on page 254](#)

Overview of the archecker utility

The **archecker** utility is useful where portions of a database, a table, a portion of a table, or a set of tables need to be recovered. It is also useful in situations where tables need to be moved across server versions or platforms.

Use archecker utility in the following situations:

- Restore data

You can use the archecker utility to restore a specific table or set of tables that have previously been backed up with ON-Bar or ontape. These tables can be restored to a specific point in time. This is useful, for example, to restore a table that has accidentally been dropped.

The following restrictions, however, limit the functionality of the archecker utility in some table-level restore operations:

- The archecker utility is not JSON compatible. If you try to use the utility with target tables that contain columns of JSON or BSON (binary JSON) data types, the utility will abort and return an error message. (Besides data-restore contexts, this limitation affects all archecker operations on tables with JSON or BSON columns.)
- You cannot logically restore a table when restoring smart large objects. Only a physical restore of BLOB or CLOB objects is supported for table-level restore operations using the archecker utility.
- You cannot restore data from a remote device.
- You cannot use a shared memory connection when performing a table-level restore.

- Copy data

The archecker utility can also be used as a method of copying data. For example, you can move a table from the production system to another system.

The archecker utility is more efficient than other mechanisms for copying data. Because archecker extracts data as text, it can copy data between platforms or server versions.

- Migrate data

You can also use the archecker utility as a migration tool to move a table to other HCL® Informix® servers.

The archecker utility is designed to recover specific tables or sets of tables. Other situations require that you use different utilities. For example, use ON-Bar or ontape in the following data recovery scenarios:

- Full system restore
- Recovery from disk failure

To configure the behavior of the archecker utility, use the archecker configuration file. To define the schema of the data that archecker recovers, use the archecker schema command file. These files are described in the following sections.

Related reference

[ON-Bar script on page 24](#)

The archecker configuration file

The archecker utility uses a configuration file to set certain parameters.

Set the **AC_CONFIG** environment variable to the full path name of the archecker configuration file. By default, the **AC_CONFIG** environment variable is set to `$INFORMIXDIR/etc/ac_config.std`. If you set **AC_CONFIG** to a user-defined file, you must specify the entire path including the file name.

For information about the configuration parameters used in this file, see [The archecker utility configuration parameters and environment variable on page 254](#).

Schema command file

The archecker utility uses a schema command file to specify the following:

- Source tables
- Destination tables
- Table schemas
- Databases
- External tables
- Point in time the table is restored to
- Other options

This file uses an SQL-like language to provide information archecker uses to perform data recovery. For complete information about the supported statements and syntax, see [The archecker schema reference on page 212](#).

There are two methods to set the schema command file:

- Set the **AC_SCHEMA** configuration parameter in the archecker configuration file. For more information, see [AC_SCHEMA configuration parameter on page 258](#).
- Use the `-f cmdname` command-line option. For more information, see [Syntax for archecker utility commands on page 208](#).

If both methods are specified, the `-f` command-line option takes precedence.

Table-level restore and locales

For table-level restore, if the table being restored (table on the archive) has a locale code set different from the default locale (en_US.8859-1) the **DB_LOCALE** environment variable must be set to have the same code set as the locale of the archived table being restored.

No code set conversion is performed during a table-level restore; the locale code set of the database or table being restored must match the locale code set of the database or table that the data is being restored to. In addition, the same **DB_LOCALE** information is used for all of the tables being restored by using the same table-level restore command schema file.

Data restore with archecker

Use the **archecker** utility to perform to types of restore operations.

The two types of restores that the **archecker** utility performs are:

- A physical restore that is based on a level-0 archive.
- A physical restore followed by a logical restore, which uses both a level-0 archive and logical logs to restore data to a specific point in time.

When reading the command file, archecker determines whether to perform a physical restore only or a physical restore followed by a logical restore. By default, archecker performs a physical and logical restore. If you use the WITH NO LOG clause, archecker does not perform a logical restore.

The procedures and resources that archecker uses differ between a physical-only restore and a physical and logical restore. These procedures are outlined in the following sections.

Physical restore

When the archecker utility performs a physical restore, the utility extracts data from a level-0 archive.

When performing a physical restore, archecker performs the following tasks:

- Disables all constraints (including foreign constraints that reference the target table), indexes, and triggers until the data is restored. Restore performance is better if the table has no constraints, indexes, or triggers.
- Reads the schema command file to determine the following:
 - The source tables
 - The destination tables
 - The schema of all tables
 - The dbspace names of where tables are located
 - The specific archive to extract data from
- Scans the archive for pages belonging to the tables being restored
- Processes each row from the data page and determines if the row is complete or partial.

If the row is a partial row, then archecker determines if the remaining portion of the row has been staged, and if not, it stages the row for later processing.

- For a physical-only restore, applies filters to the row and rejects rows that are not required.
- Inserts the row into the destination table.

To restore a table with the original schema, the source schema must be specified. To restore a table with a different schema, the table name in the target schema must be different from the table name in the source schema. After restoring by using a different schema, the table can be renamed with the **rename table** statement.

Logical restore

After a physical restore, logical recovery can further restore tables to a user-specified point in time. To do this, the archchecker utility reads backed-up logical logs, converts them to SQL statements, and then replays these statements to restore data.

Before performing a logical recovery, ensure that all transactions you want to restore are contained in backed-up logical logs. The archchecker utility cannot replay transactions from the current log. You cannot perform a logical restore on an external table.

If a table is altered, dropped, or truncated during a logical restore, the restore terminates for that table. Termination occurs at the point that the alter was performed. A message in the archchecker message log file records that an alter operation occurred.

The archchecker utility cannot process compression dictionaries during a logical restore of compressed tables in non-logged databases. A logical restore stops for a table if it finds that a new compression dictionary was created for that table.

When performing a logical restore, archchecker uses two processes that run simultaneously:

Stager

Assembles the logical logs and saves them in tables.

Applier

Converts the log records to SQL statements and executes the statements.

The stager

To collect the pertinent logical log records, the stager performs the following steps:

1. Scans only the backed-up logical logs

The stager reads the backed-up logical log files and assembles complete log records.

2. Tests the logical log records

Any log record that is not applicable to the tables being restored is rejected.

3. Inserts the logical log information in to a table

If the logical log record is not rejected, it is inserted into a stage table.

The applier

The *applier* reads data from the control table created by the stager. It begins processing the required transaction and updates the control table to show that this transaction is in process. Next, it operates on each successive log record, row by row, until the transaction commits.

All updates to the control table occur in the same transaction as the log record modification. This allows all work to be completed or undone as a single unit, maintaining integrity at all times. If an error occurs, the transaction is rolled back and the error is recorded in the control table entry for this transaction.

When data is being restored and the DBA has elected to include a logical restore, two additional work columns and an index are added to the destination table. These columns contain the original rowid and original part number. These columns provide a unique key which identifies the location of the row on the original source archive. To control the storage of the index, use the SET WORKSPACE command (see [The SET statement on page 217](#)). Otherwise, the index is stored in the same space as the table.

After the applier has finished and the restore is complete, these columns, and any indexes created on them, are dropped from the destination table.

Syntax for archecker utility commands

The archecker utility provides a command-line interface for restoring data from an archive. To use archecker, you must specify both a configuration file and a schema command file.

```
archecker {{{ -b | -t } <Table-level restore> [-d] | -D | -i } [-v] [-s] | -V | -version }
```

Table-level restore

```
" [-x { [[ -f cmd_file ] [-1 { phys | stage | apply } ] ] | -D } ]"
```

Table 36. Options for the archecker command

Element	Description
-b	Provides direct XBSA access for backups created with ON-Bar.
-d	Deletes previous archecker restore files, except the archecker message log. For more information, see When to delete restore files on page 212 .
-D	Deletes previous archecker restore files, except the archecker message log, and then exits.

The -D option can be used with the -X option to delete previous restore files plus any table-level-restore working tables in the **sysutils** database. For more information, see [When to delete restore files on page 212](#).

Table 36. Options for the archecker command

(continued)

Element	Description
-f <i>cmdfile</i>	Specifies that archecker use the command file specified by <i>cmdfile</i> . This option overrides the value of the AC_SCHEMA configuration parameter. For more information, see Schema command file on page 205 .
-i	Manually initializes the system.
-	
-lphys,stage,apply	Specifies the level of logical restore: <p>phys</p> <p>Starts a logical restore of the system, but stops after physical recovery is complete. The backed up logical logs must be available.</p> <p>stage</p> <p>After physical recovery is complete, extracts the logical logs from the storage manager and stages them in their corresponding tables, and starts the stager.</p> <p>apply</p> <p>Starts the applier. The applier takes the transactions stored in the stage tables and converts them to SQL and replays the operations.</p> <p>The default level of logical restore if -l is not listed is -lphys,stage,apply. You can specify any combination of the logical restore levels, separated with commas. Spaces are not allowed between -l and levels.</p> <p>For more information, see Manually control a logical restore on page 209.</p>
-s	Prints a status message to the screen.
-t	Specifies ontape as the backup utility.
-v	Specifies verbose mode.
-X	Specifies a table-level restore.
-V	Displays HCL Informix® version information.
-version	Displays additional version information about the build operation system, build number, and build date for HCL Informix®.

When you use ON-Bar, you can use an ON-Bar command to access archecker information to verify a backup. For information on the syntax for this command, see [onbar -v syntax: Verifying backups on page 58](#).

Manually control a logical restore

You can manually control the stager and applier with the -l command-line option.

The following examples show how to perform a logical restore. In all examples, the name of the schema command file is `cmdfile`.

The following example is a typical usage:

```
archecker -bvs -f cmdfile
```

This command is equivalent to the following command:

```
archecker -bvs -f cmdfile -lphys,stage,apply
```

After the physical restore is complete, the `archecker` utility starts the stager. After the stager has started, the applier is automatically started.

In the following example, the `-lphys` option performs a physical-only restore:

```
archecker -bvs -f cmdfile -lphys
```

In the following example, the `-lstage` option starts the `archecker` stager. The stager extracts the logical log records from the storage manager and saves the applicable records to a table.

```
archecker -bvs -f cmdfile -lstage
```

The stager should only be started after physical recovery has completed.

In the following example, the `-lapply` option starts the `archecker` applier. It looks in the `acu_control` table for the transaction to recover. The applier should only be started after the stager has been started.

```
archecker -bvs -f cmdfile -lapply
```

Performing a restore with multiple storage managers

About this task

If you use multiple storage managers, you can perform a table-level restore with `archecker` by configuring `archecker` on every node.


To perform a table-level restore that involves multiple storage managers:

1. Create an `archecker` configuration file on every node.
2. Create a schema command file on every node.
3. Remove old restores by executing the `archecker -DX` command on a single node.
4. Start the physical restore by executing the `archecker -bX -lphys` command on each node.



Restriction: Do not use the `-d` option.

5. After the physical restore completes, start the logical restore by executing the `archecker -bX -lstage` command on each node that contains logical log records.

 **Restriction:** Do not use the -d option.

6. After starting all stagers, complete the restore by executing the `archecker -bX -lapply` command on a single node.

Perform a parallel restore

If you have a fragmented table that resides in separate dbspaces, you can perform a physical table-level restore in parallel by executing multiple `archecker` commands with different schema command files for each dbspace.

During a level-0 archive, there cannot be any open transactions that would change the schema of the table. The table or table fragments being recovered must exist in the level-0 archive. The table or fragment cannot be created or added during the logical recovery. Tables created or fragments added during the logical recovery are ignored.

Because a detached fragment is no longer part of the original table, the applier does not process the detached fragment log record or any other log records for this fragment from this point forward. A message in the `archecker` message log file indicates a detach occurred.

In this example, the table is fragmented across three dbspaces. The corresponding schema command files are named `cmdfile1`, `cmdfile2`, `cmdfile3`. The following commands delete previous restores and then perform physical restores on each dbspace in parallel:

- `archecker -DX`
- `archecker -bvs -f cmdfile1 -lphys`
- `archecker -bvs -f cmdfile2 -lphys`
- `archecker -bvs -f cmdfile3 -lphys`

You cannot perform a logical restore in parallel.

Restore tables with large objects

ON-Bar supports table-level restores of smart large objects and binary large objects.

- Smart large objects

Table-level restore also supports smart large objects for physical restore only (restore from level-0 archive).

The storage location of the smart large object columns being restored must be specified with the `PUT` clause in the `CREATE TABLE` statement. The restored smart large objects are created with the create-time flags `LO_NOLOG` and `LO_NOKEEP_LASTACCESS_TIME`. These flags override the `LOG` and `KEEP ACCESS TIME` column attributes if they are specified in the target table for the smart large object column.

- Binary large objects

Table-level restore supports restoring `tblspace` binary large objects, but not `blobspace` binary large objects. If you attempt to restore a `blobspace` binary large object, the value is set to `NULL` and a warning is issued.

When to delete restore files

If you repeatedly run the same archecker table-level restore, you must clean up the archecker table-level restore working files and tables from the previous runs. These working tables refer to `acu_` tables in the **sysutils** database that are created during an archecker table-level restore. The archecker table-level restore working files and tables are kept after an archecker table-level restore completes in case these files and tables are needed for diagnosing problems.

You can remove the working files and tables by explicitly running the command `archecker -DX` or by using the `-d` option when you run the next archecker table-level restore command. The `-d` option indicates that all files and tables from the previous run of archecker table-level restore are removed before the new restore begins.

- ontape example: `archecker -tdvs -fschema_command_file`
- onbar example: `archecker -bdvs -fschema_command_file`

The archecker schema reference

The topics in this section describe the SQL-like statements used by the archecker schema command file. This file provides information that the archecker utility uses to perform data recovery.

Use the schema command file to specify the source and destination tables and to define the table schema.

For information about specifying which command file archecker uses, see [Schema command file on page 205](#).

The following are statements supported by archecker:

- CREATE TABLE
- DATABASE
- INSERT INTO
- RESTORE
- SET



Important: Standard SQL comments are allowed in the archecker utility file and are ignored during processing.

The syntax of these statements is described in the following topics.

The CREATE TABLE statement

The CREATE TABLE statement describes the schema of the source and target tables. If the target table is external, use the CREATE EXTERNAL TABLE statement described in the section [The CREATE EXTERNAL TABLE statement on page 213](#).

Syntax

The syntax of the CREATE TABLE used in the archecker schema command file is identical to the corresponding HCL® Informix® SQL statement. For a description of this syntax, see the *HCL® Informix® Guide to SQL: Syntax*.

Usage

You must include the schema for the source table in the archecker schema command file. This schema must be identical to the schema of the source table at the time the archive was created.

The schema of the source table is not validated by archecker. Failing to provide an accurate schema leads to unpredictable results.

The source table cannot be a synonym or view. The schema of the source table only needs the column list and storage options. Other attributes such as extent sizes, lock modes, and so on are ignored. For an ON-Bar archive, archecker uses the list of storage spaces for the source table to create its list of objects to retrieve from the storage manager. If the source table is fragmented, you must list all dbspaces that contain data for the source table. The archecker utility only extracts data from the dbspaces listed in the schema command file.

If the source table contains constraints, indexes, or triggers, they are automatically disabled during the restore. Foreign constraints that reference the target table are also disabled. After the restore is complete, the constraints, indexes, and triggers are enabled. For better performance, remove constraints, indexes, and triggers prior to performing a restore.

You must also include the schema of the target table in the command file. If the target table does not exist at the time the restore is performed, it is created using the schema provided.

If the target table exists, its schema must match the schema specified in the command file. Data is then appended to the existing table.

Examples

The schema of the source and target tables do not have to be identical. The following example shows how you can repartition the source data after performing the data extraction:

```
CREATE TABLE source (col1 integer, ...) IN dbspace1;
CREATE TABLE target (col1 integer, ...)
  FRAGMENT BY EXPRESSION
  MOD(col1, 3) = 0 in dbspace3,
  MOD(col1, 3) = 1 in dbspace4,
  MOD(col1, 3) = 2 in dbspace5;
INSERT INTO target SELECT * FROM source;
```

The CREATE EXTERNAL TABLE statement

The CREATE EXTERNAL TABLE statement describes the schema of an external target table.

Syntax

The syntax of the CREATE EXTERNAL TABLE statement for the archecker schema file is not identical to the SQL CREATE EXTERNAL TABLE statement.

```
CREATE EXTERNAL TABLE name ( columndata_type ) USING ( "filename" [ , { DELIMITED | INFORMIX } ] );
```

Element	Description
<i>column</i>	The name of the column. Must conform to SQL identifier syntax rules. For more information, see the <i>HCL® Informix® Guide to SQL: Syntax</i> .
<i>data_type</i>	The built-in data type of the column. For more information about data types, see the <i>HCL® Informix® Guide to SQL: Reference</i> .
<i>filename</i>	Either the name of the file in which to place the data or a pipe device. The pipe device must exist before starting the archecker utility.
<i>name</i>	The name of the table to store the external data. Must be unique among names of tables, views, and synonyms in the current database. Must conform to SQL database object name rules. For more information, see the <i>HCL® Informix® Guide to SQL: Syntax</i> .

Usage

When you use the CREATE EXTERNAL TABLE statement to send data to an external table, the data is only extracted from a level-0 archive. Logical logs are not rolled forward on an external table.

You can specify either of the following formats for external files:

- DELIMITED: ASCII delimited file. This is the default format.
- INFORMIX: internal binary representation. To optimize performance, filters are not applied to external tables. If filters exist, a warning indicates that they are ignored.

For an example of using the CREATE EXTERNAL TABLE statement, see [Restore to an external table on page 219](#).

The DATABASE statement

In the **archecker** utility, the DATABASE statement sets the current database.

Syntax

```
DATABASE dbname [LOG MODE ANSI] ;
```

Element	Description
<i>dbname</i>	The name of the current database.

Usage

Multiple DATABASE statements can be used. All table names referenced following this statement are associated with the current database.

If the logging mode of the source database is ANSI and default decimal columns are used in the table schemas, then the logging mode of the database must be declared.

If the logging mode of the source database is not declared no error will be returned, but unexpected results and data can occur.

Examples

In the following example, both the source and target tables reside in the same database **db**.

```
DATABASE db;
CREATE TABLE source (...);
CREATE TABLE target (...);
INSERT INTO target SELECT * from source;
```

You can use multiple database statements to extract a table from one database into another database.

```
DATABASE db1;
CREATE TABLE source (...) IN dbspace1;
DATABASE db2;
CREATE TABLE target (...) IN dbspace2;
INSERT INTO db2:target SELECT * FROM db1:source;
```

The INSERT statement

The INSERT statement tells the archecker utility what tables to extract and where to place the extracted data.

Syntax

```
INSERT INTO target_table [ ( target_column ) ] SELECT { src_column | * } FROM src_table [ WHERE filter ]
```

Element	Description
<i>filter</i>	<p>The following filters are supported by the INSERT statement:</p> <ul style="list-style-type: none"> • =, !=, <> • >, >=, <, <= • [NOT] MATCHES, [NOT] LIKE • IS [NOT] NULL • AND, OR • TODAY, CURRENT <p>The following operators are not supported by the archecker utility:</p> <ul style="list-style-type: none"> • Aggregates • Functions and procedures • Subscripts • Subqueries • Views • Joins <p>Filters can only be applied to physical-only restore.</p>

Element	Description
<i>src_column</i>	A list of columns to be extracted.
<i>src_table</i>	The source table on the archive where the data is restored from.
<i>target_column</i>	The destination column or columns where the data will be restored.
<i>target_table</i>	The destination table where the data will be restored.

Examples

The following example demonstrates the simplest form of the INSERT statement. This statement extracts all rows and columns from the source to the target table.

```
INSERT INTO target SELECT * FROM source;
```

You can also extract a subset of columns. In the following example, only two columns from the source table are inserted into the destination table.

```
CREATE TABLE source (col1 integer, col2 integer, col3 integer, col4 integer);
CREATE TABLE target (col1 integer, col2 integer);
INSERT INTO target (col1, col2) SELECT col3, col4 FROM source;
```

The RESTORE statement

The RESTORE statement is an optional command to restore tables to a specific point in time.

Syntax

```
RESTORE [ TO { "time" | CURRENT } [ WITH NO LOG ] ]
```

Element	Description
"time"	The date and time the table is to be restored to.

Usage

The TO clause is used to restore the table to a specific point in time, which is specified by a date and time or the reserved word CURRENT.

Only one RESTORE statement can be specified in a command file. If this statement is not present in the command file, then the system will be restored to the most current time using logical logs.

If the WITH NO LOG clause is present, only a physical restore is performed. In addition, the two extra columns and the index are not added to the destination table. Physical-only restores are based on level-0 archives only.



Tip: Use this option when you do not have logical logs. You will not receive any messages about logical recovery.

Example

```
RESTORE TO CURRENT WITH NO LOG;
```


The SET statement

The SET statement controls the different features in the table-level unload library.

Syntax

```
SET { COMMIT TO number | WORKSPACE TO { / dbspace } }
```

Element	Description
<i>number</i>	Sets the number of records to insert before committing during a physical restore. The default is 1000.
<i>dbspace</i>	The dbspaces to use for the working storage space. The default is the root dbspace. You cannot use temporary dbspaces for the working storage space.

The archecker utility creates several tables for the staging of logical log records during a logical restore. These tables are created in the **sysutils** database and stored in the working storage space.

Examples

```
SET COMMIT TO 20000;
SET WORKSPACE to dbspace1;
```

Schema command file examples

This section contains examples that show different command file syntax for different data recovery scenarios.

Simple schema command file

The schema command file in this example extracts a table from the most recent level-0 backup of **dbspace1**. The data is placed in the **table test1:tlr** and the logs are applied to bring the table **tlr** to the current point in time.

```
database test1;
create table tlr (
  a_serial serial,
  b_integer integer,
  c_char char,
  d_decimal decimal
) in dbspace1;
insert into tlr select * from tlr;
```

Restore a table from a previous backup

The schema command file in this example extracts a table from the level-0 backup of **dbspace1**. The logical logs are used to bring the table to the time of "2003-01-01 01:01:01". The data is placed in the table **test1:tlr**.

```
database test1;
create table tlr (
  a_serial serial,
  b_integer integer,
  c_char char,
  d_decimal decimal
```

```

) in dbspace1;
insert into tlr select * from tlr;
restore to '2003-01-01 01:01:01';

```

Restore to a different table

The schema command file in this example extracts a table called **test1:tlr** from the most recent backup of **dbspace1** and places the data in the table **test1:tlr_dest**.

```

database test1;
create table tlr (
  a_serial serial,
  b_integer integer,
  c_char char(20),
  d_decimal decimal,
) in dbspace1;
create table tlr_dest (
  a_serial serial,
  b_integer integer,
  c_char char(20),
  d_decimal decimal
) in dbspace2;
insert into tlr_dest select * from tlr;

```

Extract a subset of columns

The schema command file in this example extracts a table **test1:tlr** from the most recent backup of **dbspace1** and places a subset of the data into the table **test1:new_dest**


```

database test1;
create table tlr (
  a_serial serial,
  b_integer integer,
  c_char char(20),
  d_decimal decimal
) in dbspace1;
create table new_dest (
  X_char char(20),
  Y_decimal decimal,
  Z_name char(40)
) in dbspace2;
insert into new_dest (X_char, Y_decimal) select c_char,d_decimal from tlr;

```

Use data filtering

The schema command file in this example extracts a table **test1:tlr** from the most recent backup of **dbspace1** and places the data in the table **test1:tlr** only where the list conditions are true.

 **Important:** Filters can only be applied to a physical restore.

```

database test1;
create table tlr (
  a_serial serial,
  b_integer integer,

```

```

c_char      char(20),
d_decimal   decimal,
) in dbpace1;
insert into tlr
select * from tlr
where c_char matches 'john*'
and d_decimal is NOT NULL
and b_integer > 100;
restore to current with no log;

```

Restore to an external table

The schema command file in this example extracts a table called **test1:tlr** from the most recent backup of **dbpace1** and places the data in a file called `/tmp/tlr.unl`.

```

database test1;
create table tlr
(a_serial serial,
 b_integer integer
) in dbpace1;
create external table tlr_dest
(a_serial serial,
 b_integer integer
) using ("/tmp/tlr.unl", delimited );
insert into tlr_dest select * from tlr;
restore to current with no log;

```

Restore multiple tables

The schema command file in this example extracts a table **test1:tlr_1** and **test1:tlr_2** from the most recent backup of **dbpace1** and places the data in **test1:tlr_1_dest** and **test1:tlr_2_dest**. This is an efficient way of restoring multiple tables because it requires only one scan of the archive and logical log files.

```

database test1;
create table tlr_1
( columns ) in dbpace1;
create table tlr_1_dest ( columns );
create table tlr_2
( columns ) in dbpace1;
create table tlr_2_dest ( columns );
insert into tlr_1_dest select * from tlr_1;
insert into tlr_2_dest select * from tlr_2;

```

Perform a distributed restore

The schema command file in this example extracts a table **test:tlr_1** from the most recent backup of **dbpace1** and places the data on the database server **rem_srv** in the table **rem_dbs:tlr_1**.

```

database rem_dbs
create table tlr_1
( columns );
database test1;
create table tlr_1
( columns ) in dbpace1;

```

```
insert into rem_dbs@rem_srv.tlr_1
select * from tlr_1;
```

Backup and restore configuration parameter reference

Backup and restore configuration parameters

These topics describe the configuration parameters that you use with the ON-Bar, ontape, and archecker utilities.

You set most of these configuration parameters in the `onconfig` file. However, you set some of the archecker configuration parameters in the **AC_CONFIG** file.

Be sure to configure your storage manager. Depending on the storage manager that you choose, you might set different ON-Bar configuration parameters. If you are using a third-party storage manager, see [Configuring a third-party storage manager on page 36](#), before you start ON-Bar.

The following table describes the following attributes (if relevant) for each parameter.

Attribute	Description
<code>ac_config</code> . <i>std value</i>	For archecker configuration variables. The default value that appears in the <code>ac_config.std</code> file.
<code>onconfig</code> . <i>std value</i>	For <code>onconfig</code> configuration variables. The default value that appears in the <code>onconfig.std</code> file.
<i>if value not present</i>	The value that the database server supplies if the parameter is missing from your <code>onconfig</code> file. If this value is present in <code>onconfig.std</code> , the database server uses the <code>onconfig.std</code> value. If this value is not present in <code>onconfig.std</code> , the database server uses this value.
<i>units</i>	The units in which the parameter is expressed
<i>range of values</i>	The valid values for this parameter
<i>takes effect</i>	The time at which a change to the value of the parameter affects ON-Bar operation. Except where indicated, you can change the parameter value between a backup and a restore.
<i>refer to</i>	Cross-reference to further discussion

Related reference

[ON-Bar backup and restore system on page 20](#)

ON-Bar and ontape configuration parameters and environment variable

Many properties of the ON-Bar and ontape utilities are controlled by configuration parameters in the `onconfig` file. ON-Bar also has an environment variable.

! **Important:** ON-Bar does not use the `TAPEDEV`, `TAPEBLK`, `TAPESIZE`, `LTAPEBLK`, and `LTAPESIZE` configuration parameters. ON-Bar checks if `LTAPEDEV` is set to `/dev/null` on UNIX™ or `NUL` on Windows™.

Related information

[Configuring Informix Primary Storage Manager on page 191](#)

ARCHIVE_UNLOGGED_OBJECTS configuration parameter

`ARCHIVE_UNLOGGED_OBJECTS` configuration parameter specifies if a backup is allowed on an RS secondary server if non-logging objects exist in the primary database server.

onconfig.std value

0 – The default value which will result in the legacy behavior. The archive on the RSS will be blocked if the instance contains non-logged objects.

1 – Quiet mode. Non-logged objects will be searched. If found, a general message will be printed in the log and archiving will continue.

2 – Verbose mode. A separate descriptive message will be printed in the log identifying every non-logged object found, then the archiving will continue.

takes effect

After you edit your `onconfig` file of the RSS server and restart the database server.

Usage

By default, a backup on an RSS server will fail if non-logging objects (i.e. RAW tables) exist in the primary database server. If you want the backup to complete even if non-logging objects exist, you can set this configuration parameter to 1 or 2. While restoring these backups using `onbar/ontape`, a warning message will be printed in the `online.log` if they were taken with this configuration parameter set to 1 or 2. The message will read: *Warning: This backup is incomplete, the source database instance contains RAW/non-log objects that are missing from this backup, restore it at your own risk.*

BACKUP_FILTER configuration parameter

Use the `BACKUP_FILTER` configuration parameter to specify the path name and any options for an external filter program that you use with the ON-Bar or `ontape` utility.

onconfig.std value

Not set. Backup data is not filtered.

values

The path name of a command and any options. By default, the path name is relative to the `$INFORMIXDIR/bin` directory, otherwise, the path name must be the absolute path of the program. If you include command-line options, both the filter name and the options must be surrounded by single quotation marks.

takes effect

After you edit your `onconfig` file and ON-Bar or **ontape** starts.

Usage

This filter transforms data before backing it up, such as compressing it. The transformed data is then backed up and stored as a single file. When you perform a restore, you must transform the data back to its original format. Specify the appropriate program to transform data before a restore by setting the `RESTORE_FILTER` configuration parameter.

For security purposes, filters should not have write permission to non-privileged users. Permission on the filters is the same as that of permission on other executable files that are called by the HCL Informix® server or utilities.



Note: If the `BACKUP_FILTER` parameter is set in the `onconfig` file, the `LTAPESIZE` configuration parameter cannot be set to 0. Otherwise the ON-Bar or **ontape** utility returns an error when backing up logical logs to a directory on disk. The error message is:

```
The LTAPESIZE configuration parameter cannot be set to 0 when the BACKUP_FILTER
configuration parameter is set; change the value of LTAPESIZE.
Program over.
```

A workaround is to set the `LTAPESIZE` configuration parameter to a high value. Log files are not much higher than the `LOGSIZE` configuration parameter. Use the value in the `LOGSIZE` as the upper limit for this database.

When you specify filter information in the `BACKUP_FILTER` configuration parameter, specify the path name of a filter program, and any options, as shown in this example:

```
BACKUP_FILTER /bin/compress
```

Output produced by this filter is saved as a single object in the storage manager.

The `BACKUP_FILTER` configuration parameter can include command-line options as well as the filter name. For example, specify:

```
BACKUP_FILTER 'my_encrypt -file /var/adm/encryption.pass'
```

Related reference

[RESTORE_FILTER configuration parameter on page 249](#)

BAR_ACT_LOG configuration parameter

Use the `BAR_ACT_LOG` configuration parameter to specify the full path name of the ON-Bar activity log.

onconfig.std value

UNIX™: BAR_ACT_LOG \$INFORMIXDIR/tmp/bar_act.log

Windows™: BAR_ACT_LOG %INFORMIXDIR%\tmp\bar_act.log

range of values

Full path name

takes effect

When onbar-driver starts

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` or equivalent SQL administration API command.

Usage

You should specify a path to an existing directory with an appropriate amount of space available or use `$INFORMIXDIR/bar_act.log`.

Whenever a backup or restore activity or error occurs, ON-Bar writes a brief description to the activity log. The format of the file resembles the format of the database server message log. You can examine the activity log to determine the results of ON-Bar actions.

The file specified by the `BAR_ACT_LOG` configuration parameter is created if it does not exist. If the ON-Bar command (or any ON-Bar-related utility such as the `onsmsync` utility) never ran on the system, then the file does not exist.

The `sysbaract_log` table is a system monitoring interface pseudo table that reads data from the file specified by `BAR_ACT_LOG`. The following errors are returned if you attempt to query the `sysbaract_log` on a system where the `BAR_ACT_LOG` file does not exist:

```
244: Could not do a physical-order read to fetch next row.
101: ISAM error: file is not open.
```

Usage when you specify a file name only

If you specify a file name only in the `BAR_ACT_LOG` configuration parameter, ON-Bar creates the ON-Bar activity log in the working directory in which you started ON-Bar. For example, if you started ON-Bar from `/usr/mydata` on UNIX™, the activity log is written to that directory.

For UNIX™, if the database server launches a continuous logical-log backup, ON-Bar writes to the ON-Bar activity log in the working directory for the database server.

For Windows™, if the database server launches a continuous logical-log backup, ON-Bar writes to the activity log in the `%INFORMIXDIR%\bin` directory instead.

Related reference

[View ON-Bar backup and restore performance statistics on page 114](#)

[PSM_ACT_LOG configuration parameter on page 261](#)

Related information

[bar_act.log file: ON-Bar activity log on page 23](#)

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

BAR_BSALIB_PATH configuration parameter

Use the BAR_BSALIB_PATH configuration parameter to specify the path name and file name of the XBSA shared library for the storage manager that you use.

default value

UNIX™: \$INFORMIXDIR/lib/ibsad001.*extension*

Windows™: %INFORMIXDIR%\lib\ibsad001.*extension*

The *extension* is platform specific.

onconfig.std value

Not set.

takes effect

When onbar-driver starts

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` or equivalent SQL administration API command.

Usage

ON-Bar and the storage manager rely on a shared library to integrate with each other. Configure the BAR_BSALIB_PATH configuration parameter for your storage-manager library. Support for BAR_BSALIB_PATH is platform-specific. Check your machine notes to determine whether you can use this configuration parameter with your operating system. You can change the value of BAR_BSALIB_PATH between a backup and restore.

To ensure that this integration takes place, specify the shared-library path name. Set one of the following options:

UNIX™:

- Place the storage-manager library in the default directory.

For example, the suffix for Solaris is `so`, so you specify `$INFORMIXDIR/lib/ibsad001.so` on a Solaris system.

For example, the suffix for Solaris is `so`, so you specify `$INFORMIXDIR/lib/libbsapsm.so` on a Solaris system.

- Place the storage-manager library in any directory and create a symbolic link from `$INFORMIXDIR/lib/ibsad001.platform_extension` to it.

If you use IBM® Informix® Storage Manager (ISM), create a symbolic link to `$INFORMIXDIR/lib/libbsa.platform_extension` or set `BAR_BSALIB_PATH` to this absolute path value.

If you use IBM® Spectrum Protect, create a symbolic link to `$INFORMIXDIR/lib/libtxbsa.platform_extension` or set `BAR_BSALIB_PATH` to this absolute path value.

- Set the **LD_LIBRARY_PATH** environment variable. For example, set **LD_LIBRARY_PATH** to `$INFORMIXDIR/lib`.

Windows™:

- Place the storage-manager library in the default directory.
- Set the path name for the ISM shared library to the installation directory for ISM: `%ISMDIR%\bin\libbsa.dll`.

The `%ismDIR%` variable includes a version or release number. For example: set `ISMDIR=C:\program files\informix\ism\2.20`. This directory is set when the database server is installed on Windows™. This path name is different if you use a different storage manager.

If the parameter `BAR_BSALIB_PATH` is missing or has no value and the database server cannot open the XBSA shared library for your platform, ON-BAR tries to use the HCL Informix® Primary Storage Manager as the storage manager in all platforms.



Tip: Be sure that the shared library can access the backup data in the storage manager in a restore. You cannot back up on to one storage manager and restore from a different storage manager.

Related information

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

BAR_CKPTSEC_TIMEOUT configuration parameter

The `BAR_CKPTSEC_TIMEOUT` configuration parameter specifies the amount of time, in seconds, that an RS secondary server should wait for a checkpoint to arrive from the primary server while performing an external backup.

onconfig.std value

UNIX™: 15

Windows™: 16

15

default value

UNIX™: 15

Windows™: 16

15

units

seconds

range of values

5 through twice the value of the CKPTINTVL configuration parameter

takes effectAfter you edit your `onconfig` file and restart the database server.When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` command.When you reset the value in memory by running the `onmode -wm` command.**Usage**

When an external backup is performed on an RS secondary server, the secondary server must wait until a checkpoint arrives in the logical logs from the primary server. A checkpoint flushes buffers to disk, and blocks user transactions that involve temporary tables. If the checkpoint on the primary does not complete in the time-out period, the backup on the RS secondary server fails. You can set the `BAR_CKPTSEC_TIMEOUT` configuration parameter to a longer amount of time, in seconds, that an RS secondary server should wait for a checkpoint to arrive from the primary server while performing an external backup.

Related information

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

[CKPTINTVL configuration parameter on page](#)

BAR_DEBUG configuration parameter

Use the `BAR_DEBUG` configuration parameter to specify the amount of debugging information that the database server captures in the ON-Bar activity log.

onconfig.std value`BAR_DEBUG 0`**values**

0 = Do not display debugging information.

1 = Print a small amount of information

2 = Print a message every time ON-Bar:

- Enters a function.
- Exits a function. The message includes the return code for the function.

3 = Print exit and entry information with additional details.

4 = Also print information about ON-Bar parallel operations.

5 = Also print information about:

- Objects that are being backed up or restored.
- The `act_node` structure corresponding with the `bar_action` table.

6 = Print additional information about:

- Objects that are being backed up or restored.
- The `act_node` structure corresponding with the `bar_action` table.

7 = Also print:

- Information about the contents of the `ins_node` structure corresponding with the `bar_instance` table.
- Information about modifications to the `bar_action` table.
- Information about logical logs and objects that are restored.
- SQL statements done on the **sysutils** database and SQLCODES that were returned.

8 = Also print page headers of all pages archived and restored. This setting requires a large amount of space.

9 = Print the contents of:

- The `bar_ins` structure after it was initialized.
- The object descriptors that are cold restored.

takes effect

Immediately after you edit your `onconfig` file for any currently executing ON-Bar command and any subsequent commands. Any ON-Bar command that is currently executing when you update `BAR_DEBUG` reads the new value of `BAR_DEBUG` and prints debug messages at the new level.

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` or equivalent SQL administration API command.

Usage

The default value of 0 displays no debugging information. Set the `BAR_DEBUG` configuration parameter to a higher value to display more detailed debugging information in the ON-Bar activity log.

You can dynamically update the value of `BAR_DEBUG` in the `onconfig` file during a session.

Related reference

[BAR_DEBUG_LOG configuration parameter on page 230](#)

Related information

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

BAR_ENCRYPTION configuration parameter

Use the BAR_ENCRYPTION parameter to encrypt the backups.

onconfig.std value

Not set. Backup data is not encrypted.

takes effect

When ON-Bar or ontape starts.

Usage

Use the BAR_ENCRYPTION configuration parameter to enable backup encryption, set the path to the keystore containing the credentials to access a remote key server or the path of a keyfile containing the backup encryption key, and specify the encryption cipher.

```
Syntax for the BAR_ENCRYPTION configuration parameter

>>BAR_ENCRYPTION--keystore-----keystore_name----->
      '--keyfile-----keyfile_name--'
>----->
  '-,--cipher-----+--aes128--+'
      +-aes192+-
      '-aes256-'
```

Table 37. Options for the BAR_ENCRYPTION configuration parameter value

Field	Value
keystore	<p>The <i>keystore</i> specifies the name of the keystore and stash file names. The files are created in the <code>INFORMIXDIR/etc</code> directory:</p> <ul style="list-style-type: none"> • <i>keystore.p12</i> = The keystore file that contains the security certificates. • <i>keystore.sth</i> = The stash file that contains the encryption password. <p>You must manually back up the keystore and password stash files. These files are not backed up when you run a back up with the ON-Bar or ontape utilities.</p>

Table 37. Options for the BAR_ENCRYPTION configuration parameter value (continued)

Field	Value
keyfile	The <i>keyfile</i> specifies the full path of a text file that contains a backup encryption key of suitable size for the cipher chosen. The key must be encoded in base64 format.
cipher	Specifies the encryption cipher: <ul style="list-style-type: none"> • aes128 = Default. Advanced Encryption Standard cipher with 128-bit keys. • aes192 = Advanced Encryption Standard cipher with 192-bit keys. • aes256 = Advanced Encryption Standard cipher with 256-bit keys.

BAR_DECRYPTION configuration parameter

Use the BAR_DECRYPTION parameter to decrypt the backups.

onconfig.std value

Not set. If BACKUP_DECRYPTION is not set, backup data that will not be decrypted.

takes effect

When ON-Bar, ontape, onlog and archecker starts.

Usage

Use the BAR_DECRYPTION to override the BAR_ENCRYPTION configuration parameter during a restore. Like with BAR_ENCRYPTION, you can set the path to a keystore or to a keyfile. You cannot set the cipher in this parameter as the cipher is set by the object to be restored.

Syntax for the BAR_DECRYPTION configuration parameter

```
>>BAR_DECRYPTION--keystore----keystore_name----->
      '--keyfile-----keyfile_name--'
```

Table 38. Option for the BAR_DECRYPTION configuration parameter value

Field	Value
keystore	The <i>keystore</i> specifies the name of the keystore and stash file names. The files are created in the <code>INFORMIXDIR/etc</code> directory:

Table 38. Option for the BAR_DECRYPTION configuration parameter value (continued)

Field	Value
	<ul style="list-style-type: none"> • <i>keystore.p12</i> = The keystore file that contains the security certificates. • <i>keystore.sth</i> = The stash file that contains the decryption password. <p>You must manually back up the keystore and password stash files. These files are not backed up when you run a back up with the ON-Bar or ontape utilities.</p>
keyfile	The <i>keyfile</i> specifies the full path of a text file that contains a backup decryption key. The key must be encoded in base64 format.

BAR_DEBUG_LOG configuration parameter

Use the BAR_DEBUG_LOG parameter to specify the location and name of the ON-Bar debug log.

onconfig.std value

```
/usr/informix/bar_dbug.log
```

if value not present

```
UNIX™: /tmp/bar_dbug.log
```

```
Windows™: \tmp\bar_dbug.log
```

takes effect

When ON-Bar starts

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` or equivalent SQL administration API command.

Usage

For security reasons, set the BAR_DEBUG_LOG configuration parameter to a directory with restricted permissions, such as the `$INFORMIXDIR` directory.

Related reference

[BAR_DEBUG configuration parameter on page 226](#)

[PSM_DEBUG_LOG configuration parameter on page 265](#)

Related information

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

BAR_HISTORY configuration parameter

Use the BAR_HISTORY configuration parameter to specify whether the **sysutils** database maintains a backup history when you use onsmsync to expire old backups.

onconfig.std value

Not in the `onconfig.std` file.

default value

0

range of values

0 = Remove records for expired backup objects from the **sysutils** database

1 = Keep records for expired backup objects in the **sysutils** database

takes effect

When onsmsync starts

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` or equivalent SQL administration API command.

Usage

If you set the value to 0, onsmsync removes the **bar_object**, **bar_action**, and **bar_instance** rows for the expired backup objects from the **sysutils** database. If you set the value to 1, onsmsync sets the **act_type** value to 7 in the **bar_action** row and keeps the **bar_action** and **bar_instance** rows for expired backup objects in the **sysutils** database. If you do not set BAR_HISTORY to 1, the restore history is removed.

Regardless of the value of BAR_HISTORY, onsmsync removes the line that describes the backup object from the emergency boot file and removes the object from the storage manager when the storage manager expires the object.

For more information about onsmsync, see [The onsmsync utility on page 107](#).

Related information

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

BAR_IXBAR_PATH configuration parameter

Use the BAR_IXBAR_PATH configuration parameter to change the path and name of the ON-Bar boot file.

onconfig.std value

Not set in `onconfig.std`.

default value

UNIX™ or Linux™: `$INFORMIXDIR/etc/ixbar.servernum`

Windows™: `%INFORMIXDIR%\etc\ixbar.servernum`

range of values

Full path name for the ON-Bar boot file

takes effect

When ON-Bar or `onsmsync` starts

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` or equivalent SQL administration API command.

Usage

By default, the ON-Bar boot file is created in the `%INFORMIXDIR%\etc` folder on Windows™ and in the `$INFORMIXDIR/etc` folder on UNIX™ or Linux™. The default name for this file is `ixbar.servernum`, where `servernum` is the value of the `SERVERNUM` configuration parameter.

For example, in an instance with the `SERVERNUM` configuration parameter equal to 41, the ON-Bar boot file is created by default with this path and name in UNIX™:

```
BAR_IXBAR_PATH $INFORMIXDIR/etc/ixbarboot.41
```

You can change the path to create the file in another location. For example, if you want to create the ON-Bar boot file in the directory `/usr/informix` with the name `ixbar.new`, specify:

```
BAR_IXBAR_PATH=/usr/informix/ixbar.new
```

Related information

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

BAR_MAX_BACKUP configuration parameter

Use the `BAR_MAX_BACKUP` parameter to specify the maximum number of parallel processes that are allowed for each ON-Bar command.

onconfig.std value

0

if value not present

4

units

ON-Bar processes

values

0 = Maximum number of processes allowed on system

1 = Serial backup or restore

n = Specified number of processes created

takes effect

When ON-Bar starts

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` or equivalent SQL administration API command.

Although the database server default value for `BAR_MAX_BACKUP` is 4, the `onconfig.std` value is 0.

Both UNIX™ and Windows™ support parallel backups.

Specify serial backups and restores

To perform a serial backup or restore, including a serial whole system backup or restore, set `BAR_MAX_BACKUP` to 1.

Specify parallel backups and restores

To specify parallel backups and restores, including parallel whole system backups and restores, set `BAR_MAX_BACKUP` to a value higher than 1. For example, if you set `BAR_MAX_BACKUP` to 5 and execute an ON-Bar command, the maximum number of processes that ON-Bar creates concurrently is 5. Configure `BAR_MAX_BACKUP` to any number up to the maximum number of storage devices or the maximum number of streams available for physical backups and restores. ON-Bar groups the dbspaces by size for efficient use of parallel resources.

If you set `BAR_MAX_BACKUP` to 0, the system creates as many ON-Bar processes as needed. The number of ON-Bar processes is limited only by the number of storage spaces or the amount of memory available to the database server, whichever is less.

The amount of memory available is based on `SHMTOTAL`. ON-Bar performs the following calculation where *N* is the maximum number of ON-Bar processes that are allowed:

```
N = SHMTOTAL / (# transport buffers * size of transport buffers / 1024)
```

If `SHMTOTAL` is 0, `BAR_MAX_BACKUP` is reset to 1. If *N* is greater than `BAR_MAX_BACKUP`, ON-Bar uses the `BAR_MAX_BACKUP` value. Otherwise, ON-Bar starts *N* backup or restore processes.

Related information

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

BAR_MAX_RESTORE configuration parameter

Use the BAR_MAX_RESTORE parameter to specify the maximum number of parallel restore processes that are allowed during an ON-Bar restore operation.

onconfig.std value

0

if value not present

The value of the BAR_MAX_BACKUP configuration parameter

units

ON-Bar processes

values

0 = Maximum number of restore processes allowed on system

1 = Serial restore

n = Specified number of restore processes created

takes effect

When ON-Bar starts

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` or equivalent SQL administration API command.

Both UNIX™ and Windows™ support parallel restores.

Specify serial restores

To perform a serial restore, including a serial whole system restore, set BAR_MAX_RESTORE to 1.

Specify parallel restores

To specify parallel restores, including parallel whole system restores, set BAR_MAX_RESTORE to a value higher than 1. For example, if you set BAR_MAX_RESTORE to 5 and start a restore, the maximum number of restore processes that ON-Bar creates concurrently is 5. Configure BAR_MAX_RESTORE to any number up to the maximum number of storage devices or the maximum number of streams available for physical restores. ON-Bar groups the dbspaces by size for efficient use of parallel resources.

If you set BAR_MAX_RESTORE to 0, the system creates as many ON-Bar restore processes as needed. The number of restore processes is limited only by the number of storage spaces or the amount of memory available to the database server, whichever is less.

The amount of memory available is based on SHMTOTAL. ON-Bar performs the following calculation where N is the maximum number of ON-Bar processes that are allowed:

```
N = SHMTOTAL / (# transport buffers * size of transport buffers / 1024)
```

If SHMTOTAL is 0, BAR_MAX_RESTORE is reset to 1. If N is greater than BAR_MAX_RESTORE, ON-Bar uses the BAR_MAX_RESTORE value. Otherwise, ON-Bar starts N restore processes.

Related information

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

BAR_NB_XPORT_COUNT configuration parameter

Use the BAR_NB_XPORT_COUNT configuration parameter to specify the number of data buffers that each onbar_d process can use to exchange data with the database server.

onconfig.std value

20

if value not present

20

units

Buffers

range of values

3 to unlimited

takes effect

When ON-Bar starts

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` or equivalent SQL administration API command.

The value of this parameter affects ON-Bar performance. For example, if you set BAR_NB_XPORT_COUNT to 5 and then issue five ON-Bar commands, the resulting 25 ON-Bar processes use a total of 125 buffers.

To calculate the amount of memory that each onbar_d process requires, use the following formula. For information about the page size for your system, see the release notes:

```
required_memory = (BAR_NB_XPORT_COUNT * BAR_XFER_BUF_SIZE
                  * page_size) + 5 MB
```

Related information

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

BAR_PERFORMANCE configuration parameter

Use the BAR_PERFORMANCE configuration parameter to specify the type of performance statistics to report to the ON-Bar activity log for backup and restore operations.

onconfig.std value

0

units

Levels of statistics

values

0 = Does not collect performance statistics

1 = Reports time spent transferring data between the database server and the storage manager.

2 = Reports ON-Bar processing performance, in microseconds, in the timestamps in the activity log and the error log

3 = Reports both microsecond timestamps and transfer statistics.

takes effect

When ON-Bar starts

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` or equivalent SQL administration API command.

Usage

For example, if you set BAR_PERFORMANCE to 3, ON-Bar reports the time spent transferring data between the HCL Informix® instance and the storage manager, in the activity log. If you set BAR_PERFORMANCE to 0 or do not set it, ON-Bar does not report performance statistics.

- To turn performance monitoring off, set the value to 0. This is the default.
- To display the time spent transferring data between the Informix® instance and the storage manager, set the parameter to 1.
- To display timestamps in microseconds, set the parameter to 2.
- To display both timestamps and transfer statistics, set the parameter to 3.

Related reference

[View ON-Bar backup and restore performance statistics on page 114](#)

Related information

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

BAR_PROGRESS_FREQ configuration parameter

Use the BAR_PROGRESS_FREQ configuration parameter to specify, in minutes, the frequency of the progress messages in the ON-Bar activity log for backup and restore operations.

onconfig.std value

0

if value not present

0

units

minutes

range of values

0, then 5 to unlimited

takes effect

When ON-Bar starts

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` or equivalent SQL administration API command.

Usage

Example: If you set BAR_PROGRESS_FREQ to 5, ON-Bar reports the percentage of the object backed up or restored every 5 minutes. If you set BAR_PROGRESS_FREQ to 0 or do not set it, ON-Bar does not write any progress messages to the activity log.

Specify a value 5 minutes or over. Do not set BAR_PROGRESS_FREQ to 1, 2, 3, or 4, ON-Bar automatically resets it to 5 to prevent overflow in the ON-Bar activity log.

If ON-Bar cannot determine the size of the backup or restore object, it reports the number of transfer buffers sent to the database server instead of the percentage of the object backed up or restored.

Related information

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

BAR_RETRY configuration parameter

Use the BAR_RETRY configuration parameter to specify how many times onbar should try a data backup, logical-log backup, or restore operation if the first attempt fails.

onconfig.std value

1

if value not present

1

units

integer

range of values

0 = BAR_ABORT, stop the rest of the backup/restore

1 = BAR_CONT, continue the rest of the backup/restore

n = 2 to 32766

takes effect

When ON-Bar starts

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` or equivalent SQL administration API command.

Usage

The setting of the `BAR_RETRY` parameter determines ON-Bar behavior in the following ways:

- If set to 0 (`BAR_ABORT`), ON-Bar stops the backup or restore session when an error occurs for a storage space or logical log, returns an error, and quits. If ON-Bar is running in parallel, the already running processes finish but no new ones are started.
- If set to 1 (`BAR_CONT`), ON-Bar stops the backup or restore attempt for that particular storage space, returns an error, and attempts to back up or restore any storage spaces or logical logs that remain.
- If set to a specific number (retry backup and restore operations 2 to 32766 times), ON-Bar attempts to back up or restore this storage space or logical log the specified number of times before it gives up and moves on to the next one.

Related reference

[onbar -RESTART syntax: Restarting a failed restore on page 87](#)

[Resolve a failed restore on page 90](#)

Related information

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

BAR_SEC_ALLOW_BACKUP configuration parameter

Use the `BAR_SEC_ALLOW_BACKUP` configuration parameter on an RSS node to enable the taking of archives and log backups on that node using either `onbar` or `ontape`.

onconfig.std value

0

default value

0

range of values

0 = Do not allow archives or log backups to be taken on this RSS node.

1 = Allow archives and log backups to be taken on this RSS node.

takes effect

When the RSS starts

Usage

The setting of the BAR_SEC_ALLOW_BACKUP parameter on an RSS node determines whether the ontape or onbar utilities may take archives or log backups on that node.

- If set to 0, ontape and onbar will return an error if one attempts to archive spaces or back up logical logs.
- If set to 1 and the RSS node has been properly configured, either ontape or onbar may be used to take archives and log backups locally.

The value of BAR_SEC_ALLOW_BACKUP is ignored on any non-RSS node.

Related information

[Backup and restore a Remote Secondary Server\(RSS\) on page 179](#)

BAR_SIZE_FACTOR configuration parameter

Use the BAR_SIZE_FACTOR configuration parameter to augment the estimate for the size of a backup object, before the backup.

onconfig.std value

Not in onconfig.std.

default value

0

range of values

0 = The estimated size of the backup is not augmented.

Positive integers = The percentage of the original backup size.

takes effect

When the database server starts

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` or equivalent SQL administration API command.

Usage

The estimate is handled before the backup and is calculated so that the storage manager can allocate the storage media appropriately. Because the backup is done online, the number of pages to back up can change during the backup. Some storage managers are strict and if the backup estimate is too low, the backup results in an error.

The value of `BAR_SIZE_FACTOR` is taken as percentage of the original backup object size, and then added to the estimate, before communicating it to the storage manager. `BAR_SIZE_FACTOR` is used only for dbspace backup objects, not for logical log backup objects.

The formula used for calculating the new estimated backup object size is:

$$\text{new_estimate} = \text{original_estimate} \times (1 + (\text{BAR_SIZE_FACTOR} / 100))$$

The value to which you set this parameter in a specific server environment depends on the activity on the system during backup or archive. Therefore, determining the value needs to be based on the individual experience with that system.

Related information

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

BAR_XFER_BUF_SIZE configuration parameter

Use the `BAR_XFER_BUF_SIZE` configuration parameter to specify the size of each transfer buffer.

onconfig.std value

31 if the page size is 2 KB

15 if the page size is 4 KB

units

pages

range of values

For storage managers that support long transfer buffers:

- 1 - 16383 pages when the page size is 4 KB
- 1 - 32766 pages when the page size is 2 KB

For storage managers that do not support long transfer buffers:

- 1 - 15 if the database server base page size is 4 KB
- 1 - 31 if the database server base page size is 2 KB

takes effect

When ON-Bar starts

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` or equivalent SQL administration API command.

Usage

The database server passes the transfer buffer to ON-Bar and the storage manager.

The value of `BAR_XFER_BUF_SIZE` is in database server base page sizes. For Linux™, Solaris, and HP, the database server base page size is 2 KB; and for AIX®, Windows™, and Mac, the database server base page size is 4 KB. To calculate the size of the transfer buffer in a storage space or logical-log backup, multiply the value of the `BAR_XFER_BUF_SIZE` configuration parameter by the system page size, as shown in the following formula:

```
one transfer buffer KB = BAR_XFER_BUF_SIZE * pagesize
```

You can determine the system page size by running the `onstat -b` command.

The maximum size of the transfer buffer for most storage managers is 64 KB. IBM® Spectrum Protect supports long transfer buffer sizes of up to 65532 KB.

The maximum size of the transfer buffer for many storage managers is 64 KB. IBM® Spectrum Protect and HCL Informix® Primary Storage Manager support long transfer buffer sizes of up to 65532 KB.

To calculate how much memory, in KB, the database server needs for each transfer buffer, use the following formula:

```
memory KB = (BAR_XFER_BUF_SIZE * pagesize) + 500 bytes/1028
```

The extra 500 bytes is for system use. For example, if `BAR_XFER_BUF_SIZE` is 15, the transfer buffer can be 66,292 bytes, or 64.5 KB.

The number of transfer buffers per backup stream is specified by the value of the `BAR_NB_XPORT_COUNT` configuration parameter, and the number of parallel backup streams is specified by the `BAR_MAX_BACKUP` configuration parameter.



Restriction: You cannot change the buffer size between a backup and restore. The values of the `AC_TAPEBLOCK` and `AC_LTAPEBLOCK` configuration parameters must be the same value as the value of the `BAR_XFER_BUF_SIZE` configuration parameter was at the time of backup.

Example

Example

For example, for a transfer buffer size of 128*2 KB (a value of 256 KB) on Linux™, specify:

```
BAR_XFER_BUF_SIZE 128
```

Related information

[Configuring ON-Bar for optional Spectrum Protect features on page 36](#)

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

[onstat -b command: Print buffer information for buffers in use on page](#)

IFX_BAR_NO_BSA_PROVIDER environment variable

Set the **IFX_BAR_NO_BSA_PROVIDER** environment variable to force ON-Bar to use the `sm_versions` file as the source of information about the XBSA library for the storage manager.

```
setenvIFX_BAR_NO_BSA_PROVIDER1
```

By default, ON-Bar communicates directly with the XBSA library for some storage managers. ON-Bar does not require that the `sm_versions` file is updated to contain information about the XBSA library for those storage managers.

Set the **IFX_BAR_NO_BSA_PROVIDER** environment variable if you are instructed to do so by IBM® Software Support.

To unset the **IFX_BAR_NO_BSA_PROVIDER** environment variable, run the following command:

```
unset IFX_BAR_NO_BSA_PROVIDER
```

IFX_BAR_NO_LONG_BUFFERS environment variable

Set the **IFX_BAR_NO_LONG_BUFFERS** environment variable to prevent the size of transfer buffers from exceeding 64 KB when the `BAR_XFER_BUF_SIZE` configuration parameter is set to a long transfer buffer size value.

```
setenvIFX_BAR_NO_LONG_BUFFERS1
```

IBM® Spectrum Protect supports long transfer buffer sizes of up to 65532 KB.

IBM® Spectrum Protect and HCL Informix® Primary Storage Manager support long transfer buffer sizes of up to 65532 KB.

Set the **IFX_BAR_NO_LONG_BUFFERS** environment variable if you are instructed to do so by IBM® Software Support.

To unset the **IFX_BAR_NO_LONG_BUFFERS** environment variable, run the following command:

```
unset IFX_BAR_NO_LONG_BUFFERS
```

Related information

[Configuring ON-Bar for optional Spectrum Protect features on page 36](#)

IFX_BAR_USE_DEDUP environment variable

Set the **IFX_BAR_USE_DEDUP** environment variable to optimize the deduplication capabilities of storage managers.

```
setenvIFX_BAR_USE_DEDUP
```

You are not required to set the **IFX_BAR_USE_DEDUP** environment variable to a value. You can set it to any value or to no value.

Deduplication is a storage manager feature that reduces the size of backups by removing data that exists in older backups. To use deduplication, enable deduplication for your storage manager and set the **IFX_BAR_USE_DEDUP** environment variable in the database server environment and then restart the database server. When you set the **IFX_BAR_USE_DEDUP** environment variable, the backup format has fewer unique pages, which optimizes the deduplication process.

Deduplication is incompatible with incremental backups. Do not make incremental backups while the **IFX_BAR_USE_DEDUP** environment variable is set.



Important: Backups that you take while the **IFX_BAR_USE_DEDUP** environment variable is set must be restored while the **IFX_BAR_USE_DEDUP** environment variable is set.

HCL Informix® Primary Storage Manager does not support deduplication.

To unset the **IFX_BAR_USE_DEDUP** environment variable, run the following command:

```
unset IFX_BAR_USE_DEDUP
```

After you unset the **IFX_BAR_USE_DEDUP** environment variable, you must perform a level-0 backup.

Related reference

[Editing the Spectrum Protect client system options file on page 32](#)

[Configure ontape on page 132](#)

Related information

[Configuring ON-Bar for optional Spectrum Protect features on page 36](#)

[Configuring a third-party storage manager on page 36](#)

IFX_TSM_OBJINFO_OFF environment variable

Set the **IFX_TSM_OBJINFO_OFF** environment variable to disable support for restoring backup objects that are replicated, imported, or exported between IBM® Spectrum Protect servers.

```
setenv IFX_TSM_OBJINFO_OFF 1
```

By default, ON-Bar stores unique IDs in the metadata of backup objects that are created by Spectrum Protect. During a restore, ON-Bar checks the metadata to identify the object. Therefore, ON-Bar can restore a backup whose original ID, which is stored as CopyID columns in the ON-Bar catalog, changed because the object was replicated, exported, or imported between Spectrum Protect servers. To prevent the ability to restore backup objects that are moved between Spectrum Protect servers, set the **IFX_TSM_OBJINFO_OFF** environment variable to 1.

To unset the **IFX_TSM_OBJINFO_OFF** environment variable, run the following command:

```
unset IFX_TSM_OBJINFO_OFF
```

Related information

[Configuring ON-Bar for optional Spectrum Protect features on page 36](#)

ISM_DATA_POOL configuration parameter

The ISM_DATA_POOL configuration parameter, when listed in the `onconfig` file for the database server, specifies the volume pool that you use for backing up storage spaces.

onconfig.std value

ISMDData

takes effect

When ON-Bar starts

Usage

The value for this parameter can be any volume pool that IBM® Informix® Storage Manager (ISM) (ISM) recognizes. If this parameter is not present, ISM uses the ISMDData volume pool. For details, see the *Informix® Storage Manager Administrator's Guide*.

ISM_LOG_POOL configuration parameter

The ISM_LOG_POOL parameter, when listed in the `onconfig` file for the database server, specifies the volume pool that you use for backing up logical logs.

onconfig.std value

ISMLogs

takes effect

When ON-Bar starts

Usage

The value for this parameter can be any volume pool that IBM® Informix® Storage Manager (ISM) (ISM) recognizes. If this parameter is not present, ISM uses the ISMLogs volume pool. For details, see the *Informix® Storage Manager Administrator's Guide*.

LTAPEBLK configuration parameter

Use the LTAPEBLK configuration parameter to specify the block size of the device to which the logical logs are backed up when you use **ontape** for dbspace backups.

LTAPEBLK also specifies the block size for the device to which data is loaded or unloaded when you use the `-l` option of **onload** or **onunload**. If you are using **onload** or **onunload**, you can specify a different block size at the command line.

onconfig.std value

- On UNIX™: 32
- On Windows™: 16

units

Kilobytes

range of values

Values greater than (*page size*/`1024`)

To obtain the page size, run the `onstat -b` command.

takes effect

For **ontape**:

- When you execute **ontape**.
- When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` command.

For **onload** and **onunload**: When the database server is shut down and restarted

Usage

Specify LTAPEBLK as the largest block size permitted by your tape device. The database server does not check the tape device when you specify the block size. Verify that the LTAPEDEV tape device can read the block size that you specify. If not, you might not be able to read from the tape.

UNIX™ only: The UNIX™ `dd` utility can verify that the LTAPEDEV tape device can read the block size. It is available with most UNIX™ systems.

If you specify a LTAPEBLK value, ON-Bar ignores the value.

Related reference

[Changing your ontape configuration on page 138](#)

[LTAPEDEV configuration parameter on page 246](#)

[LTAPESIZE configuration parameter on page 247](#)

[ontape backup and restore system on page 132](#)

[TAPEBLK configuration parameter on page 250](#)

Related information

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

[The onunload and onload utilities on page](#)

LTAPEDEV configuration parameter

Use the LTAPEDEV configuration parameter to specify the device or directory file system to which the logical logs are backed up when you use **ontape** for backups.

The LTAPEDEV configuration parameter also specifies the device to which data is loaded or unloaded when you use the **-l** option of **onload** or **onunload**. If you are using LTAPEDEV to specify a device for **onunload** or **onload**, the same information for TAPEDEV is relevant for LTAPEDEV.

onconfig.std value

On UNIX™: **/dev/tapedev** On Windows™: **NUL**

if not present

On UNIX™: **/dev/null** On Windows™: **NUL**

takes effect

For **ontape**:

- When you execute **ontape**, if set to a tape device.
- When the database server is shut down and restarted, if set to **/dev/null** on UNIX™ or **nul** on Windows™.
- When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` command.
- When you reset the value in memory by running the `onmode -wm` command.

For **onload** and **onunload**: When the database server is shut down and restarted

Usage



Warning: Do not set LTAPEDEV to **/dev/null** or **nul** when you use ON-Bar to back up logical logs.

If you specify a tape device in the LTAPEDEV configuration parameter, ON-Bar ignores the value.



Important: Set LTAPEDEV to `/dev/null` or leave it blank on UNIX™ or `NUL` on Windows™ only if you do not want to back up the logical logs. You must take the database server offline before you change the value of LTAPEDEV to `/dev/null`.

When you set LTAPEDEV to `/dev/null`:



- The database server frees the logical logs without requiring that you back up those logs. The logical logs do not get marked as free, but the database server can reuse them.
- The ON-Bar activity log shows a warning and return code 152. Because the database server marks the logical logs as backed up when they are no longer current, ON-Bar cannot find logical logs to back up. All transactions in those logs are lost, and you are not able to restore them.

If you performed a whole-system backup with LTAPEDEV set to null, you must use the `onbar -r -w -p` command during restore to notify ON-Bar that you do not want to restore the logs. .

Related reference

[Changing your ontape configuration on page 138](#)

[LTAPEBLK configuration parameter on page 244](#)

[ontape backup and restore system on page 132](#)

[TAPEDEV configuration parameter on page 251](#)

Related information

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

[The onunload and onload utilities on page](#)

LTAPESIZE configuration parameter

Use the LTAPESIZE configuration parameter to specify the maximum tape size of the device to which the logical logs are backed up when you use ontape for backups.

The LTAPESIZE configuration parameter also specifies the maximum tape size of the device to which data is loaded or unloaded when you use the `-l` option of the onload or onunload utility. If you are using onload or onunload, you can specify a different tape size on the command line. If you want to use the full capacity of a tape, set LTAPESIZE to 0.

onconfig.std value

LTAPESIZE 0

units

KB

range of values

0 - 4294967296 (4 TB)

0 - 9223372036854775807 (9 ZB)

takes effect

For the ontape utility:

- When you run an `ontape` command.
- When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` command.

For the `onload` and `onunload` utilities: When the database server is shut down and restarted

Usage

`LTAPESIZE` specifies the maximum tape size of the device to which the logical logs are backed up when you run the `ontape` utility for backups. `LTAPESIZE` also specifies the maximum tape size of the device to which data is loaded or unloaded when you use the `-l` option of the `onload` or `onunload` utility. If you are using the `onload` or `onunload` utility, you can specify a different tape size on the command line. If you want to use the full capacity of a tape, set `LTAPESIZE` to 0.



Note: If the `BACKUP_FILTER` parameter is set in the `ONCONFIG` file, the `LTAPESIZE` cannot be set to 0. Otherwise the `ontape` utility returns an error when backing up logical logs to a directory on disk. The error message is:

```
The LTAPESIZE configuration parameter cannot be set to 0 when the BACKUP_FILTER
configuration parameter is set; change the value of LTAPESIZE.
Program over.
```

A workaround is to set the `LTAPESIZE` configuration parameter to a very high value. Log files are not much higher than the `LOGSIZE` configuration parameter. Use the value in the `LOGSIZE` as the upper limit for this database.

If you specify a `LTAPESIZE` value, `ON-Bar` ignores the value.

Related reference

[Changing your `ontape` configuration on page 138](#)

[Specify the tape size on page 137](#)

[`LTAPEBLK` configuration parameter on page 244](#)

[`ontape` backup and restore system on page 132](#)

[`TAPESIZE` configuration parameter on page 253](#)

Related information

[`onmode -wf, -wm`: Dynamically change certain configuration parameters on page](#)

[The `onunload` and `onload` utilities on page](#)

`RESTARTABLE_RESTORE` configuration parameter

Use the `RESTARTABLE_RESTORE` configuration parameter to enable or disable restartable restores.

`onconfig.std` value

```
RESTARTABLE_RESTORE ON
```


values**OFF**

Disables restartable restore. If a restore fails and `RESTARTABLE_RESTORE` is OFF, you are not able to restart it.

ON

Enables restartable restore. Set `RESTARTABLE_RESTORE` to ON before you begin a restore. Otherwise, you will be unable to restart the restore after a failure.

takes effect

After you edit your `onconfig` file. If you need to restart a physical restore, you do not need to restart the database server before you can use `RESTARTABLE_RESTORE`. If you need to restart a logical restore, you must restart the database server before you can use restartable restore.

Turning on `RESTARTABLE_RESTORE` slows down logical restore performance. For more information, see [onbar -RESTART syntax: Restarting a failed restore on page 87](#).

Related reference

[onbar -RESTART syntax: Restarting a failed restore on page 87](#)

[Resolve a failed restore on page 90](#)

RESTORE_FILTER configuration parameter

Use the `RESTORE_FILTER` configuration parameter to specify the path name of a filter program, and any options.

onconfig.std value

Not set. Restored data is not filtered.

values

The path name of a command and any options. By default, the path name is relative to the `$INFORMIXDIR/bin` directory, otherwise, the path name must be the absolute path of the program. If you include command-line options, both the filter name and the options must be surrounded by single quotation marks.

takes effect

After you edit your `onconfig` file and ON-Bar or **ontape** starts.

Usage

This filter transforms data that was transformed during backup to its original format prior to a restore. The filter specified by the `RESTORE_FILTER` configuration parameter must match the filter specified by the `BACKUP_FILTER` configuration parameter. For example, if data was compressed during backup, data must be uncompressed during a restore.

For security purposes, filters should not have write permission to non-privileged users. Permission on the filters are the same as that of permission on other executable files that are called by the HCL Informix® server or utilities.

For example, if you want to compress backed up data, you could set the `BACKUP_FILTER` and `RESTORE_FILTER` configuration parameters to the following values:

```
BACKUP_FILTER /bin/compress
RESTORE_FILTER /bin/uncompress
```

The `RESTORE_FILTER` configuration parameter can include command-line options as well as the filter name. For example, specify:

```
RESTORE_FILTER 'my_decrypt -file /var/adm/encryption.pass'
```

In this example, the command in quotation marks is used as the filter.

Related reference

[BACKUP_FILTER configuration parameter on page 221](#)

TAPEBLK configuration parameter

Use the `TAPEBLK` configuration parameter to specify the block size of the device to which **ontape** writes during a storage-space backup.

onconfig.std value

- On UNIX™: 32
- On Windows™: 16

units

Kilobytes

range of values

Values greater than `pagesize/1024`

To obtain the page size, run the `onstat -b` command.

takes effect

For **ontape**:

- When you execute **ontape**.
- When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` command.
- When you reset the value in memory by running the `onmode -wm` command.

For **onload** and **onunload**: When the database server is shut down and restarted

Usage

`TAPEBLK` also specifies the default block size of the device to which data is loaded or unloaded when you use the **onload** or **onunload** utilities. If you are using **onload** or **onunload**, you can specify a different block size on the command line.

The database server does not check the tape device when you specify the block size. Verify that the TAPEBLK tape device can read the block size that you specify. If not, you might not be able to read from the tape.

If you specify a TAPEBLK value, ON-Bar ignores the value.

Related reference

[Changing your ontape configuration on page 138](#)

[TAPEDEV configuration parameter on page 251](#)

[TAPESIZE configuration parameter on page 253](#)

[LTAPEBLK configuration parameter on page 244](#)

[ontape backup and restore system on page 132](#)

Related information

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

[The onunload and onload utilities on page](#)

TAPEDEV configuration parameter

Use the TAPEDEV configuration parameter to specify the device or directory file system to which the ontape utility backs up storage spaces.

onconfig.std value

On UNIX™: **/dev/tapedev**

On Windows™: **\\.\TAPE0**

if not present

On UNIX™: **/dev/null**

On Windows™: **NUL**

units

Path name

takes effect

For the ontape utility:

- If it is set to `/dev/null` on UNIX™ or `NUL` on Windows™, when the database server is shut down and restarted
- If it is set to a tape device, when you run the ontape utility
- When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` command.
- When you reset the value in memory by running the `onmode -wm` command.

For the onload and onunload utilities: When the database server is shut down and restarted

Usage

The `ontape` utility reads the value of the `TAPEDEV` parameter at the start of processing. If you set `TAPEDEV` to `/dev/null`, you must do it before you start `ontape` to request the backup. When you set `TAPEDEV` to `/dev/null` and request a backup, the database server bypasses the backup but still updates the `dbspaces` with the new backup time stamps.

You can set the `TAPEDEV` configuration parameter to `STDIO` to direct **ontape** utility back up and restore operations to standard I/O instead of to a device.

The `TAPEDEV` configuration parameter also specifies the default device to which data is loaded or unloaded when you use the `onload` or `onunload` utilities. However, if `TAPEDEV` is set to `STDIO`, the `onunload` utility will not be able to unload data.

If you change the tape device, verify that the `TAPEBLK` and `TAPESIZE` configuration parameter values are correct for the new device.

If you specify a `TAPEDEV` value, ON-Bar ignores the value.

Remote devices (UNIX™)

You can perform a storage-space backup across your network to a remote device attached to another host computer on UNIX™ and Linux™ platforms. The remote device and the database server computer must have a trusted relationship so that the `rsh` or the `rlogin` utility can connect from the database server computer to the remote device computer without asking for password. You can establish a trusted relationship by configuring the `/etc/hosts.equiv` file, the user's `~/.rhosts` files, or any equivalent mechanism for your system on the remote device computer. If you want to use a different utility to handle the remote session than the default utility used by your platform, you can set the `DBREMOTECMD` environment variable to the specific utility that you want to use.

Symbolic links to remote devices (UNIX™)

The `TAPEDEV` configuration parameter can be a symbolic link, enabling you to switch between tape devices without changing the path name that the `TAPEDEV` configuration parameter specifies.

Use the following syntax to specify a tape device attached to another host computer:

```
host_machine_name:tape_device_pathname
```

The following example specifies a tape device on the host computer **kyoto**:

```
kyoto:/dev/rmt01
```

Rewinding tape devices before opening and on closing

The tape device that The `TAPEDEV` configuration parameter specifies must perform a rewind before it opens and when it closes. The database server requires this action because of a series of checks that it performs before it writes to a tape.

When the database server attempts to write to any tape other than the first tape in a multivolume `dbspace` or `logical-log` backup, the database server first reads the tape header to make sure that the tape is available for use. Then the device is closed and reopened. The database server assumes the tape was rewound when it closed, and the database server begins to write.

Whenever the database server attempts to read a tape, it first reads the header and looks for the correct information. The database server does not find the correct header information at the start of the tape if the tape device did not rewind when it closed during the write process.

Related reference

[Changing your ontape configuration on page 138](#)

[TAPEBLK configuration parameter on page 250](#)

[LTAPEDEV configuration parameter on page 246](#)

[ontape backup and restore system on page 132](#)

Related information

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

[The onunload and onload utilities on page](#)

TAPESIZE configuration parameter

Use the TAPESIZE parameter specifies the size of the device to which the ontape utility backs up storage spaces.

onconfig.std value

TAPESIZE 0

units

KB

range of values

0 - 4294967296 (4 TB)

0 - 9223372036854775807 (9 ZB)

takes effect

For ontape:

- When you run an ontape command.
- When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` command.
- When you reset the value in memory by running the `onmode -wm` command.

For the onload and onunload utilities: When the database server is shut down and restarted.

Usage

The TAPESIZE also specifies the size of the default device to which data is loaded or unloaded when you use the onload or onunload utility. When you run the onload or onunload utility, you can specify a different tape size on the command line. If you want to use the full physical capacity of a tape, set TAPESIZE to 0.



Note: Tape size is irrelevant if TAPEDEV is set to `STDIO`.

If you specify a TAPESIZE value, ON-Bar ignores the value.

Related reference

[Changing your ontape configuration on page 138](#)

[Specify the tape size on page 137](#)

[TAPEBLK configuration parameter on page 250](#)

[ontape backup and restore system on page 132](#)

[LTAPESIZE configuration parameter on page 247](#)

Related information

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

[The onunload and onload utilities on page](#)

The archecker utility configuration parameters and environment variable

These topics describe the **AC_CONFIG** environment variable and the configuration parameters that you use with the archecker utility.

The **archecker** utility uses the configuration parameters in the `ac_config.std` template to verify a backup or perform a table-level restore. If you need to change these parameters, copy the `ac_config.std` template to the AC_CONFIG file. The AC_CONFIG environment variable specifies the location of the AC_CONFIG file.

Because ON-Bar calls the archecker utility to verify backups, you must configure the archecker environment variable and parameters before you can use the `onbar -v` option.

You can also use other archecker configuration parameters that do not have default values in the `ac_config.std` file, but are valid in that file.

Table 39. Configuration parameters that the archecker utility uses

Configuration parameter	Description
AC_DEBUG	Prints debugging messages in the archecker message log.
AC_IXBAR	Specifies the path name to the IXBAR file. If not set in the <code>ac_config</code> file, the value of the <code>BAR_IXBAR_PATH</code> configuration parameter is used.
AC_LTAPEBLOCK	Specifies the ontape block size for reading logical logs.

Table 39. Configuration parameters that the archecker utility uses (continued)

Configuration parameter	Description
	If not set in the <code>ac_config</code> file, the value of the LTAPEBLOCK configuration parameter is used.
AC_LTAPEDEV	Specifies the local device name used by ontape for reading logical logs. If not set in the <code>ac_config</code> file, the value of the LTAPEDEV configuration parameter is used.
AC_MSGPATH	Specifies the location of the archecker message log. This configuration parameter is in the default <code>ac_config</code> file.
AC_SCHEMA	Specifies the path name to the archecker schema command file.
AC_STORAGE	Specifies the location of the temporary files that archecker builds. This configuration parameter is in the default <code>ac_config</code> file.
AC_TAPEBLOCK	Specifies the tape block size in kilobytes. If not set in the <code>ac_config</code> file, the value of the TAPEBLOCK configuration parameter is used.
AC_TAPEDEV	Specifies the local device name used by the ontape utility. If not set in the <code>ac_config</code> file, the value of the TAPEDEV configuration parameter is used.
AC_TIMEOUT	Specifies the timeout value for the onbar and the archecker processes if one of them exits prematurely.
AC_VERBOSE	Specifies either verbose or terse mode for archecker messages. This configuration parameter is in the default <code>ac_config</code> file.
BAR_BSALIB_PATH	Identical to the BAR_BSALIB_PATH server configuration parameter that is in the <code>onconfig.std</code> file. For more information, see BAR_BSALIB_PATH configuration parameter on page 224 .

If you use the **ontape** utility and the AC_TAPEDEV, AC_TAPEBLK, AC_LTAPEDEV and AC_LTAPEBLK configuration parameters are not set in the AC_CONFIG file, the archecker utility will use the values specified in the TAPEDEV, TAPEBLK, LTAPEDEV, LTAPEBLK configuration parameters specified in the `onconfig` file.

Related reference[archecker table level restore utility on page 203](#)

AC_CONFIG file environment variable

Set the **AC_CONFIG** environment variable to the full path name for the archecker configuration file (either `ac_config.std` or user defined).

```
setenvAC_CONFIG pathname
```

default value

UNIX™: `$INFORMIXDIR/etc/ac_config.std`

Windows™: `%INFORMIXDIR%\etc\ac_config.std`

takes effect

When ON-Bar starts

The following are examples of valid **AC_CONFIG** path names:

- UNIX™: `/usr/dbserver/etc/ac_config.std` and `/usr/local/my_ac_config.std`
- Windows™: `c:\dbserver\etc\ac_config.std` and `c:\dbserver\etc\my_ac_config.std`

If **AC_CONFIG** is not set, the archecker utility sets the default location for the archecker configuration file to `$INFORMIXDIR/etc/ac_config.std` on UNIX™ or `%INFORMIXDIR%\etc\ac_config.std` on Windows™.

Important: If you do not specify the entire path, including the configuration file name in the **AC_CONFIG** file, the archecker utility might not work correctly.

AC_DEBUG configuration parameter

The **AC_DEBUG** configuration parameter causes debugging messages to be printed in the archecker message file. Use this parameter only as directed by technical support.

The use of this configuration parameter can cause the archecker message log file to grow very large and can substantially slow down archecker processing.

Default value

Off

Range

1-16

AC_IXBAR configuration parameter

Use the **AC_IXBAR** configuration parameter to specify the location of the **IXBAR** file.

Default value

None

Range

Any valid path name

AC_LTAPEBLOCK configuration parameter

Use the AC_LTAPEBLOCK configuration parameter to the **ontape** block size for reading logical logs.

Default value

32 kilobytes

Range

0 - 2,000,000,000

Usage

When you perform an archive with:

- onbar -b, the value of AC_TAPEBLOCK should be the value the BAR_XFER_BUF_SIZE configuration parameter multiplied by the current page size. For more information, see [BAR_XFER_BUF_SIZE configuration parameter on page 240](#).
- ontape -t, the value of AC_LTAPEBLOCK should be the value that the TAPEBLK ONCONFIG configuration parameter was set to at the time of the archive. For more information, see [Specify the tape-block-size on page 137](#).

AC_LTAPPEDEV parameter

Use the AC_LTAPPEDEV configuration parameter to specify the local device name that is used by the ontape utility.

If the tape device is set to STDIO, archecker receives input from standard input.

Default value

None

Range

Any valid path name or STDIO

AC_MSGPATH configuration parameter

Use the AC_MSGPATH parameter in the **AC_CONFIG** file to specify the location of the archecker message log (`ac_msg.log`).

ac_config.std value

UNIX™: `AC_MSGPATH /tmp/ac_msg.log`

Windows™: `AC_MSGPATH c:\temp\ac_msg.log`

takes effect

When ON-Bar starts

Usage

You must specify the entire path of the message log in the **AC_CONFIG** file or else the archecker utility might not work correctly.

When you verify backups with `onbar -v`, the archecker utility writes summary messages to the `bar_act.log` and indicates whether the verification succeeded or failed. It writes detailed messages to the `ac_msg.log`. If the backup fails verification, discard the backup and try another backup, or give the `ac_msg.log` to IBM® Software Support. For sample messages, see [onbar -v syntax: Verifying backups on page 58](#).

AC_SCHEMA configuration parameter

Use the AC_SCHEMA configuration parameter to specify the path name to the archecker schema command file.

Default value

None

Range

Any valid path name

This configuration parameter is overridden by the `-f cmdfile` command line option.

AC_STORAGE configuration parameter

Use the AC_STORAGE configuration parameter in the **AC_CONFIG** file to specify the location of the directory where archecker stores its temporary files.

ac_config.std value

UNIX™: /tmp

Windows™: c:\temp

takes effect

When ON-Bar starts

Usage

You must specify the entire path of the storage location in the **AC_CONFIG** file or else the archecker utility might not work correctly.

The following table lists the directories and files that archecker builds. If verification is successful, these files are deleted.

Table 40. The archecker temporary files

Directory	Files
CHUNK_BM	Bitmap information for every backed up storage space.
INFO	Statistical analysis and debugging information for the backup.
SAVE	Partition pages in the <code>PT.#####</code> file. Chunk-free pages in the <code>FL.#####</code> file. Reserved pages in the <code>RS.#####</code> file. Blob-free map pages in the <code>BF.#####</code> file

To calculate the amount of free space that you need, see [Temporary space for backup verification on page 60](#). It is recommended that you set `AC_STORAGE` to a location with plenty of free space.

AC_TAPEBLOCK configuration parameter

Use the `AC_TAPEBLOCK` configuration parameter to specify the size of the tape block in kilobytes when an archive is performed either the `onbar -b` command or the `ontape -t` command.

Default value

32 kilobytes

Range

0 - 2,000,000,000

Usage

When you perform an archive with:

- `onbar -b`, the value of `AC_TAPEBLOCK` should be the value the `BAR_XFER_BUF_SIZE` configuration parameter multiplied by the current page size. For more information, see [BAR_XFER_BUF_SIZE configuration parameter on page 240](#).
- `ontape -t`, the value of `AC_TAPEBLOCK` should be the value that the `TAPEBLK_ONCONFIG` configuration parameter was set to at the time of the archive. For more information, see [Specify the tape-block-size on page 137](#).

AC_TAPEDEV configuration parameter

Use the `AC_TAPEDEV` configuration parameter to specify the local device name that is used by the `ontape` utility.

If the tape device is set to `STDIO`, archecker receives input from standard input.

Default value

None

Range

Any valid path name or STDIO

AC_TIMEOUT configuration parameter

Use the AC_TIMEOUT configuration parameter to specify the timeout value for the onbar and the archecker processes if one of them exits prematurely.

ac_config.std value

UNIX™: 300

Windows™: 300

units

seconds

takes effect

When the onbar-v command starts

The AC_TIMEOUT configuration parameter was introduced to avoid onbar and archecker processes waiting for each other indefinitely if one of them exits prematurely, thus avoiding the creation of an orphan and zombie process during data server initialization.

AC_VERBOSE configuration parameter

Use the AC_VERBOSE parameter in the **AC_CONFIG** file to specify either verbose or terse output in the archecker message log (`ac_msg.log`).

ac_config.std value

1

range of values

1 = verbose messages in `ac_msg.log`

0 = terse messages in `ac_msg.log`

takes effect

When ON-Bar starts

Informix® Primary Storage Manager configuration parameters

The HCL Informix® Primary Storage Manager uses the information in some specific configuration parameters.

Related information

[HCL Informix Primary Storage Manager on page 181](#)

[Setting up Informix Primary Storage Manager on page 190](#)

[Configuring Informix Primary Storage Manager on page 191](#)

[Examples: Manage storage devices with Informix Primary Storage Manager on page 183](#)

PSM_ACT_LOG configuration parameter

Use the PSM_ACT_LOG configuration parameter to specify the location of the HCL Informix® Primary Storage Manager activity log if you do not want the log information included in the ON-Bar activity log.

onconfig.std value

none

if value not present

The value of the BAR_ACT_LOG configuration parameter is used

range of values

Full path name

takes effect

When the **onpsm** utility starts

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` command.

Usage

Specify a path to an existing directory with an appropriate amount of space available or use `$(INFORMIXDIR)/psm_act.log`. If you specify a file name only, the storage manager creates the activity log in the working directory in which you started the storage manager.

If the PSM_ACT_LOG configuration parameter is not set, the Informix® Primary Storage Manager puts activity information in the directory specified with the BAR_ACT_LOG configuration parameter. To clearly distinguish ON-Bar and Informix® Primary Storage Manager activity information, use the PSM_ACT_LOG to specify a different location for the storage manager activity log.

The format of the file resembles the format of the database server message log. You can examine the activity log to determine the results of storage manager actions.

The file specified by the PSM_ACT_LOG configuration parameter is created if it does not exist.

You can also use the **PSM_ACT_LOG** environment variable to specify the location of the Informix® Primary Storage Manager activity log for your environment, for example, for a single session.

Related reference

[Message logs for Informix Primary Storage Manager on page 203](#)

[BAR_ACT_LOG configuration parameter on page 222](#)

Related information

[PSM_ACT_LOG environment variable on page](#)

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

PSM_CATALOG_PATH configuration parameter

Use the PSM_CATALOG_PATH configuration parameter to specify the full path to the directory that contains the HCL Informix® Primary Storage Manager catalog tables. These catalog tables contain information about the pools, devices, and objects managed by the storage manager.

onconfig.std value

Not set. The default value is used.

default value

UNIX™ or Linux™: `$INFORMIXDIR/etc/psm`

Windows™: `%INFORMIXDIR%\etc\psm`

range of values

Full path name for the directory that contains the Informix® Primary Storage Manager catalog tables

takes effect

When the ON-Bar or **onpsm** utility starts

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` command.

When you reset the value in memory by running the `onmode -wm` command.

Usage

You can change the default path to another location. The Informix® Primary Storage Manager places the file that information about the devices and objects in whatever directory you specify.

If you move the backup file to another location, change the value of the PSM_CATALOG_PATH configuration parameter.

You can back up the contents of the file whenever you want.

If you have multiple instances, all instances contain the same catalog tables if the PSM_CATALOG_PATH in each instance is set to the same path. You can specify a different path for each instance.

The storage manager automatically creates the catalog tables the first time you run an `onpsm` utility command or the first time that the XBSA shared library is used.

You can also use the **PSM_CATALOG_PATH** environment variable to specify the location of the Informix® Primary Storage Manager catalog tables for your environment, for example, for a single session.

Related information

[PSM_CATALOG_PATH environment variable on page](#)

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

PSM_DBS_POOL configuration parameter

Use the PSM_DBS_POOL configuration parameter to change the name of the pool in which the HCL Informix® Primary Storage Manager places backup and restore dbspace data.

onconfig.std value

DBSPool

takes effect

When the **onpsm** utility starts

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` command.

When you reset the value in memory by running the `onmode -wm` command.

Usage

The storage manager automatically places the dbspace data in the `DBSPool` or the pool you specify. You can use any combination of letters and digits.

You can also use the **PSM_DBS_POOL** environment variable to change the name of the pool for your environment, for example, for a single session.

Related information

[Device pools on page 201](#)

[PSM_DBS_POOL environment variable on page](#)

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

PSM_DEBUG configuration parameter

Use the PSM_DEBUG configuration parameter to specify the amount of debugging information that prints in the Informix® Primary Storage Manager debug log if you want to use a debug level that is different from the one used by ON-Bar.

onconfig.std value

Not set. The default value is used.

default value

The value of the `BAR_DEBUG` configuration parameter is used.

units

One digit to represent the level of debugging information that you want

range of values

0 = No debugging messages.

1 = Prints only internal errors.

2 = Prints information about the entry and exit of functions and prints internal errors.

3 = Prints the information specified by 1-2 with additional details.

4 = Prints information about parallel operations and the information specified by 1-3.

5 = Prints information about internal states in the Informix® Primary Storage Manager.

6 = Prints the information specified by 1-5 with additional details.

7 = Prints information specified by 1-6 with additional details.

8 = Prints information specified by 1-7 with additional details.

9 = Prints all debugging information.

takes effect

When the onpsm utility starts

When the ON-Bar utility executes commands and reads information that is specified in the BAR_DEBUG configuration parameter

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` command.

When you reset the value in memory by running the `onmode -wm` command.

Usage

If you set the PSM_DEBUG configuration parameter to a valid value that is higher than 0, the Informix® Primary Storage Manager logs debug messages to its debug log.

You can experiment with the debug values to find the right amount of information. Generally, if the PSM_DEBUG configuration parameter is set to 5, the storage manager prints enough information for tracing and debugging purposes.

Settings of 8 and 9 require a large amount of space.

You can also use the **PSM_DEBUG** environment variable to specify the amount of debugging information that prints in the storage manager debug log for your environment, for example, for a single session.

Related reference

[Message logs for Informix Primary Storage Manager on page 203](#)

Related information

[PSM_DEBUG environment variable on page](#)

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

PSM_DEBUG_LOG configuration parameter

Use the PSM_DEBUG_LOG configuration parameter to specify the location of the debug log to which the HCL Informix® Primary Storage Manager writes debugging messages if you do not want the log information included in the ON-Bar debug log.

onconfig.std value

Not set. The default value is used.

default value

The value of the BAR_DEBUG_LOG configuration parameter is used.

takes effect

When the onpsm utility starts

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` command.

When you reset the value in memory by running the `onmode -wm` command.

Usage

If the PSM_DEBUG_LOG configuration parameter is not set, the Informix® Primary Storage Manager puts activity information in the directory specified with the BAR_DEBUG_LOG configuration parameter. To clearly distinguish ON-Bar and Informix® Primary Storage Manager activity information, use the PSM_DEBUG_LOG to specify a different location for the Informix® Primary Storage Manager activity log.

For security reasons, set the PSM_DEBUG_LOG configuration parameter to a directory with restricted permissions, such as the `$INFORMIXDIR` directory.

If the directory that holds the debug file becomes too large, you can erase the file. You need to retain information only if there are problems that need to be debugged.

You can also use the **PSM_DEBUG_LOG** environment variable to specify the location of the debug log for your environment, for example, for a single session.

Related reference

[Message logs for Informix Primary Storage Manager on page 203](#)

[BAR_DEBUG_LOG configuration parameter on page 230](#)

Related information

[PSM_DEBUG_LOG environment variable on page](#)

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

PSM_LOG_POOL configuration parameter

Use the PSM_LOG_POOL configuration parameter to change the name of the pool in which the HCL Informix® Primary Storage Manager places backup and restore log data.

onconfig.std value

LOGPOOL

takes effect

When the **onpsm** utility starts

When you reset the value dynamically in your `onconfig` file by running the `onmode -wf` command.

When you reset the value in memory by running the `onmode -wm` command.

Usage

The storage manager automatically places the log data in the `LOGPOOL` or the pool you specify. You can use any combination of letters and digits.

You can also use the **PSM_LOG_POOL** environment variable to change the name of the pool for your environment, for example, for a single session.

Related information

[Device pools on page 201](#)

[PSM_LOG_POOL environment variable on page](#)

[onmode -wf, -wm: Dynamically change certain configuration parameters on page](#)

Event alarm configuration parameters

When you set configuration parameters for use with the ON-Bar and **ontape** utilities, also determine if you need to adjust the `ALARMPROGRAM` and `ALRM_ALL_EVENTS` configuration parameters.

Use the [ALARMPROGRAM configuration parameter on page](#) to set the **log_full.sh** script to automatically back up log files when they become full.

Use the [ALRM_ALL_EVENTS configuration parameter on page](#) to cause ALARMPROGRAM to execute every time an alarm event is invoked.

Cloud Backup

These topics provide information about storing and retrieving the backups directly to the ecosystem of selected cloud providers, namely Amazon S3 and Softlayer Object Storage.

Back up to Amazon Simple Storage Service using ON-Bar and the PSM

You can use ON-Bar and the Primary Storage Manager to back up and restore data to or from the Amazon Simple Storage Service (S3). You are responsible for terms and any charges associated with your use of the Amazon Simple Storage Service.

Before you begin

Prerequisites:

- You must have an Amazon account to perform cloud storage backups. See the Amazon website for instructions about setting up an account.

About this task

The following steps show how to back up data to the Amazon Simple Storage Service (S3) System and restore from it by using ON-Bar and the PSM. In this context, cloud storage refers to an online storage service over the Internet. If you choose to back up to cloud storage, you do not need to provide local devices. Instead, you back up the data to a virtual device, most likely located on the Internet.

1. Create a group to access S3
 - a. Using a web browser, navigate to the Amazon S3 website and log on.
 - b. Select **Groups** tab and click **Create New Group**.
 - c. Specify the Group name and click **Next Step**.
 - d. In the Attach Policy page, attach a policy to the group. In this case we will select **AmazonS3FullAccess** . As the name implies, this policy will allow any member of this group to do anything in all the containers in S3. Click **Next Step**.



Note:

If you want to use S3 for other purposes or for multiple instances, you can change this by going in to the **Policies** tab before creating the group and creating a customized policy that suits your needs. Ensure to be careful before granting full access to the Amazon S3, as it will provide the user to access all the S3 data in the account .

- e. In the Review page, review your entries and click **Create Group**.

Result

The new group is created.

2. Create a User to Access S3

- a. Select **Users** tab and click **Create New Users**.
- b. Specify the User name and select **Generate and access key for each user**.
- c. Click **Create**.

Result

Access Key and the Secret Access Key are generated. These keys are equivalent to username and password that can be used to store and retrieve data from S3 using APIs.



Note: Only authorized user can access these keys. You need to download these credentials into a text file and store in a safe location. If you lose them, you need to create a new user again.

- d. Click **Close**.

Result

User page with newly created user will be displayed.

3. Assign the user to the group.

- a. Click the check box next to the user and select **User Actions**.
- b. Click **Add User to Groups**.

4. Create a bucket.

- a. Select **Services** tab and click **Amazon S3**.
- b. Click **Create Bucket**.
- c. Specify a bucket name and select the appropriate region.

5. Configure Primary Storage Manager to use the Bucket and Credentials you just created.

Create a device in PSM of type CLOUD using S3 as provider.

```
onpsm -D add S3_CLOUD_DEV \  
-g DBSPool \  
-p HIGHEST \  
-t CLOUD \  
--url https://ifmx-s3-dev.s3.amazonaws.com \  
--provider S3 \  
--user AKIAIT111115555X4PA \  
--password A2nB2111115555nvTI0X9ZxGzUJNJivoBQY9MrD \  
--container ifmx-s3-dev \  
--region us-east-1 \  
--max_part_size 25600
```

In this command line:

- a. **S3_CLOUD_DEV**, is the arbitrary name given to this device. With FILE type devices, this is actually the full path of the directory that will store the data, but in the case of CLOUD type devices, it is just any name that will help you organize your devices. This name and the pool (DBSPool in this case) must be unique.
- b. **-t CLOUD**, is the device type that enable PSM to store/retrieve the data to/from a CLOUD infrastructure.
- c. **--provider S3**, is the target cloud provider (Amazon S3) in this case. At present, only S3 and SWIFT (OpenStack SWIFT) are supported.
- d. **--url <https://ifmx-s3-dev.s3.amazonaws.com>** is the URL where your backups will go. In the specific case of S3, the bucket name is part of the URL provided.
- e. **--user AKIAIT111155555X4PA**, for S3 is the Access Key provided to you when your user was created.
- f. **--password A2nB2111155555nvTIOX9ZxGzUJNJivoBQY9MrD**, for S3 is the Secret Key provided to you when the user was created.
- g. **--container ifmx-s3-dev** is the amazon bucket.
- h. **--region us-east-1** is used for V4 authentication.



Note: If **--region** is not specified, PSM will use V2 authentication.

- i. **--max_part_size 25600** will fragment your objects in 25MB pieces, in the case of S3, size between 25 and 100 MB is recommended.

6. Check the created device:

```
$ onpsm -D list

Informix Primary Storage Manager Device List

Type   Prio    Block/Size (MB)  Pool Name    Device Name
CLOUD  HIGHEST  --/--           DBSPool     S3_CLOUD_DEV
CLOUD  HIGHEST  --/--           LOGPool     S3_CLOUD_DEV
```

7. Take a Level Zero Backup:

```
$ onbar -b -L 0 -w

$ echo $?

$ 0
```

8. Check Backup Data in your Bucket using S3 management console.

Back up to Softlayer using ON-Bar and the PSM

You can use ON-Bar and the Primary Storage Manager to back up and restore data to or from the Softlayer Object Storage. You are responsible for terms and any charges associated with your use of the Softlayer.

Before you begin

Prerequisites:

- You must have a Softlayer account to perform cloud storage backups. See the Softlayer website for instructions about setting up an account.

About this task

The following steps show how to back up data to the Softlayer System and restore from it by using ON-Bar and the PSM. In this context, cloud storage refers to an online storage service over the Internet. If you choose to back up to cloud storage, you do not need to provide tapes. Instead, you back up the data to a virtual device, most likely located on the Internet.

1. Create an Account Name.
 - a. Using a web browser, navigate to the Softlayer website and log on.
 - b. Select **Storage** tab and click **Object Storage**.
 - c. Click **Order Object Storage**.
 - d. Select **Storage type** and click **Continue**.
 - e. In the Confirm Order page, accept the Master Service Agreement and click **Place Order**.

Result

The account is created. Now, you can retrieve the credentials which will include an authentication endpoint (URL) both private and public, a username, and a Key or password.



Note: If you have a machine with access to the internal private network to Softlayer, you should use the private URL if not, you should use the public one.

2. The PSM command line will request you to provide a container, you can provide any name you want, the container does not need to be created in the SOFTLAYER GUI, it will be created automatically by PSM.
3. Configure Primary Storage Manager.

Create a device in PSM of type CLOUD using SWIFT as provider.

```
onpsm -D add SOFTLAYER1 \
-g DBSPPOOL \
-p HIGHEST \
-t CLOUD \
--url https://dal05.objectstorage.softlayer.net/auth/v1.0 \
--provider SWIFT \
--user IBM33456-2:arturo_cavero \
--password bc076837ba9959859c306051c1cda7fb2fe3f012cba15a60c \
--container ifmx_dev \
--max_part_size 25600
```

In this command line:

- a. **SOFTLAYER1**, is the arbitrary name given to this device. With FILE type devices, this is actually the full path of the directory that will store the data, but in the case of CLOUD type devices, it is just any name that will help you organize your devices. This name and the pool (DBSPPOOL in this case) must be unique.
- b. **'-t CLOUD'**, is the device type that enable PSM to store/retrieve the data to/from a CLOUD infrastructure.
- c. **'--provider SWIFT'**, is the target cloud provider (Softlayer) in this case. At present, only SWIFT (OpenStack SWIFT) is supported.
- d. **'--url https://dal05.objectstorage.softlayer.net/auth/v1.0'** is the URL where your backups will go.

- e. `--user IBM33456-2:arturo_cavero`, for SWIFT is the User Name provided to you when your user was created.
- f. `--password bc076837ba9959859c306051c1cda7fb2fe3f012cba15a60c`, for SWIFT is the Secret Key provided to you when the user was created.
- g. `--container ifmx_dev` is the Softlayer container.
- h. `--max_part_size 25600` will fragment your objects in 25MB pieces, in the case of SWIFT, size between 25 and 100 MB is recommended.

4. Check the created device:

```
$ onpsm -D list

Informix Primary Storage Manager Device List

Type   Prio   Block/Size (MB) Pool Name   Device Name
CLOUD  HIGHEST  --/--      DBSPPOOL   SOFTLAYER1
CLOUD  HIGHEST  --/--      LOGPOOL    SOFTLAYER1
```

5. Take a Level Zero Backup:

```
$ onbar -b -L 0 -w

$ echo $?

$ 0
```

6. Check Backup Data in your Bucket using Softlayer management console.

Appendixes

Troubleshooting some backup and restore errors

This appendix lists some error and informational messages that you can receive during a backup or restore, describes under what circumstances the errors might occur or the message might appear, and provides possible solutions or workarounds.

Find errors by viewing the ON-Bar activity log

The database server does not show errors in standard output (stdout) if an error occurs when you use `onbar -b` to back up storage spaces or `onbar -r` to restore storage spaces. Therefore, when you use `onbar -b` or `onbar -r`, you must check information in the ON-Bar activity log (`bar_act_log`). As ON-Bar backs up and restores data, it writes progress messages, warnings, and error messages to the `bar_act.log`.

Corrupt page during an archive

The message `Archive detects that page is corrupt` indicates that page validation failed. If you receive this message, you can identify the table that has the corrupt page.

During an archive, the database server validates every page before writing it to the archive device. This validation checks that the elements on the page are consistent with the expected values. When a page fails this validation, a message similar to the following is written to the `online.log` file:

```
16:27:49 Assert Warning: Archive detects that page 1:10164 is corrupt.
16:27:49 Who: Session(5, informix@cronus, 23467, 10a921048)
16:27:49 Thread(40, arcbackup1, 10a8e8ae8, 1)
16:27:49 File: rsarcbu.c Line: 2915
16:27:49 stack trace for pid 23358 written to /tmp/af.41043f4
16:27:49 See Also: /tmp/af.41043f4
16:27:49 Archive detects that page 1:10164 is corrupt.
16:27:50 Archive on rootdbs Completed with 1 corrupted pages detected.
```

The archive stops after detecting 10 corrupt pages. The `online.log` file displays the full error message, including the page address, for the first 10 errors. Subsequently, only the count of the number of corrupt pages is put in to the `online.log`.

After you receive this message, identify which table the corrupt page belongs to by examining the output of the `oncheck -pe` command. To determine the extent of the corruption, execute the `oncheck -clD` command for that table.

A corrupt page is saved onto the backup media. During a restore, the corrupt page is returned in its corrupt form. No error messages are written to the `online.log` when corrupt pages are restored, only when they are archived.

Log backup already running

When using ON-Bar to create a backup, the informational messages `log backup is already running` in the `bar_act.log` file and `Process exited with return code 152` in the `online.log` file might appear under some circumstances.

These messages can appear under the following circumstances:

- When the `ALARMPROGRAM` configuration parameter is set to `log_full.sh`.

Periodically, events cause `log_full.sh` to trigger the `onbar -b -l` command. If a log fills while the `onbar -b -l` command is running, then ON-Bar backs up that log as well. If the backup has not completed by the time of the next event trigger, it generates a warning in the `bar_act.log` file. At the time of the next event trigger, the log backup can continue.

- When the `onbar -b -l` command is started automatically.

A level-0 archive (especially when started with the `-w` option) first archives the database and then automatically start the `onbar -b -l` command to back up any logical logs that are currently full and not yet backed up. There might not be a `log_full.sh` message in `online.log`, because the `onbar -b -l` command is started directly.

- When you mount a new tape after filling a previous tape, a `log_full.sh` event is scheduled but not triggered.

As soon as the next log fills and generates an event trigger in the `log_full.sh` file, all available logs are archived.

You can force the archive by running `onbar -b -l` or force `log_full.sh` to be triggered by running `onmode -l`.

No server connection during a restore

During a whole system restore with ON-Bar, the error `archive api error: no server connection` might appear in the `bar_act.log` file. ON-Bar then connects to the storage manager successfully, but eventually fails with the error `archive api error: not yet open`. If you receive these message, you can take steps to solve the problem.

The `bar_act.log` file contains information similar to the following messages:

```
2000-03-09 10:51:06 19304 19303 /usr/informix/bin/onbar_d -r -w
2000-03-09 10:51:09 19304 19303 ERROR: Unable to start the physical restore:
Archive API error: no server connection.
2000-03-09 10:51:09 19304 19303 Successfully connected to Storage Manager.
2000-03-09 10:51:36 19304 19303 Process 19304 received signal 3. Process will
exit after cleanup.
2000-03-09 10:59:13 19811 19810 /usr/informix/bin/onbar_d -r -w
2000-03-09 10:59:16 19811 19810 ERROR: Unable to start the physical restore:
Archive API error: no server connection.
2000-03-09 10:59:16 19811 19810 Successfully connected to Storage Manager.
2000-03-09 11:01:12 19811 19810 Begin cold level 0 restore llog1.
2000-03-09 11:01:12 19811 19810 ERROR: Unable to write restore data to the
database server: Archive API error: not yet open.
```

To solve this problem, check if the database server is still running. If it is, shut down the database server and run the command again.

Drop a database before a restore

If you perform a level-0 archive using ON-Bar and a storage manager, then drop a database, and then perform a restore with the `onbar -r` command, the database remains dropped. The restore salvages the logs and the logs contains the `DROP DATABASE` statement. When the logs are salvaged, or replayed, the database is dropped. If you receive these message, you can take steps to solve the problem.

To prevent this situation, perform a physical restore using the `onbar -r -p` command, and then a logical restore using the `onbar -r -l` command. This sequence does not salvage the logs and does restore the database.

No dbspaces or blobspaces during a backup or restore

If the emergency boot file, `ixbar.servernum`, does not have the correct entries for objects in the backup, the message `There are no DB/BLOBspaces to backup/restore` appears in `bar_act.log` file during a restore started with the `onbar -r` or `onbar -r -w` command.

This error can appear under the following circumstances:

- During an external restore, if the emergency boot file was not copied from the source system.
- If the emergency boot file was recreated after the archive backup was made. The previous file is saved in the form: `ixbar.xx.xxxx`.
- An attempt to execute the `onbar -r -w` command with a backup that is not a full system backup.

Restore blob space BLOBs

You can use table-level restore to restore a BLOB that is stored in a table. However, restoring a BLOB that is stored in a blob space is not supported. If you attempt to restore a blob space BLOB, the column is set to NULL.

Changing the system time on the backup system

In some circumstances when there is a problem with the system time, ON-Bar fails with the message `There are no storage spaces or logical logs to backup or restore`. If this occurs, you can take steps to solve the problem.

About this task

Time lines use the UNIX™ time as the archive checkpoint time for dbspaces and the closing time for logical logs. If logs are not automatically backed up and the system clock is changed, the time line can get corrupted.

For example, if you have logical logs that were closed before the archive checkpoint time, they have a timestamp that is higher than the archive checkpoint time. The dbspace does not need the logs and ON-Bar will try to restore the backup immediately. If a log cannot be found, ON-Bar fails with the following message: `There are no storage spaces or logical logs to backup or restore`.

To restore the storage space and logical logs:

1. Change the clock back to its original value.
2. Recover the system from backup.
3. Change the clock back to the new time.

Migrate data, servers, and tools

Backing up before a database server or storage-manager upgrade

Before you upgrade to a new version of the database server, you must perform a complete backup.

About this task



Important: The database server conversion software automatically recreates the **sysutils** database when you upgrade to the latest version of the database server. All backup and restore information from the old database server version is lost. Backups that you make under the older version of the database server are not compatible with the newer version of the database server.

To prepare for an upgrade:

1. Use ON-Bar to perform a level-0 backup of all your data before you upgrade your database server, ISM, or change storage managers.
2. Save these backups so that you can restore the data in case you need to revert to the old database server version.
3. Before you upgrade, back up the administrative files.

4. After you upgrade the database server, back up all storage spaces and logical logs.

For complete information about database server migration, see the *IBM® Informix® Migration Guide*.

If you change storage manager vendors, do not remove the old storage manager until you verify that the new storage manager works for both backup and restore operations.

Related reference

[Upgrading a third-party storage manager on page 275](#)

[Changing storage-manager vendors on page 275](#)

Upgrading a third-party storage manager

If you upgrade a third-party storage manager vendors, do not remove the old storage manager until you verify that the new storage manager works for both backup and restore operations.

Prerequisite: Before upgrading, perform a complete backup of the database server.

To use a new version of a third-party storage manager with the database server:

1. Install the new storage manager before you bring up the database server.
2. Update the `sm_versions` file with the new storage-manager definition.

Verify that the new storage manager operates correctly before you use it in a production environment:

- Make sure that the storage manager can find the backup objects that ON-Bar requests.
- Make sure that the new storage-manager version is able to read media written with your old version.

If you have continuous logical-log backup set up on the database server, ON-Bar can start backing up the logical logs soon after the database server comes online.

Use the `onsmsync` utility to expire old backup history in the **sysutils** database and emergency boot files.

Related information

[Backing up before a database server or storage-manager upgrade on page 274](#)

Changing storage-manager vendors

If you change storage manager vendors, do not remove the old storage manager until you have proof that the new storage manager works for both backup and restore operations. You can use the old storage manager as a backup storage manager to use in case the new storage manager does not meet your needs.

ON-Bar supports working with multiple storage managers at the same time. To set up to test one storage manager and keep the other as a backup storage manager, specify information for both of the storage managers in the `BAR_BSALIB_PATH` configuration parameter and in the `$INFORMIXDIR/etc/sm_versions` file.

If you cannot use the old and new storage managers at the same time, use ON-Bar and the HCL Informix® Primary Storage ManagerISM or `ontape` as an alternative for backups while you check that backup and restore operations work correctly with the new storage manager. Until you confirm that the new storage manager works correctly, perform all your backups as whole system level-0 backups (`onbar -b -L 0 -w`)

If you change physical connectivity, such as moving a storage device from a local connection to a network server, make sure the new storage manager can move the data across the network. Also ensure that the new storage manager can send multiple data streams to storage devices. It also might use a different version of XBSA.

Related information

[Backing up before a database server or storage-manager upgrade on page 274](#)

Switching from `ontape` to ON-Bar

You cannot back up data with `ontape` and restore it with ON-Bar, or conversely because the data storage formats and backup capabilities are different. However, you can back up data with `ontape`, prepare to use ON-Bar, and then back up with ON-Bar.

About this task

To switch from `ontape` to ON-Bar:

1. Use `ontape` to perform a full backup.
2. Take the backup media offline to prevent possible reuse or erasure.
3. Configure the storage manager to be used with ON-Bar.
4. Configure your environment:
 - a. Set the configuration parameters that you use with ON-Bar and the storage manager.
 - b. Create the `sm_versions` file with the storage-manager definition. If you use a storage manager other than the HCL Informix® Primary Storage Manager, create the `sm_versions` file with the storage-manager definition. The Informix® Primary Storage Manager does not use the `sm_versions.std` file.
5. Use ON-Bar (`onbar -b` or `onbar -b -w`) to perform a full backup.
6. Verify the backup with the `onbar -v` command.

GLS support

This appendix contains information about using Global Language Support (GLS) with ON-Bar.

Use GLS with the ON-Bar utility

The ON-Bar utility supports Global Language Support (GLS), which allows users to work in their native language. The language that the client application uses is called the *client locale*. The language that the database uses for its server-specific files is called the *server locale*.

ON-Bar must run on the same computer as the database server. However, you can run ON-Bar in any locale for which you have the supporting message and globalization files. For example, if the server locale is English and the client locale is French, you can issue ON-Bar commands in French.

The following command performs a level-0 backup of the dbspaces specified in the file, `tomb`: `onbar -b -L 0 -f tomb`

On Windows™, you cannot use multibyte file names in backup or restore commands because they are not supported.

The **sysutils** database, the emergency boot files, and the storage-manager boot file are created with the `en_us.8859-1` (default English) locale. The ON-Bar catalog tables in the **sysutils** database are in English. Change the client and database locales to `en_us.8859-1` before you attempt to connect to the **sysutils** database with DB-Access or third-party utilities.

Identifiers that support non-ASCII characters

You can use non-ASCII characters in the database names and filenames with the ON-Bar and `ondblog` commands, and for file names in the `onconfig` file.

The *HCL® Informix® GLS User's Guide* describes the SQL identifiers that support non-ASCII characters. Non-ASCII characters include both 8-bit and multibyte characters.

For example, you can specify a non-ASCII file name for the ON-Bar activity login `BAR_ACT_LOG` and a non-ASCII path name for the storage-manager library in `BAR_BSALIB_PATH`.

Identifiers that require 7-bit ASCII characters

You must use 7-bit ASCII characters for storage space names and database server names.

Locale of ON-Bar messages

All ON-Bar messages appear in the activity log in the client locale except the messages that the database server issues.

For example, the part of the message that tells you that a database server error occurred appears in the client locale, and the server-generated part appears in the server locale.

Use the `GL_DATETIME` environment variable with ON-Bar

The database server must know how to interpret and convert the end-user formats when they appear in date or time data that the client application sends. You can use the **GL_DATE** and **GL_DATETIME** environment variables to specify alternative date and time formats.

If you do not set these environment variables, ON-Bar uses the date and time format of the client locale.

If you perform a point-in-time restore, enter the date and time in the format specified in the **GL_DATETIME** environment variable if it is set.

Use GLS with the `ontape` utility

The `ontape` utility supports GLS in the same way as ON-Bar does. You can specify the database name in the national locale.

Index

Special Characters

- /dev/null, ontape
 - as a tape device 137
 - for logical-log backup 156

A

- AC_CONFIG
 - environment variable 256
 - file 256
 - setting archecker parameters 205, 254
- AC_CONFIG environment variable 62, 257
- ac_config.std file 256, 257, 260
- AC_DEBUG configuration parameter 256
- AC_IXBAR configuration parameter 256
- AC_LTAPBLOCK configuration parameter 257
- AC_LTAPPEDEV configuration parameter 257
- AC_MSGPATH configuration parameter 257
- AC_SCHEMA configuration parameter 258
- AC_STORAGE configuration parameter 62, 258
- AC_TAPEBLOCK configuration parameter 259
- AC_TAPEDEV configuration parameter 259
- AC_TIMEOUT configuration parameter 260
- AC_VERBOSE configuration parameter 260
- Activity log, ON-Bar
 - defined 23
 - diagnosing failed verification 62
 - message numbers 124
 - monitoring backups 46
 - overview 23
 - return codes 124
 - specifying location 222
 - specifying performance monitoring 236
 - specifying performance statistics 39
 - specifying progress messages 39, 237
 - verification message 58
- Administrative files
 - backing up 45
 - copying before cold restore 75, 80
 - external backup 93, 95, 174, 175
- Administrative tasks
 - backing up after changing schema 16, 141
- ALARM_ALL_EVENTS configuration parameter 39
- ALARMPROGRAM configuration parameter
 - BAR_RETRY 39
 - BAR_XFER_BUF_SIZE 39
 - defined 39
 - LTAPPEDEV 39
 - RESTARTABLE_RESTORE 39
 - RESTORE_FILTER 39
- Altering table type
 - restoring data 16
- API
 - backup services 22
- Applier, manually controlling logical restore
 - using 207, 208
- archecker
 - AC_LTAPPEDEV configuration parameter 257
 - AC_TAPEDEV configuration parameter 259
- archecker utility
 - AC_CONFIG parameters 257
 - blobspaces 58
 - configuration parameters 254
 - defined 203
 - described 204
 - estimating temporary space 60
 - fixing verification problems 63

- message log 257
- messages 58, 257
- onsmsync 62
- point-in-time verification 58
- sbspaces 58
- sbspaces and blobspaces 60
- syntax
 - diagram 58, 208
 - elements 208
- table-level restore 210
- temporary files 60, 258
- uses 204
- using 208
- verification failure, causes of 61, 61
- whole-system backup 58
- ARCHIVE_UNLOGGED_OBJECTS configuration parameter 221
- Autochanger 25

B

- Back up
 - ON-Bar 43
 - Backing up compressed data 19
 - Backup activity, viewing for ON-Bar 53
 - Backup and restore a Remote Secondary Server(RSS) 179
 - Backup Services API 22
 - BACKUP_FILTER
 - about 19
 - configuration parameter 39, 132, 221
 - ON-Bar
 - utility
 - 221
 - ontape utility 132
 - Backups
 - advantages 144
 - catalogs
 - expiring and synchronizing 105
 - sysutils 105
 - directory 145
 - examples 148
 - expiring, on ISM 111
 - filled logical-log files 152
 - levels 4
 - oncheck, monitoring history 153
 - ontape utility
 - creating 143
 - IFX_ONTAPE_FILE_PREFIX environment variable 145
 - logical log filled 152
 - LTAPPEDEV configuration parameter 145
 - TAPEBLK configuration parameter 144
 - TAPPEDEV configuration parameter 145
 - TAPESIZE configuration parameter 144
 - overview 4
 - planning 15, 16
 - schedule 17, 27
 - Raw tables 149
 - setting TAPPEDEV to STDIO 144
 - standard output
 - defined 144
 - syntax 146
 - tapes 143
 - troubleshooting error messages 271
 - viewing error messages 23, 271
- Backups, ON-Bar
 - O option 52
- changing database logging 16
- checking data consistency 44
- current log 95, 175
- defined 46
- external 93, 174
- imported restore 83, 83
- ISM
 - catalog
 - 52
 - list of storage spaces 46, 51
 - logical log 5, 46
 - physical-only 46
 - salvaging logs 6
 - table types 16
 - verification 63
- Backups, ontape
 - before you create 142
 - interrupted 153
 - labelling the tape 144
 - levels 141
 - monitoring 153
 - one device 135
 - premature termination 153
- bar_act_log 23
- BAR_ACT_LOG configuration parameter 39, 222
- bar_act.log file 123
- bar_action table 116
- BAR_BSALIB_PATH configuration parameter 38, 224
- BAR_CKPTSEC_TIMEOUT configuration parameter 225
- BAR_DEBUG configuration parameter 39, 226
- BAR_DEBUG_LOG configuration parameter 39, 230
- BAR_DECRYPTION configuration parameter 229
- BAR_ENCRYPTION configuration parameter 228
- BAR_HISTORY configuration parameter 39, 231
- bar_instance table 116
- bar_ixbar table 118
- BAR_IXBAR_PATH configuration parameter 39, 231
- BAR_MAX_BACKUP configuration parameter 39, 232
- BAR_MAX_RESTORE configuration parameter 39, 234
- BAR_NB_XPORT_COUNT configuration parameter 39, 235
- bar_object table 120
- BAR_PERFORMANCE configuration parameter 39, 236
- BAR_PROGRESS_FREQ configuration parameter 39, 237
- BAR_RETRY configuration parameter 39, 237
- BAR_SEC_ALLOW_BACKUP configuration parameter 238
- bar_server table 121
- BAR_SIZE_FACTOR configuration parameter 239
- BAR_XFER_BUF_SIZE configuration parameter 39, 240
 - page size 240
- bargroup group 40
- batch files
 - onbar.bat 103

- bldutil.process_id file 41
- Blobspaces
 - availability for backing up 52, 154
 - backing up 4
 - backing up offline 52
 - optical platters 45
 - temp space for archecker 60
- Block size, ontape
 - parameter 137
- Blocking, database server 93, 174, 175
- Boot file
 - ixbar 23

C

- Catalog tables
 - ON-Bar utility 22
- Chunks
 - creating new chunks 169
 - files
 - recreating during restore 79
 - renaming during restore 66, 81
 - renaming 169
 - during a restore 66, 81
 - listfile 170
 - nonexistent device 82, 170
 - nonexistent device, procedure 170
 - overview with
 - ON-Bar
 - 66, 81
 - overview with ontape 168
 - specifying other options 170
 - with a file 170
- Client locale 276
- cloning 92, 173
- Cold restore
 - automatic log salvage 75
 - defined 7, 8
 - ON-Bar
 - utility
 - example 75
 - renaming chunks during 66, 81
 - ontape utility
 - defined 158
 - mixed restore 159
 - performing 164
 - renaming chunks during 168
- Complete restore
 - onbar -r commands 66
 - procedure 66
- Components
 - ON-Bar 20
- Compressing row data
 - effect on backup and restore 19
- Configuration file
 - AC_CONFIG 257
 - archecker 205
- configuration parameters
 - valid in archecker 254
- Configuration parameters
 - AC_MSGPATH 257
 - AC_STORAGE 258
 - AC_TIMEOUT 260
 - AC_VERBOSE 260
 - ALARMPROGRAM 39
 - ARCHIVE_UNLOGGED_OBJECTS 221
 - BACKUP_FILTER 39, 221
 - BAR_ACT_LOG 39, 222
 - BAR_BSALIB_PATH 38, 224
 - BAR_CKPTSEC_TIMEOUT 225
 - BAR_DEBUG 39, 226
 - BAR_DEBUG_LOG 39, 230

- BAR_DECRYPTION 229
- BAR_ENCRYPTION 228
- BAR_HISTORY 39, 231
- BAR_IXBAR_PATH 39, 231
- BAR_MAX_BACKUP 39, 232
- BAR_MAX_RESTORE 234
- BAR_NB_XPORT_COUNT 39, 235
- BAR_PERFORMANCE 39, 236
- BAR_PROGRESS_FREQ 39, 237
- BAR_RETRY 237
- BAR_SEC_ALLOW_BACKUP 238
- BAR_SIZE_FACTOR 239
- BAR_XFER_BUF_SIZE 240
- DBSERVERNAME 83, 121
- Environment variables
 - IFX_BAR_USE_DEDUP 39
- for
- Informix
 - Primary Storage Manager
 - 260
 - IFX_BAR_USE_DEDUP environment
 - variable 39
 - ISM_DATA_POOL 244
 - ISM_LOG_POOL 244
 - LTAPEBLK 244
 - LTAPEDEV 246
 - LTAPESIZE 247
 - OFF_RECVRVY_THREADS 66
 - ON_RECVRVY_THREADS 66
 - ontape 132
 - PSM_ACT_LOG 261
 - PSM_CATALOG_PATH 262
 - PSM_DBS_POOL 263
 - PSM_DEBUG 263
 - PSM_DEBUG_LOG 265
 - PSM_LOG_POOL 266
 - RESTARTABLE_RESTORE 87, 248
 - RESTORE_FILTER 249
 - SERVERNUM 83
 - TAPEBLK 250
 - TAPEDEV 251
 - TAPESIZE 253
- Configuring
 - ISM
 - 29
 - Spectrum Protect
 - 30
 - third-party storage manager 36
- Continuous log backup
 - specifying 5, 46
- Continuous log restore
 - configuring with ON-Bar 76
 - configuring with ontape 167
 - defined 10
 - versus (High Availability Replication (HDR) 10
- Cooked chunks
 - backing up 46
 - restoring 79
- Copying data 204
- CREATE EXTERNAL TABLE statement
 - archecker schema file 213
 - syntax 213
- CREATE TABLE statement, syntax 212
- Critical files
 - restoring 66
- cron command 24

D

- Data
 - filtering

- example, schema command file 218
- physical restores 218
- recovery
 - defined 3
 - usage 15
 - verifying consistency 44
- Database logging
 - backups 16
 - log backups 46
- Database logging status, changing with ontape 140
- Database servers
 - blocking 93, 174, 175
 - evaluating 16
 - imported restore 83
 - migrating
 - back up first 274
 - unblocking 94, 175
 - upgrading 83
 - back up first 274
- DATABASE statement
 - declaring logging mode 214
 - syntax 214
- DBSERVERNAME configuration parameter 121
- DBSERVERNAME Configuration parameter 83
- dbspaces
 - backing up 4
 - rename
 - cold restore 158
 - warm restore 158
 - restore selected, ontape 157
- Debugging
 - AC_DEBUG 256
 - BAR_DEBUG configuration parameter 226
 - BAR_DEBUG_LOG configuration
 - parameter 230
 - levels of messages in
 - ON-Bar
 - debug log
 - 226
 - messages 256
 - ON-Bar activity log 23
 - PSM_DEBUG configuration parameter 263
 - PSM_DEBUG_LOG configuration
 - parameter 265
 - with archecker 256
 - Deleting a bad backup, ON-Bar
 - 112
 - Device, ontape
 - /dev/null 156
 - logical-log backup 156
 - LTAPEDEV parameter 156
- Disaster recovery
 - imported restore 83
 - Spectrum Protect
 - administration services
 - 26
- Distributed restore
 - performing 219

E

- Emergency boot file
 - backing up 45
 - ixbar 23
 - regenerating 111
- en_us.8859-1 locale 276
- Encrypted storage spaces
 - restoring 66, 159
- Environment variables
 - AC_CONFIG 62, 256, 257

- GL_DATE 277
- GL_DATETIME 277
- IFX_BAR_NO_BSA_PROVIDER 242
- IFX_BAR_NO_LONG_BUFFERS 242
- IFX_BAR_USE_DEDUP 242
- IFX_ONTAPE_FILE_PREFIX 145
- IFX_TSM_OBJINFO_OFF 243
- INFORMIXSQLHOSTS 121
- TSM
 - setting the interface 34
 - setting the interface, example of 34
- Error messages
 - On-Bar activity log 271
 - troubleshooting backup problems 271
 - troubleshooting restore problems 271
- Errors
 - ON-Bar activity log 23
- Evaluating
 - backup and restore time 17
 - hardware and memory usage 17
 - logging and transaction activity 19
- Examples
 - storage manager 183
- Examples, CREATE TABLE statement 212
- Expiration policy 106
 - retention date 106
 - retention generation 106
 - retention interval 106
- Expiring backups
 - all 112
 - generation of 112
 - multiple Point-in-Time restores 112
 - on ISM 111
 - retention date 112
 - retention interval 112
 - using onsm sync 112
- External backup
 - blocking database server 94, 175
 - defined 93, 174
 - procedure 95, 175
 - RS secondary server 96
 - rules for 94
 - tracking backup objects 95, 176
 - unblocking database server 94, 175
- External backup and restore
 - cloning 92, 173
 - defined 92
 - disk mirroring 92, 173
 - initializing HDR 102
 - ON-Bar
 - 92
 - ontape 173
 - recovering data 92, 173
 - recovering data, procedure 92
- External programs
 - filters 19
 - transforming data with 19
- External restore
 - cold restore 176
 - cold restore procedure 100, 178
 - examples 101, 178
 - initializing HDR during 102, 178
 - ON-Bar
 - , renaming chunks 101
 - ontape utility
 - e command 176
 - l command 176
 - p command 176
 - performing 178
 - prerequisites 177
 - restored components 176
 - rules 177
 - salvaging logical logs 100
 - syntax diagram 98, 177
 - warm restore procedure 101
- External tables
 - CREATE EXTERNAL TABLE statement 213
 - restoring
 - schema command file example 219

F

- File directory, ontape
 - syntax to specify 136
- Files
 - ixbar 23
 - logical log 5
 - ON-Bar
 - boot 23
 - onbar.bat 24
 - Files for ON-Bar 41
 - Filters
 - for compression 19
 - transforming data 19
 - finderr utility 122

G

- GL_DATE environment variable 277
- GL_DATETIME environment variable 277
- Global Language Support (GLS) 276

H

- Hardware resources, evaluating 17
- HDR.
 - external backup and restore 102
- High-Availability Data Replication
 - imported restore 83
 - initializing 83
 - initializing with external restore 102, 178

I

- I/O
 - simultaneous backup and restore 173
 - environment variables 173
 - example 173
 - secondary host 173

IBM

- Informix Storage Manager (ISM)
 - backup requests 52
 - configuring 29
 - ISM
 - catalog 52
 - overview 25
 - requirements 25
 - IBM Spectrum Protect 26
 - IFX_BAR_NO_BSA_PROVIDER environment variable 242
 - IFX_BAR_NO_LONG_BUFFERS environment variable 242
 - IFX_BAR_USE_DEDUP environment variable 242
 - IFX_ONTAPE_FILE_PREFIX environment variable 145
 - IFX_TSM_OBJINFO_OFF environment variable 243
 - ifxbkpcld.jar utility 151
 - Imported restore 83

- ixbar.51 83
- ixbar.52 83
- performing 83
- Incremental backup
 - defined 15
- Informix Primary Storage Manager 180
 - activity log 261
 - catalog tables 262
 - configuration parameters 260
 - configuring 191
 - debug log 263, 265
 - device configuration 190
 - device pools 201
 - Device pools
 - for Informix Primary Storage Manager 201
 - device-configuration file 201
 - file-naming conventions 202
 - managing storage devices 191
 - message log 203
 - onpsm syntax 192
 - onpsm utility
 - managing storage devices 191
 - overview 181
 - setting up 190
 - viewing catalog details 198
 - viewing devices 199, 200
- INFORMIXSQLHOSTS environment variable 121
- INSERT statements
 - archecker syntax 215
 - physical-only restores 215
- ISM
 - environment variables 29
 - onconfig 29
 - sm_versions file 30
 - storage-manager definition 30
- ISM catalog
 - backing up 52
 - ism_catalog command 52, 52, 66, 66
 - ISM_DATA_POOL configuration parameter 244
 - ISM_LOG_POOL configuration parameter 244
 - ism_watch command 52, 66
 - ixbar boot file 23
 - IXBAR file, specifying location of 256

J

- Jukebox 25

L

- Level-0 archive, open transactions during 211
- Level-0 backup 4
- Level-1 backup 4
- Level-2 backup 4
- Locales
 - using GLS with ON-Bar 276
- Logging activity, evaluating 19
- Logging mode, declaring in DATABASE Statement 214
- Logical log
 - b option 55
 - l option 55
 - n option 55
 - P option 55
 - q option 55
 - t option 55
 - u option 55

- x option 55
- backing up files 167
- backup 46
 - O option 52
 - current log 95, 175
 - defined 5, 46
- files
 - restoring 166
 - restoring, examples of 166
- ON-Bar
 - utility
 - blobspace issues 52
 - checking available space 44
 - continuous backup 5, 46
 - manual backup 46
 - replaying records in parallel 66
 - salvaging 75
 - status of backup 46
 - when to back up 46
- onbar -P 55
- ontape utility
 - automatic backup, starting 155
 - backed-up status 154
 - backup, device to use 156
 - backup, if tape fills 152
 - backup, on another computer 136, 136
 - backup, procedure 153, 153
 - backup, separate devices 135
 - backup, to /dev/null 156
 - backup, when 154
 - blobspace blobs 154
 - continuous backup 155
 - importance of backing up 5
 - used status 154
- overview
 - defined 5
 - manual backup 5
 - salvaging 6
 - salvage, example of 165
 - viewing backed-up logs 55
- Logical restore 207
 - altering or adding tables during 207
 - defined 7
 - ontape utility
 - cold restore 158
 - warm restore 158
 - procedure 66
 - tables and fragments not processed by
 - applier during 211
- LTAPDEV configuration parameter
 - ontape utility
 - continuous backup to a directory 155
- LTAPEBLK configuration parameter 244
 - ontape utility 137
- LTAPEDEV configuration parameter 246
 - /dev/null 137
 - ON-Bar 39
 - ontape utility 133
 - changing to /dev/null 246
 - if two tape devices 152
 - purpose 156
- LTAPESIZE configuration parameter 133, 137, 138, 247

M

- Manual log backup
 - specifying 5, 46
- Memory resources, evaluating 17
- Message file 45
- Message log
 - archecker 211

- Migrating
 - data with archecker 204
- Migration
 - storage managers 275
- Mirroring configuration
 - backup state 164
 - setting 164
- Mixed restore
 - defined 7, 8
 - ontape utility 159
 - point-in-time 77
 - restoring data 77
 - strategies for using 78
- Moving data 204
- Multiple tables, restoring 219

O

- Offline storage spaces, restoring 66
- ON-Bar
 - activity log 114
 - activity, viewing 53
 - backup procedure 43
 - bar_act.log file 114
 - boot files 23
 - files 41
 - Performance statistics 114
 - preparing for back ups 44
 - registered backups 54
 - restore 64
 - Statistics
 - performance 114
 - verifying expired backups 63
- ON-Bar
 - activity log.
 - 66, 81
 - ON-Bar message log
 - format 123
 - ON-Bar return codes 124
 - ON-Bar tables
 - bar_action 116
 - bar_instance 118
 - bar_ixbar 120
 - bar_object 121
 - map 122
- ON-Bar
 - utility
 - O option
 - onbar 46
 - restore 52, 66, 66
 - whole-system restore 66
 - activity log 124
 - archecker
 - setting configuration parameters 254
 - backup and restore time 17
 - catalog tables 22
 - compared to ontape 11
 - components 20
 - configuration parameters 39, 248
 - configuration parameters for 221
 - deleting bad backups 112
 - external backup and restore 92
 - monitoring restores
 - onstat -d 65
 - operations supported 11
 - planning recovery strategy
 - creating backup plan 11
 - data loss 15
 - data usage 14, 14
 - failure severity 15
 - pre-recovery checklist 64
 - progress feedback 23
- restore
 - e option 66
 - f option 66
 - i option 66
 - l option 66
 - n option 66
 - O option 66
 - q option 66
 - r option 66
 - t option 66
 - T option 66
 - w option 66
 - storage management 25
 - switching from ontape 276
 - XBSA 22
- onbar -b 46
- onbar script 103
 - updating 104
 - usage and examples 116
- onbar syntax
 - RESTART 87
- onbar_m utility
 - defined 24
- onbar-driver
 - defined 24
- oncfg file 45
- Online storage spaces, restoring 66
- onmode -c
 - block command 94, 175
 - unblock command 94, 175
- onmode -ky command 75
- onmode -l command 52
- onpsm utility
 - C detail option 198
 - D list option 199
 - O list option 200
 - command reference 192
 - syntax 192
- onmsmsync utility
 - archecker 62
 - commands 107
 - syntax 107
 - using 105
- onmsmsync, expiring old backups 112
- onstat -d command 65
- onstat -l command 52
- ontape utility
 - d option 146
 - D option 11, 159
 - F option 146
 - FILE option 146, 159
 - L option 146, 159
 - s option 144
 - t option 146
 - v option 146
 - X option 159
 - backing up logical log 153
 - backup levels 141
 - backups and modes 143
 - changing
 - database logging status 140
 - LTAPEDEV to /dev/null 246
 - TAPEDEV to /dev/null 251
 - checking configuration parameters 138
 - compared to
 - ON-Bar
 - 11
 - configuration parameters 132
 - configuration parameters for 221
 - creating an archive 141
 - device name 204

- ensuring adequate memory 144
- example 259
- exit codes 148
- external backup and restore 173
- fake backup 146
- labelling backup tapes 143
- LTAPEBLK, uses 244
- LTAPEDEV, uses 246
- LTAPESIZE, uses 247
- migrating to ON-Bar 143
- option
 - r 146
- overview 176
- parameters
 - changing 146
 - overriding 146
- parameters, changing 138
- physical restore, choosing type 157
- precautions, one tape device 157
- premature backup termination 140
- read to end of tape 155
- restore, choosing mode 135, 158
- restoring data 158
- standard output 146
- starting continuous backup 153
- syntax
 - full backup 155
 - logical-log backup 139
- TAPEBLK, uses 250
- TAPEDEV, uses 251
- TAPESIZE, uses 253
- writing to end of tape 155

P

- Parallel backup
 - specifying 232
- Parallel restore
 - performing 211
 - specifying 234
- Performing backup 144
- Performing imported restore 83
- Physical backup 46
- Physical restore 7
 - ON-Bar
 - utility, procedure 157
 - ontape utility, cold restore 7
- Planning a backup system 16
- Point-in-log restore 66
- Point-in-time restore
 - example 66
 - mixed 77
- Primary Storage Manager 180
- PSM_ACT_LOG configuration parameter 261
- PSM_CATALOG_PATH configuration parameter 262
- PSM_DBS_POOL configuration parameter 263
- PSM_DEBUG configuration parameter 263
- PSM_DEBUG_LOG configuration parameter 265
- PSM_LOG_POOL configuration parameter 266

R

- Raw chunks
 - backing up 46
 - restoring 80
- Raw table
 - backing up
 - ontape utility 149
 - restoring 167
- Recovering data 206
- Recovery strategy planning
 - creating backup plan 15
 - data loss 14, 14
 - data usage 15
 - failure severity 14
- Recovery system
 - comparing
 - ON-Bar
 - and ontape 11
 - defined 3
 - Recovery, tables or fragments added or created during 211
- Recreating chunk files 79
- Remote device, ontape
 - interrupt key 156
 - syntax to specify 136
 - tape size 138
- Rename chunks restore
 - defined 158
 - ON-Bar
 - , external 101
 - ontape syntax 168
 - overview with ontape 168
- Replaying log records 6, 66
- Restartable restore
 - using 87
- RESTARTABLE_RESTORE configuration parameter 39, 87, 248
- Restore
 - Bringing the database back online 166
 - considerations 168
 - failed restore 90
 - resolving 90
 - scenarios 90
 - logical 207
 - logical, manually controlling 209
 - mounting tapes 165
 - multiple storage managers 210
 - physical 206
 - restarting 87
 - standard input 171
 - example 171
- Restore activity, viewing for ON-Bar 53
- restore error messages 271
- RESTORE statement, syntax 216
- RESTORE_FILTER
 - about 19
 - configuration parameter 39, 132, 249
 - ON-Bar
 - utility 249
 - ontape utility 132
- Restore, logical 216
- Restoring 7
 - critical files 66
 - ON-Bar
 - utility
 - O option 66, 66
 - cold, example 66

- cold, setting mode 75
- cooked chunks 79
- external 98, 101
- mixed 77
- offline storage spaces 66
- online storage spaces 75
- operational tables 66
- point-in-time example 66, 66
- raw tables 80
- renaming chunks 66, 81
- restartable 87
- selected spaces, example 87
- syntax 66
- table types 58
- ontape utility 159, 159, 159
 - selected storage spaces 157, 166
 - whole system, steps 162
- overview
 - defined 7, 162
 - logical phase 7, 7
 - physical phase 7
 - troubleshooting error messages 271
 - viewing error messages 23, 271
- Restoring compressed data 19
- Restoring data 66, 81, 204
 - archecker 206
 - mixed restore 77
 - point-in-time 206
- Return codes 124
- Rewindable tape device 17, 137
- Root dbspace
 - onbar -r command 124
 - when to back up 16
- RS secondary server
 - backup 96
- RS Secondary Server
 - external backup 96

S

- Salvaging logs
 - example 6
 - skipping 75
- Sample-code conventions 75
- Save sets
 - defined 52
- sbspaces
 - archecker 60
 - backing up 66
 - verifying 58
- Scheduling backups 16, 17
- Schema command file
 - example 218
 - restoring to a different a table 217
 - restoring to an external table 218, 219
 - simple 219
 - using data filtering 217
 - setting 217
 - Schema command file example
 - extracting a subset of columns 27
 - performing a distributed restore 218
 - restoring a table from previous backup 219
- Schema file
 - archecker 205
 - CREATE EXTERNAL TABLE statement 213
- Scratch table
 - backing up 205
- scripts
 - onbar 103
- Scripts
 - backup and restore 24
 - onbar 24

- Server locale 276
- SERVERNUM configuration parameter 83
- SET statement
 - examples 217
 - syntax 217
- Severity of data loss 217
- Shared libraries, XBSA
 - specifying location 38, 224
- Shared-memory parameters
 - setting to maximum assigned value 164
- Silos 38
- Skipping
 - log salvage 6
 - logical replay 75
- sm_versions file 28, 35
 - defining storage manager 45
- Smart large objects
 - Restoring 211
- Spectrum Protect
 - Administration services 26
 - Backup and restore services 26
 - client system options
 - dsm.sys 32
 - configuring 30
 - client options files 31
 - client system options 32
 - client user options 32
 - deduplication 36, 242
 - defined 26, 30
 - export backup objects 36
 - import backup objects 36
 - management class
 - assigning a backup 33
 - INCLUDE option 33
 - Managing resources 26
 - registering 34
 - replication 36, 243
 - transfer buffer 36
 - sqlhosts file
 - copying 45
 - Stager, manually controlling logical restore using 121, 207
 - Standard backup 207
 - Standard input
 - overriding TAPEDEV 159
 - restoring 171
 - example 171
 - single-user mode 171
 - setting TAPEDEV 137
 - simultaneous backup and restore 173
 - Standard output
 - backing up to 144
 - setting TAPEDEV 137
 - stdio TAPEDEV setting 137
 - storage catalog tables 181
 - Storage devices
 - ISM volume pools 52
 - Storage manager 183
 - communication with
 - ON-Bar
 - 28
 - migration 275
 - onbar -P
 - viewing backed-up logical logs 55
 - requirements 275
 - sm_versions file 36
 - Storage managers
 - ON-Bar 25
 - Storage spaces
 - backing up a list 46
 - backing up all 46, 51

- decrypted during restore 66, 159
- encrypting during restore 66, 159
- restartable restore 87
- restoring 7
 - when to back up 154
- Symbolic links
 - chunking names 168
 - ontape utility
 - specify tape devices 135
 - specify tape size 137
 - restoring chunk files 80
 - using with TAPEDEV configuration parameter 251
- Syntax diagrams
 - backup verification 58, 208
 - external restore 98, 175
 - onsmsync 107
 - restore 66
- System time changes, and backups 274
- sysutils database
 - regenerating 112

T

- Tables
 - altering
 - or adding during logical restore 211
 - restoring to user-specified point in time 211
- Tables and fragments
 - added or created during logical recovery 211
 - not processed by applier during logical restore 211
- Tape autochangers 25
- Tape block size
 - ON-Bar
 - utility
 - 207
 - ontape utility 257
- Tape device
 - ontape utility
 - before opening and on closing 137
 - block-size parameters 137, 257
 - gathering for cold restore 137
 - gathering for warm restore 137
 - parameters, setting 132
 - precautions with only one 135
 - reading to end 134
 - rewindable device 135
 - specifying symbolic links 137
 - using /dev/null 137
 - setting TAPEDEV to stdio 137
 - writing to end 135
- Tape device, block size 244
- Tape libraries 137
- TAPEBLK configuration parameter 250
 - ontape utility 137
 - defined 133
 - setting 133
- TAPEDEV configuration parameter
 - defined 251
 - ontape utility
 - changing to /dev/null 137
 - if two tape devices 152
 - using a symbolic link 251
- TAPESIZE configuration parameter 253
- Temp table
 - backing up 142
- Text editor, changing ontape parameters 138
- Third-party storage manager
 - configuring 36
 - functions 36

- onbar script 36
- Transaction activity
 - evaluating 19, 211
- Transactions
 - open during level-0 archive 27, 211
 - projecting with logical log backups 6
 - transforming data with filters 19
- Troubleshooting
 - backup error messages 271
 - failed backup verification 62
 - ON-Bar return codes 124
- TSM
 - deduplication 242, 242
 - environment variables
 - setting the interface 34
 - Informix
 - interface
 - initializing passwords 35
 - Specifying XBSA library location 38

U

- Unblocking database server 94, 94
- UNIX operating system
 - bargroup 40, 40
 - copying files 38, 40
- Upgrading database server 95
- Using CREATE EXTERNAL TABLE statement 213
- Using CREATE TABLE statement 212
- Using DATABASE statement 214
- Using mixed restore, strategies 77
- Utilities
 - archecker 78
 - onsmsync 62
- Utility
 - ifxbkpcload.jar 151

V

- Variables, in syntax diagrams 113
- Verifying raw devices 164
- Virtual processors,
 - ON-Bar
 - utility
 - 17
- Volume pools
 - backup locations 39

W

- Warm restore
 - defined 7, 8
 - ontape utility
 - critical dbspaces 167
 - defined 158, 167
 - mixed restore 159
 - part of a mixed restore 158
 - performing 163
 - steps to perform 159
 - overview 167
- Whole-system backup 46
- Whole-system restore
 - O option 66

X

- XBSA
 - ON-Bar utility interface 22
- XBSA shared library 181
 - default location 38
 - specifying location 38
- xcfg file 38