NetBackup™ for SAP Administrator's Guide

for UNIX, Windows, and Linux

Release 11.2

VERITAS
now part of Cohesity

NetBackup™ for SAP Administrator's Guide

Last updated: 2026-05-28

Legal Notice

Cohesity Support

Reach Cohesity Support

There are several ways to create a Cohesity support case.

- Go to Cohesity Support, to search in our knowledge base; or contact us by phone - United States and Canada: 1-855-9CO-HESI (926-4374), option 2.
- Log in to the Cohesity Support Portal to create a new case.
- Click the (**?**) icon on the Cohesity UI and select Support Portal.

Support/Service Assistance

First, contact the Service Provider that you have contracted for service and support. If you work directly with Cohesity and have a product warranty/entitlement, repair pricing, or technical support-related question, see your options below:

- To find solutions to your product issues or for suggestions or best practices, visit the Cohesity Knowledge Base.
- Log in to the Cohesity Support Portal to create a new case.
- To monitor your open cases, log in to the portal and click the **Cases** tab on the home page. This page should have all the case statuses and updates. You can also view individual case status.

Cohesity Software Running on Partner Hardware

For Cohesity software running on qualified third-party hardware, the following support workflow applies:

1. The customer may contact Cohesity Support first if the issue cannot be determined as a hardware issue.

   Note: Cohesity cannot process hardware replacement requests for partner hardware.

2. Cohesity Support triages the issue. If it is a software issue, Cohesity Support continues to work on it.
3. If it is a hardware/firmware issue or is suspected to be a hardware/firmware issue, Cohesity provides information about the issue to the customer and requests that the customer open a support ticket with the appropriate partner.
4. If needed, Cohesity Support can join a three-way call with the partner and the customer.
5. The customer informs Cohesity Support on the progress of the partner's case.

Contents

Chapter 1

Introduction to NetBackup for SAP

This chapter includes the following topics:

About NetBackup for SAP

NetBackup for SAP integrates the database backup and recovery capabilities of SAP with the backup and the recovery management capabilities of NetBackup.

NetBackup for SAP supports SAP environments based on the following databases:

- Oracle databases. In an Oracle environment, you can use the Oracle database either with or without the recovery manager (RMAN).
- SAP DB and MaxDB™ databases. This documentation describes using NetBackup for SAP with both SAP DB and MaxDB databases. For readability, the remainder of this manual uses only the term “MaxDB” to describe these database platforms.

Note: When necessary, differences are explained between using SAP on an Oracle database and SAP on a MaxDB database. For most tasks, NetBackup for SAP operates the same way in each environment.

Features of NetBackup for SAP

Table 1-1 shows NetBackup for SAP's main features and introduces some terms that are used in this documentation.

Table 1-1 NetBackup for SAP features and descriptions

Feature	Description
Media and device management	All the devices Media Manager supports are available to NetBackup for SAP.
Scheduling facilities	NetBackup scheduling facilities on the primary server can be used to schedule automatic and unattended SAP backups. This feature also lets you choose the times when these operations can occur. For example, to prevent interference with normal daytime operations, you can schedule your database backups to occur only at night.
Multiplexed backups and restores	NetBackup for SAP lets you take advantage of NetBackup's multiplexing capabilities. Multiplexing directs multiple data streams to one backup device, thereby reducing the time necessary to complete the operation.
Transparent SAP and regular file system backup and restore operations	All backups and restores run simultaneously and transparently without any action from the NetBackup administrator. The database administrator can run database backup and restore operations through NetBackup. An administrator or any other authorized user can use NetBackup to run database backups and restores.
Sharing the same storage units that are used for other file backups	Devices and media can be shared for other backups or to you can give SAP exclusive use of certain devices and media. NetBackup for SAP can use Media Manager, disk, and Media Server Deduplication Pool (MSDP) storage units.
Centralized and networked backup operations	From the NetBackup primary server, you can schedule database backups or start them manually for any client. The SAP databases can also reside on any hosts that are different from the devices on which NetBackup stores the backups.

Table 1-1 NetBackup for SAP features and descriptions *(continued)*

Feature	Description
User interfaces	NetBackup provides the NetBackup web UI for policy management and server-directed backups and restores.
Parallel backup and restore operations	NetBackup for SAP supports the parallel backup and restore capabilities of SAP Tools. For example, a user can run more than one tape device at a time for a single SAP backup or restore, thereby reducing the time necessary to complete the operation.
Compression	Compression increases backup performance over the network and reduces the size of the backup image that NetBackup writes to the storage unit. NetBackup for SAP supports compression for the following types of backups: ■ Stream-based backups through both MaxDB and RMAN. ■ File-based backups through `backint`. ■ Snapshot Client backups through `backint`. NetBackup for SAP does not support compression for Snapshot Client backups through RMAN.

Technical overview of NetBackup for SAP

Figure 1-1 shows the major components in a NetBackup for SAP configuration. The server that is hosting the Oracle or the MaxDB database must be a NetBackup client. It must also have NetBackup for SAP licensed.

Figure 1-1 NetBackup for SAP environment on an Oracle or MaxDB database

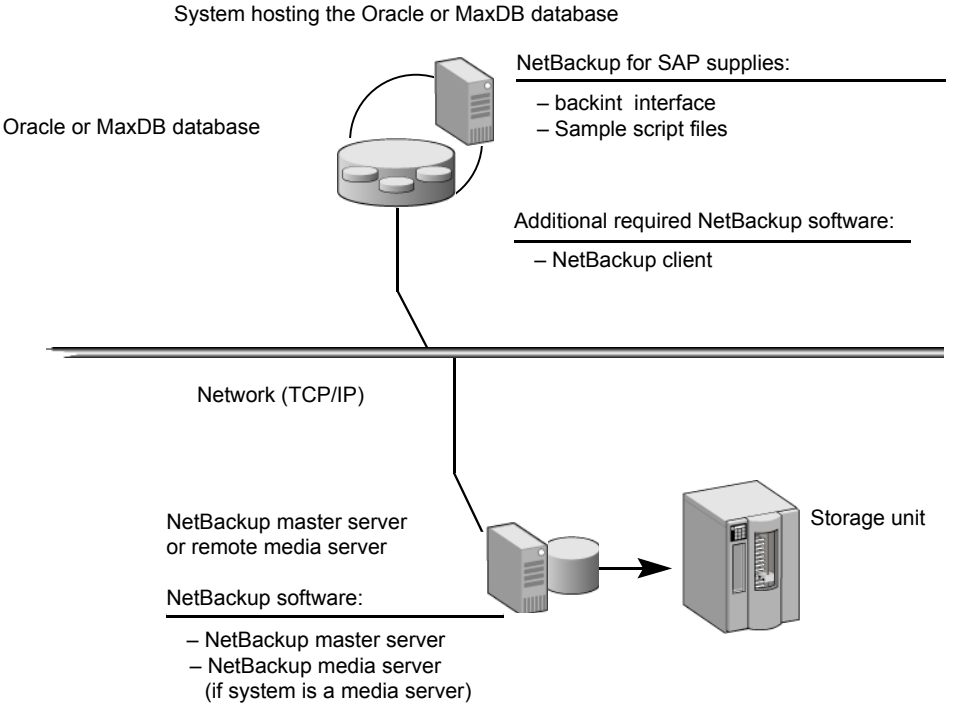

NetBackup for SAP assumes either an Oracle or a MaxDB database as the underlying database.

The following additional components are present depending on the underlying database:

- In an Oracle environment, the SAP tools perform database-related tasks. The `brbackup`, `brarchive`, and `brrestore` utilities communicate with NetBackup through the NetBackup for SAP `backint` interface. The `sapdba` component of SAP tools also communicates with NetBackup through the `backint` interface. The component queries the NetBackup image catalog and determines the status of backup images. The component also accesses the backup catalog used by NetBackup in order to determine the status of the valid backups.
- In a MaxDB environment, you initiate backups and restores through the MaxDB administrative interface. This interface can be the database manager graphical user interface (DBM GUI), the database manager command line interface (DBM CLI), and the Web database manager (Web DBM).

Sequence of operation for NetBackup for SAP

NetBackup users or schedules start database backups or restores by running the SAP script. If a policy on the primary server initiates the backup, a NetBackup process called `bphdb` starts the SAP script on the client. Alternatively a user or process on the client host can run the script. The script runs the SAP tools command line interface which, then starts the requested operation on the databases.

Figure 1-2 shows the sequence of operation.

Figure 1-2 NetBackup for SAP operations

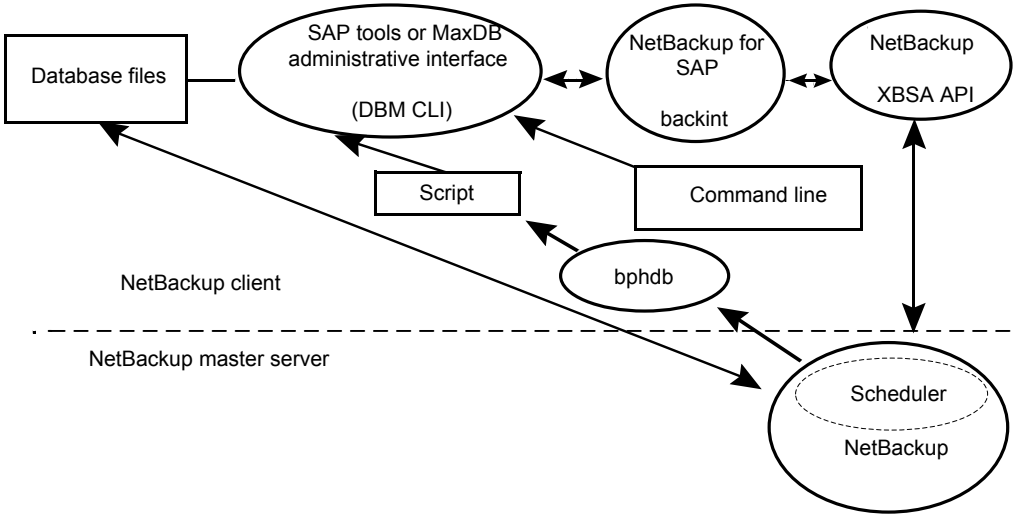

For a backup, `brbackup` calls the NetBackup for SAP `backint` interface. The `backint` interface calls the XBSA API to interface with NetBackup.

A restore works in a similar way as a backup. Except that the NetBackup for SAP `backint` interface calls the XBSA API. Therefore causes NetBackup to retrieve the data from secondary storage and send it to the client.

Using NetBackup for SAP on Oracle databases without RMAN

Figure 1-3 shows the components that are found in a NetBackup for SAP on Oracle database environment without RMAN.

Figure 1-3 Software components in a NetBackup for SAP on Oracle database environment

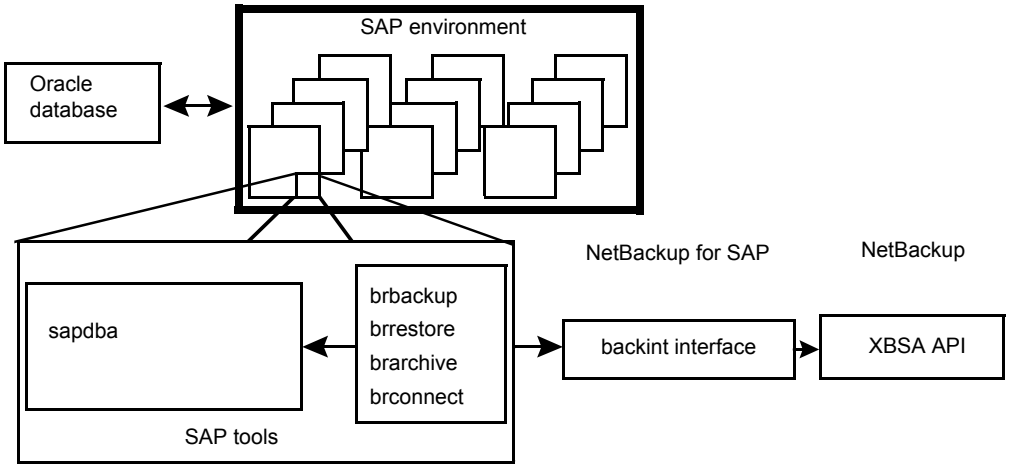

Using NetBackup for SAP on Oracle databases with RMAN

NetBackup for SAP implements the Oracle defined SBT interface, which integrates the SAP tools with RMAN.

You do not need to use RMAN to use NetBackup for SAP on Oracle databases, but the benefits of using RMAN backups are as follows:

- RMAN supports incremental backups, which are backups of only those blocks that have changed since the previous backup. This results in a reduced number of blocks being backed up. Unused and unchanged database blocks are not backed up.
- RMAN detects logical errors in database blocks during backup processing.
- The `BEGIN BACKUP` and `END BACKUP` commands are not needed for online backups because RMAN performs a block-by-block check to verify data consistency. This results in a reduced amount of redo log information.
- You can use RMAN's `verify` command to verify backups.

Figure 1-4 shows a NetBackup for SAP on Oracle database environment that includes RMAN.

Figure 1-4 Software components in a NetBackup for SAP on Oracle database environment with RMAN

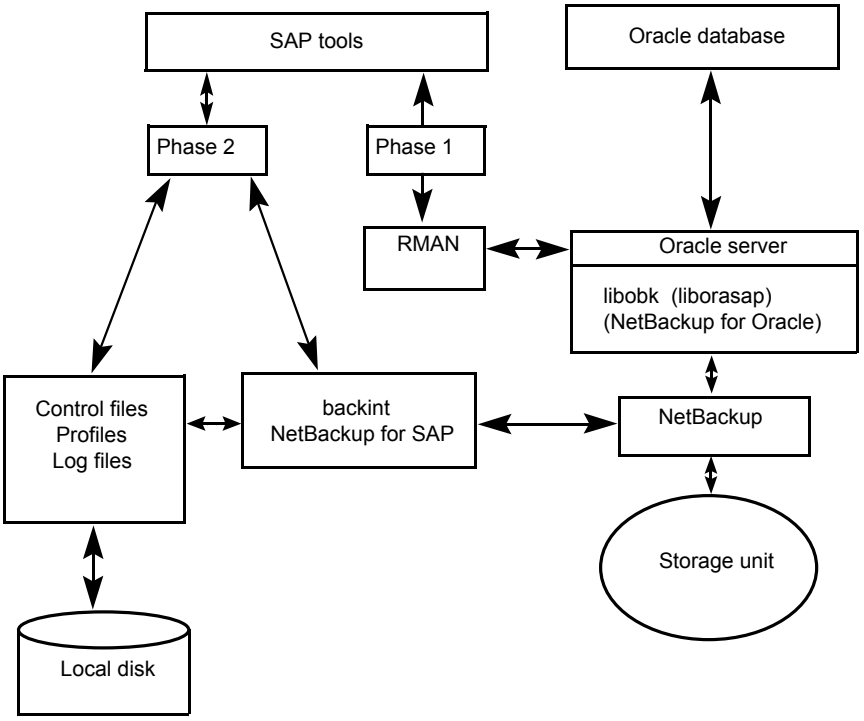

If you are in a NetBackup for SAP on Oracle database environment with RMAN, the backup consists of the following phases:

- Phase 1 backs up the Oracle database files.
- Phase 2 backs up the SAP control, log, and configuration files.

SAP tools overview

The SAP environment consists of many modules and applications. One component is the SAP tools. You can use the SAP tools whether or not you also use RMAN.

Table 1-2 lists the software that is included in the SAP tools.

Table 1-2 SAP tools software

Component	Function
`sapdba`	`sapdba` is a menu-driven utility, with menus designed to reflect the user's point of view. `sapdba` provides easy access to `brbackup`, `brarchive`, and `brrestore` for database backups and restores. The tool can restore a backup of an entire database or reset the database to a previous state.
`brbackup`	This command brings database servers online or offline, checks the status of SAP files, and places database tablespaces into `BACKUP` mode to guarantee their data consistency. The `brbackup` command provides online and offline backups. It also keeps a profile and log of each backup. `brbackup` uses the NetBackup software, through NetBackup for SAP, for the following actions: ■ SAP data file backups ■ Data file and online log backups ■ Error handling
`brarchive`	This command archives Oracle offline redo log files by communicating with the NetBackup for SAP `backint` interface. These files are copied by Oracle in its archiving directory. The `brarchive` command ensures that duplicates of these logs are available and that original files are not deleted prematurely. This command also keeps a profile and log of each archive.
`brrestore`	This command recovers database data files, control files, and online redo log files through the NetBackup for SAP `backint` interface. The `brrestore` command ensures that sufficient space is available before restoring these files, and it removes the files that are overwritten during the recovery. This command also provides a query mode.
SAP script	This script is a small script that contains SAP commands such as `brbackup` and `brrestore`.

NetBackup for SAP on MaxDB databases

Figure 1-5 shows the components that are found in a NetBackup for SAP on MaxDB database environment.

Figure 1-5 Software components for a NetBackup for SAP on MaxDB database environment

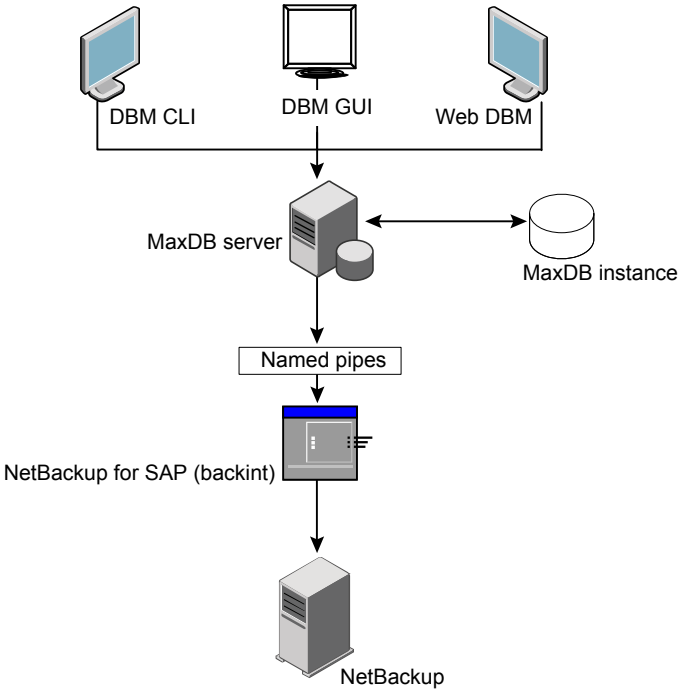

Chapter 2

Introduction to NetBackup SAP HANA

This chapter includes the following topics:

About NetBackup for SAP HANA

The NetBackup for SAP HANA Agent integrates the `backint` interface for SAP HANA along with the backup and the recovery management capabilities of NetBackup.

The software works in single as well as multi-node environments.

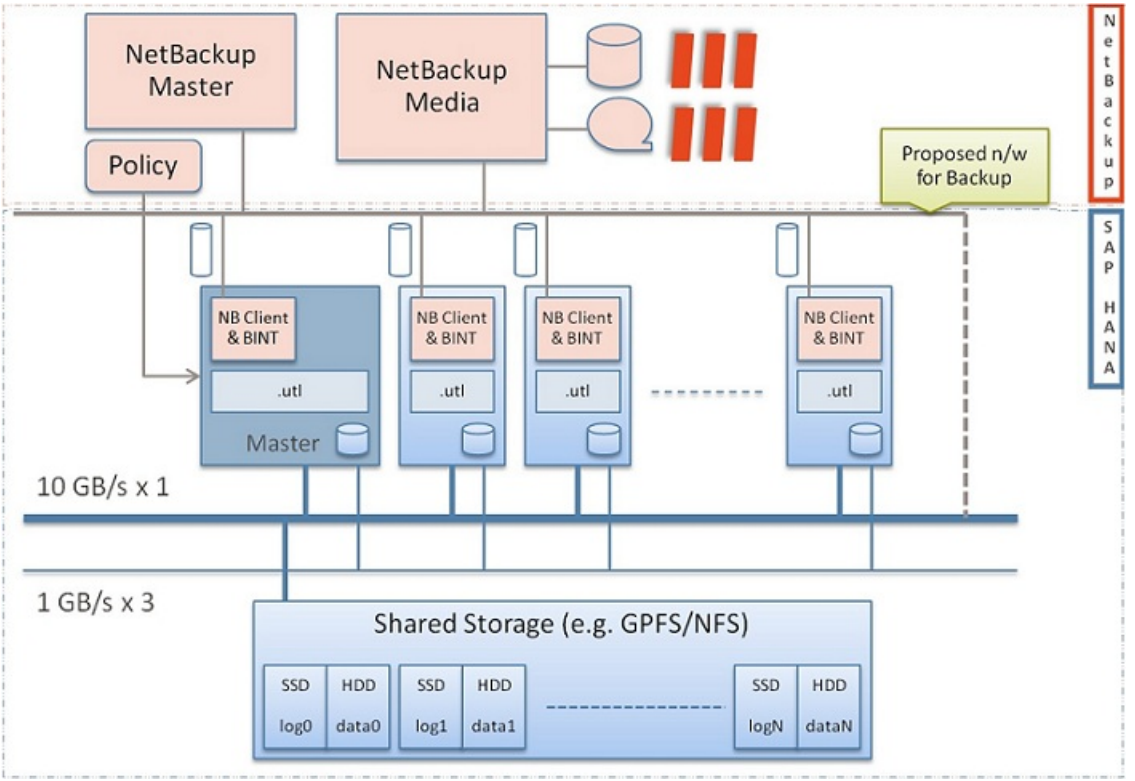

NetBackup for SAP HANA databases

The NetBackup HANA agent implementation should work on all vendor HANA appliances that have HANA version 5, revision 46, and later versions. For Hitachi and IBM appliance there is a need to add an extra library `libelf.so.0`.

For the current support, a backup and a restore for HANA are initiated from the SAP HANA studio and not from NetBackup web UI.

SAP HANA features

Table 2-1 SAP HANA features

Feature	Description
NetBackup integration	Full integration with the NetBackup primary server and media manager. Job monitoring from the primary server.
Integration with NetBackup role-based access control (RBAC)	The NetBackup web UI provides RBAC roles to control which NetBackup users can manage SAP HANA operations in NetBackup. The user does not need to be a NetBackup administrator to manage SAP HANA operations.
SAP HANA Intelligent Policy	The following benefits are included: ■ Use an intelligent policy to protect the following: Multiple SAP HANA instances or instance databases. Instances can be spread over multiple clients. ■ Include a full, differential, cumulative incremental, and log backup in the same policy. ■ You are not required to know SAP HANA commands or to write and use batch files.
Management of SAP HANA assets	NetBackup automatically discovers SAP HANA instances and databases in the environment. You can also perform manual discovery. After instances are registered, the SAP HANA workload administrator can protect the SAP HANA assets with a policy.

Table 2-1 SAP HANA features *(continued)*

Feature	Description

Authentication and credentials	SAP HANA Intelligent Policy supports the following database users created in the SYSTEMDB:
	■ SYSTEM
	■ Backup admin database user and should be granted the following privileges:
	  ■ BACKUP ADMIN
	  ■ DATABASE BACKUP ADMIN
	  ■ CATALOG READ
	■ Backup operator database user and should be granted the following privileges:
	  ■ BACKUP OPERATOR
	  ■ CATALOG READ

Backup and restore features	The following features are available for backups and restores with the NetBackup web UI: ■ Backups are managed entirely by the NetBackup server from a central location. Administrators can schedule automatic, unattended backups for instances on local or remote hosts across the network. ■ NetBackup supports the backup of databases and logs. ■ Backup schedules for full, differential, cumulative incremental, or log backups. ■ Manual backups. ■ Support for high availability (HA) environments, including SAP HANA clusters. ■ Ability to use multiple streams during a backup. ■ Configure tuning options that can improve the performance of backups.
Stream-based backups and restores	Stream-based backup and restore of SAP HANA objects.

Chapter 3

Installing NetBackup for SAP

This chapter includes the following topics:

Planning the installation of NetBackup for SAP

Table 3-1 shows the major installation steps that are needed to run NetBackup for SAP.

Table 3-1 Installation steps for NetBackup for SAP

Step	Action	Description
Step 1	Verify the installation prerequisites.	See "Verifying the operating system and platform compatibility " on page 27. See "NetBackup server and client requirements" on page 28. See "SAP server software requirements " on page 28. See "Configuration and licensing requirements for NetBackup for SAP with Snapshot Client" on page 29.
Step 2	Verify that the primary server has a valid license for NetBackup for SAP and any NetBackup options or add-ons.	See "License for NetBackup for SAP" on page 29.
Step 3	Perform the appropriate linking procedure. OR Copy the NetBackup for SAP binary file.	See "Linking NetBackup for SAP with backint (UNIX or Linux Oracle databases without RMAN only)" on page 30. See "Linking the Oracle instances with the SBT library (UNIX or Linux Oracle databases with RMAN only)" on page 31. See "Copying the SAP binary file (Windows Oracle databases without RMAN only)" on page 30.

Verifying the operating system and platform compatibility

Verify that the NetBackup for SAP agent is supported on your operating system or platform.

To verify operating system and compatibility

1. Go to the NetBackup compatibility list site.

   https://support.cohesity.com/s/article/article-100040093

2. Select the link for the following document:

   Application/Database Agent Compatibility List

3. For information on support for Snapshot Client, see the Snapshot Client section in the following document:

   Hardware and Cloud Storage Compatibility List (HCL).

NetBackup server and client requirements

Before you install NetBackup, review the requirements for the NetBackup server and the NetBackup clients.

NetBackup server requirements

Verify that the following requirements are met for the NetBackup server:

- The NetBackup server software is installed and operational on the NetBackup server.
  See the NetBackup Installation Guide.
- Make sure that you configure any backup media that the storage unit uses. The number of media volumes that are required depends on several things:
  - The devices that are used and the storage capacity of the media.
  - The sizes of the databases that you want to back up.
  - The amount of data that you want to archive.
  - The size of your backups.
  - The frequency of backups or archives.
  - The length of retention of the backup images.

  See the NetBackup Administrator's Guide, Volume I.

NetBackup client requirements

Verify that the following requirements are met for the NetBackup clients:

- The NetBackup client software is installed on the computer that has the databases you want to back up.
- If the database is clustered, you must use the same version of NetBackup on each node in the cluster.
- To use the new features that are included in NetBackup for SAP in NetBackup 11.2, you must upgrade your NetBackup for SAP clients to NetBackup 11.2. The NetBackup media server must use the same version as the NetBackup for SAP client or a higher version than the client.

SAP server software requirements

Verify the following regarding the SAP server software on the NetBackup server or client:

- The SAP server software must be installed and operational.
  Refer to the Application/Database Agent Compatibility List for supported versions of the SAP server software.

See "NetBackup server and client requirements" on page 28.

Requirements for using NetBackup for SAP MaxDB, SAP Oracle, and SAP HANA in a NetBackup cluster

If you plan to use NetBackup for SAP on a NetBackup server configured in a NetBackup cluster, verify the following requirements:

- NetBackup supports your cluster environment.
  the Software Compatibility List (SCL).
- The NetBackup server software is installed and configured to work in a NetBackup cluster.
  See the NetBackup Installation Guide.
  See the NetBackup Clustered Primary Server Administrator's Guide.
- The NetBackup client software is installed and operational on each node to which NetBackup can failover.
- A valid license for NetBackup for SAP must exist on each node where NetBackup server resides.

Configuration and licensing requirements for NetBackup for SAP with Snapshot Client

To use NetBackup for SAP with Snapshot Client, you must have a valid license for NetBackup Snapshot Client on the primary server.

License for NetBackup for SAP

The NetBackup for SAP agent is installed with the NetBackup client software. No separate installation is required. A valid license for the agent must exist on the primary server.

More information is available on how to add licenses.

See the NetBackup Web UI Administrator's Guide.

Copying the SAP binary file (Windows Oracle databases without RMAN only)

Perform the following procedures only if you use NetBackup for SAP with an Oracle database but without the Oracle recovery manager (RMAN).

SAP requires that all SAP tools be located in a predetermined directory. The directory path is as follows:

`C:\usr\sap\`*`SID`*`\SYS\exe\run`

where *SID* is the unique name for an Oracle database instance. SID is also known as the System ID.

The directory should contain the following commands:

- `brarchive`
- `brbackup`
- `brconnect`
- `brrestore`
- `brtools`
- `sapdba`

To copy the NetBackup for SAP binary file (Windows Oracle databases without RMAN only)

- Copy `backint.exe` from the NetBackup install directory to the SAP tools directory.

  For example, if the Oracle instance name is CER, enter the following command:

  ```
  copy install_path\NetBackup\bin\backint.exe
  c:\usr\sap\CER\sys\exe\run
  ```

Linking NetBackup for SAP with backint (UNIX or Linux Oracle databases without RMAN only)

Perform the following procedures only if you use NetBackup for SAP with an Oracle database but without the Oracle recovery manager (RMAN).

SAP requires that all SAP tools be located in a predetermined directory. The directory path is as follows:

`/usr/sap/`*`SID`*`/SYS/exe/run`

where *SID* is the unique name for an Oracle database instance. SID is also known as the System ID.

The directory should contain the following commands:

- `brarchive`
- `brbackup`
- `brconnect`
- `brrestore`
- `brtools`
- `sapdba`

To link the NetBackup for SAP binary file (UNIX or Linux Oracle databases without RMAN only)

- Link `backint` from the NetBackup install directory to the SAP tools directory.

```
ln -s /usr/openv/netbackup/bin/backint \
/usr/sap/CER/SYS/exe/run/backint
```

Note: Never copy the file from one directory to another. Whenever NetBackup is upgraded it replaces `/usr/openv/netbackup/bin/backint` with a new version. SAP automatically uses the updated file by the symbolic link. Copying a file can give errors.

Linking the Oracle instances with the SBT library (UNIX or Linux Oracle databases with RMAN only)

Perform the linking procedures that are described in this topic only if you use NetBackup for SAP with an Oracle database and RMAN. The exact linking command to use depends on your operating system platform.

Note: You must link Oracle with the SBT library for all Oracle instances on the host that are planned to be backed up through NetBackup for SAP using RMAN. Whenever a new Oracle instance is added it must also be similarly linked.

Link the library with NetBackup when:

- Using NetBackup for SAP for the first time.

- Adding a new SAP instance that should be backed up with NetBackup for SAP with RMAN.

To link the NetBackup for SAP binary file on AIX (64-bit) and 64-bit Oracle9i R2 or later software

- Enter the following command to link the Oracle database and the SBT library to NetBackup:

```
ln -s /usr/openv/netbackup/bin/libsapora.a64 \
$ORACLE_HOME/lib/libobk.a
```

To link the NetBackup for SAP binary file on HP-Itanium and 64-bit Oracle9i R2 or later software

- Enter the following command to link the Oracle database and the SBT library to NetBackup:

```
ln -s /usr/openv/netbackup/bin/libsapora.so \
$ORACLE_HOME/lib/libobk.so
```

To link the NetBackup for SAP binary file on HP PA-RISC (64-bit) and 64-bit Oracle9i R2 or later software

- Enter the following command to link the Oracle database and the SBT library to NetBackup:

```
ln -s /usr/openv/netbackup/bin/libsapora.sl64 \
$ORACLE_HOME/lib/libobk.sl
```

To link the NetBackup for SAP binary file on Linux IBMpSeries or Linux Itanium and 64-bit Oracle9i R2 or later software

- Enter the following command to link the Oracle database and the SBT library to NetBackup:

```
ln -s /usr/openv/netbackup/bin/libsapora.so \
$ORACLE_HOME/lib/libobk.so
```

To link the NetBackup for SAP binary file on Linux (64-bit) and 64-bit Oracle9i R2 or later software

- Enter the following command to link the Oracle database and the SBT library to NetBackup:

```
ln -s /usr/openv/netbackup/bin/libsapora.so64 \
$ORACLE_HOME/lib/libobk.so
```

To link the NetBackup for SAP binary file on Solaris SPARC (64-bit) and 64-bit Oracle 9i R2 or later software

- Enter the following command to link the Oracle database and the SBT library to NetBackup:

```
ln -s /usr/openv/netbackup/bin/libsapora.so64.1 \
$ORACLE_HOME/lib/libobk.so
```

To link the NetBackup for SAP binary file on Solaris Opteron and 64-bit Oracle 10g R2 or later supported versions

- Enter the following command to link the Oracle database and the SBT library to NetBackup:

```
ln -s /usr/openv/netbackup/bin/libsapora.so.1 \
$ORACLE_HOME/lib/libobk.so
```

Chapter 4

Installing NetBackup for SAP HANA

This chapter includes the following topics:

Planning the installation

Table 4-1 shows the major installation steps that are needed to run NetBackup for SAP HANA. Each step contains one or more links to pertinent procedures and concepts.

For details about installing the media server and the primary server, refer to the NetBackup Installation Guide.

Table 4-1 Installation steps for NetBackup for SAP HANA

Step	Action	Description
Step 1	Verify the installation prerequisites.	See "Verifying the operating system and platform compatibility for NetBackup for SAP HANA" on page 35. See "NetBackup server and client requirements for NetBackup for SAP HANA" on page 35. See "Software requirements for NetBackup for SAP HANA" on page 36.
Step 2	Add the license key for NetBackup for SAP.	See "License for NetBackup for SAP" on page 29.

Verifying the operating system and platform compatibility for NetBackup for SAP HANA

Verify that the NetBackup for SAP HANA agent is supported on your operating system or platform.

To verify operating system and compatibility

1. Go to the following webpage:

   https://support.cohesity.com/s/article/article-100040093

2. In the list of documents, click on the following document:

   Application/Database Agent Compatibility List

NetBackup server and client requirements for NetBackup for SAP HANA

Verify that the following requirements are met for the NetBackup server:

- The NetBackup server software is installed and operational on the NetBackup server. The NetBackup server platform can be any that NetBackup supports. See the NetBackup Installation Guide.
  Make sure that you configure any backup media that the storage unit uses. The number of media volumes that are required depends on several things:

- The devices that are used and storage capacity of the media
- The sizes of the databases that you want to back up
- The amount of data that you want to archive
- The size of your backups
- The length of retention of the backup images.
- The frequency of backups or archives
  See the NetBackup Administrator's Guide, Volume I.

Verify that the following requirements are met for the NetBackup clients:

- The NetBackup client software is installed on the computer that has the databases you want to back up.
- Check that `libelf.so.0` is symbolically linked to the `libelf.so.0.8.12` file. These files are located in `/usr/lib64/`.

Software requirements for NetBackup for SAP HANA

Verify the following regarding the SAP HANA software on the NetBackup client:

- SAP HANA software must be installed and operational on a single host, multiple hosts and cluster of servers that supports active-active or active-passive high availability failover configurations.
- One or more SAP HANA instances must exist.

Reviewing the auto-discovered mappings for SAP HANA

In certain scenarios, a NetBackup host shares a particular name with other hosts or has a name that is associated with a cluster. To successfully perform backups and restores with SAP HANA, you must approve each valid auto-discovered mapping that NetBackup discovers in your environment. Or, manually add the mappings.

Examples of the configurations that have multiple host names include:

- A host is associated with its fully qualified domain name (FQDN) and its short name or its IP address.
- If the SAP HANA server is clustered, the host is associated with its node name and the virtual name of the cluster.

- If the SAP HANA server is clustered, the host is associated with its node name and the virtual name of the cluster.

These mappings are configured in the **Security > Host mappings** node in the NetBackup web UI. You can also use the `nbhostmgmt` command to manage the mappings.

Auto-discovered mappings for a cluster

In SAP HANA cluster environment, you must map the node names to the virtual name of the cluster if the following apply:

- If the backup policy includes the cluster name (or virtual name)
- If the NetBackup client is installed on more than one node in the cluster, the virtual name must be mapped to each node. If the NetBackup Client is only installed on one node, then no mapping is necessary.

Approve the auto-discovered mappings for a cluster

To approve the auto-discovered mappings for a cluster

1. In the NetBackup web UI, expand **Security > Host mappings**.
2. Click the **Mappings to approve** tab.

   The list displays the hosts in your environment and the mappings or additional host names that NetBackup discovered for those hosts. A host has one entry for each mapping or name that is associated with it.

   For example, for a cluster with hosts `client01.lab04.com` and `client02.lab04.com`, you may see the following entries:

Host	Auto-discovered mapping
client01.lab04.com	client01
client01.lab04.com	clustername
client01.lab04.com	clustername.lab04.com
client02.lab04.com	client02
client02.lab04.com	clustername
client02.lab04.com	clustername.lab04.com

3. Click the name of the host.

4 Review the mappings for the host and click **Approve** if you want to use the discovered mappings.

For example, if the following mappings are valid for `client01.lab04.com`, then you approve them.

Auto-discovered mapping	Valid name for
client01	The short name of the client
clustername	The virtual name of the cluster
clustername.lab04.com	The FQDN of the virtual name of the cluster

5 When you finish approving the valid mappings for the hosts, click on the Hosts tab.

For hosts `client01.lab04.com` and `client02.lab04.com`, you see entries for Mapped host or IP address that are similar to the following:

Host	Mapped host names/IP addresses
client01	The short name of the client
clustername	The virtual name of the cluster
clustername.lab04.com	The FQDN of the virtual name of the cluster

Auto-discovered mappings for a SAP HANA server cluster in a multiple NIC environment

If you have SAP HANA server cluster in a multi-NIC environment, you need to approve each valid auto-discovered mapping for the hosts in that environment. You must map the virtual name of the SAP HANA server cluster on the private network to the private name of each SAP HANA server cluster node.

To approve the auto-discovered mappings for SAP HANA server cluster in a multiple NIC environment

1 In the NetBackup web UI, expand **Security > Host mappings**.

2 Click the **Mappings to approve** tab.

3 The list displays the hosts in your environment and the mappings or additional host names that NetBackup discovered for those hosts. A host has one entry for each mapping or name that is associated with it.

For example, for a cluster in a multi-NIC environment with hosts `client01-bk.lab04.com` and `client02-bk.lab04.com`, you may see the following entries:

Host	Auto-discovered mapping
client01-bk.lab04.com	clustername-bk.lab04.com
client02-bk.lab04.com	clustername-bk.lab04.com

4 Review the mappings for the host and click Approve if you want to use the discovered mappings. For example, if following mapping is valid for `client01-bk.lab04.com`, then you approve it.

Auto-discovered mapping	Valid name for
clustername-bk.lab04.com	The virtual name of the SAP HANA cluster on the private network

5 When you finish approving the valid mappings for the hosts, click on the Hosts tab.

For hosts `client01-bk.lab04.com` and `client02-bk.lab04.com`, you may see the following Mapped host or IP address.

Host	Mapped host or IP address
client01-bk.lab04.com	clustername-bk.lab04.com
client02-bk.lab04.com	clustername-bk.lab04.com

6 If you need to add a mapping that NetBackup did not automatically discover, you can add it manually.

Example mapped host names for SAP HANA server cluster in a multi-NIC environment

Table 4-2 Example mapped host names for SAP HANA server cluster in a multi-NIC environment

Host	Mapped host names
Private name of *Node 1*	Virtual name of the SAP HANA cluster on the private network
Private name of *Node 2*	Virtual name of the SAP HANA cluster on the private network

Manually map host names

If you need to add a mapping that NetBackup did not automatically discover, you can add it manually.

To manually map host names

1. In the NetBackup web UI, expand **Security > Host mappings**.
2. Click on the **Hosts** tab.
3. Click **Add shared or cluster mappings**.

   For example, type the name of the virtual name of the cluster. Then click **Add** to choose the hosts to which you want to map that virtual name.

Configure the number of jobs allowed for backup operations for SAP HANA

When NetBackup starts a backup of SAP HANA, a number of jobs are created. Depending on the policy configuration, additional jobs are created if you configure settings such as **Number of backup streams** and **Parallel backup operations**.

You can increase or limit the number of jobs that are created. You can also control the number of jobs that are sent to the storage unit.

Limit jobs per policy
: Sets the maximum number of instances that NetBackup can back up concurrently in each policy. This setting is configured in the policy attributes.

Maximum jobs per client
: In a policy, the maximum number of jobs per client that you want to allow. This setting applies to all clients in all policies. You can configure this property in the web UI in the Global attributes host properties for the primary server.

Maximum concurrent jobs	The maximum number of jobs that NetBackup can send to a storage unit at one time. This setting is configured in the storage unit properties.
Maximum concurrent write drives	The number of tape drives that NetBackup can use at one time for jobs to this storage unit. This setting is configured in the storage unit properties.

Configure the Maximum jobs per client setting for SAP HANA

The **Maximum jobs per client** specifies the maximum number of concurrent backups that are allowed per instance or database. Each instance or database that is specified in the policy creates a new backup job.

To configure the maximum jobs per client

1. On the left, select **Hosts > Host properties**.
2. Select the primary server.
3. If necessary, click **Connect**. Then click **Edit primary server**.
4. Click **Global attributes**.
5. Change the **Maximum jobs per client** value to the required value.

   The default is 1.

For Intelligent Policies, use the following formula to calculate a smaller value for the Maximum jobs per client setting:

Maximum jobs per client = *number_of_database_objects X number_of_streams X number_of_policies*

For batch file-based policies, use the following formula to calculate a smaller value for the **Maximum jobs per client** setting:

Maximum jobs per client = *number_of_streams X number_of_policies*

Refer to the following definitions:

number of database_objects	For intelligent policies, this number is the number of databases, filegroups, or files that you want to back up in parallel. For batch file-based policies, this number is the number of databases, filegroups, or files that you want to back up in parallel.

number_of_streams	The number of backup streams between the database server and NetBackup.
number_of_policies	The number of policies of any type that can back up this client at the same time. This number can be greater than one. For example, a client can be in two policies to back up two different databases. These backup windows can overlap.

Chapter 5

Configuring NetBackup for SAP

This chapter includes the following topics:

About configuring NetBackup for SAP

Before attempting to configure NetBackup for SAP, verify that your environment meets the prerequisites and perform any other configuration that is required.

After you complete the installation, follow the procedures in to configure your environment.

Note: If you use NetBackup for SAP on an Oracle database, all the configuration topics apply. If you use NetBackup for SAP on a MaxDB database, differences are noted in these topics. Also see the information on how to use NetBackup for SAP on MaxDB databases.

See "About NetBackup for SAP on MaxDB databases" on page 188.

Table 5-1 Steps to configure NetBackup for SAP

Step	Action	Description
Step 1	Configuring a backup policy.	See "About configuring a backup policy for SAP " on page 44.
Step 2	Create backup scripts.	See "NetBackup for SAP backup scripts" on page 53.
Step 3	(Windows) Configure the logon account for the NetBackup Client Service	See "Configuring the logon account for the NetBackup Client Service for NetBackup for SAP " on page 62.
Step 4	Edit SAP configuration files.	See "About SAP configuration files" on page 62.
Step 5	(UNIX or Linux) Configure the backups of files that are mounted with LOFS.	See "Backing up files mounted with LOFS with NetBackup for SAP (UNIX or Linux)" on page 72.
Step 6	(UNIX and Linux) Configure the permissions for log files.	See "About permissions for NetBackup for SAP log files (UNIX)" on page 72.
Step 7	Test the configuration settings.	See "Perform a manual backup" on page 77.
Step 8	Test multiple drives.	See "Testing multiple concurrent backup jobs for NetBackup for SAP for Oracle (without RMAN)" on page 77.

About configuring a backup policy for SAP

A backup policy defines the backup criteria for a specific group of one or more clients.

These criteria include the following:

- Storage unit and media to use
- Policy attributes
- Backup schedules
- The clients to back up
- The script files to run on the clients

To back up the database environment, define at least one SAP policy with the appropriate schedules. A configuration can have a single policy that includes all clients, or there can be many policies, some of which include only one client.

See "Add a policy for SAP" on page 45.

Add a policy for SAP

This topic describes how to create a policy to protect a database.

To add a policy for SAP

1. Open the NetBackup web UI.
2. On the left, select **Protection > Policies**. Then select **Add**.
3. Type a unique name for the new policy.
4. From the **Policy type** list, select **SAP**.
5. Complete the entries on the **Attributes** tab.

   See "About policy attributes" on page 46.
6. Add other policy information as follows:
   - Add schedules.
     See "Configure an application backup schedule" on page 47.
     See "Configure automatic backup schedules" on page 48.
   - Add clients.
     See "Add clients to a policy" on page 50.
   - Add scripts to the backup selections list.
     See "Adding NetBackup for SAP scripts to the backup selections list" on page 51.
7. When you have completed the policy configuration, select **Create**.

About policy attributes

With a few exceptions, policy attributes for a SAP policy are managed in the same way as for most other policy types. Certain policy attributes vary according to your specific backup strategy and system configuration.

For more information on policy attributes, see the NetBackup Administrator's Guide, Volume I.

Table 5-2 Policy attributes for NetBackup for SAP policies

Attribute	Description
Policy type	Determines the types of clients that can be backed up with the policy. For SAP databases, select the policy type **SAP**.
Keyword phrase	For NetBackup for SAP, the **Keyword phrase** entry is ignored.
Snapshot Client and Replication Director	This group contains the options that enable backups with Snapshot Client.

NetBackup for SAP backup types

Table 5-3 shows the backup schedules you can specify.

Table 5-3 SAP backup types

Backup type	Description
Application Backup	The Application Backup schedule enables user-controlled NetBackup operations from the client. These operations include those initiated from the client and those initiated from an automatic schedule on the primary server. Configure at least one Application Backup schedule for each database policy. The Default-Application-Backup schedule is configured automatically as an Application Backup schedule.
Automatic Full Backup	An automatic full backup schedule specifies the dates and times for NetBackup to automatically start backups. NetBackup runs the scripts in the order that they appear in the file list. If there is more than one client in the policy, the scripts are run on each client. The schedule name and type are passed to the scripts so they can distinguish initiation from a full vs incremental schedule and operate appropriately.

Table 5-3 SAP backup types *(continued)*

Backup type	Description
Automatic Incremental Backup	An automatic incremental backup is a backup of only those blocks that have changed since the last automatic full (baseline) backup. This kind of backup takes less time and space than a full backup because the automatic incremental backup contains only the changed data. NetBackup for SAP supports this type of backup in MaxDB environments and in Oracle environments with RMAN. This schedule type has no direct value in Oracle environments without RMAN because a file level backup, and not the block level backup, is performed.

Configure an application backup schedule

A database backup requires an application backup schedule. You cannot perform backups if this type of schedule is not included in the policy. NetBackup automatically creates this schedule and names it **Default-Application-Backup**.

The backup window for an application backup schedule must encompass the time period during which all scheduled jobs and client-initiated jobs can occur. This window is necessary because the application backup schedule accepts the backup request from NetBackup for SAP regardless of whether the backup was initiated from an automatic schedule or from the client. You can choose to set the window for the application backup schedule for 24 hours per day, seven days per week. This window ensures that your operations are never locked out due to the application backup schedule.

To configure an application backup schedule

1. Open the policy and select the **Schedules** tab.
2. Select the schedule that is named **Default-Application-Backup** and select **Edit**.
3. Specify the other properties for the schedule.

   Note: Some types of SAP backups allow for multiple application backup schedules, so that storage and retention attributes can be set appropriately. In that case, additional application schedules can be created. Specify the application backup schedule name in the `init`*`SID`*`.utl` file on the client.

   See "Schedule properties" on page 49.
4. Select **Add**.

Example application backup schedule

Assume the following:

- Users perform database backup operations during business hours, 08:00 to 13:00.
- The automatic backups that use this policy start between 18:00 and 22:00.

In this scenario, the application backup schedule must have a start time of 0800 and a duration of 14 hours. Alternatively, the schedule can have two windows each day; one with a start time of 0800 and duration of 5 hours, and another with a start time of 1800 and a duration of 4 hours.

Table 5-4 Example settings for a NetBackup for SAP application backup schedule

Schedule option	Setting
Retention	2 weeks
Backup window	Sunday through Saturday 00:08:00 - 22:00:00

Configure automatic backup schedules

If you plan to have NetBackup perform automatic backups, or if you use Snapshot Client features, you need one or more automatic backup schedules.

To configure an automatic backup schedule

1. Open the policy and select the **Schedules** tab.
2. Click **Add**.
3. Specify a unique name for the schedule.
4. Select the **Type of backup**.

   See “NetBackup for SAP backup types” on page 46.
5. Specify the other properties for the schedule.

   See “Schedule properties” on page 49.
6. Select **Add**.

Example automatic backup schedule

Table 5-5 shows example settings for an automatic backup schedule.

Table 5-5 Example settings for a NetBackup for SAP automatic backup schedule

Schedule property	Setting
Retention	2 weeks
Frequency	Every week
Backup window	Sunday, 18:00:00 - 22:00:00

Schedule properties

This topic describes the schedule properties that have a different meaning for database backups than for file system backups. Other schedule properties vary according to your specific backup strategy and system configuration. Additional information about other schedule properties is available. See the NetBackup Administrator's Guide, Volume I.

Table 5-6 Description of schedule properties

Property	Description
Type of backup	Specifies the type of backup that this schedule can control. The selection list shows only the backup types that apply to the policy you want to configure. See "NetBackup for SAP backup types" on page 46.
Schedule type	You can schedule an automatic backup in one of the following ways: ■ **Calendar** The **Calendar** option lets you schedule the backup operations that are based on specific dates, recurring week days, or recurring days of the month. ■ **Frequency** The **Frequency** specifies the period of time that can elapse until the next backup operation begins on this schedule. For example, assume that the frequency is 7 days and a successful backup occurs on Wednesday. The next full backup does not occur until the following Wednesday. Typically, incremental backups have a shorter frequency than full backups.

Table 5-6 Description of schedule properties *(continued)*

Property	Description
Retention	Specifies a retention period to keep backup copies of files before they are deleted. The retention level also denotes a schedule priority within the policy. A higher level has a higher priority. Set the time period to retain at least two full backups of your database. In this way, if one full backup is lost, you have another full backup to restore. For example, if your database is backed up once every Sunday morning, you should select a retention period of at least 2 weeks. The retention period for an application backup schedule refers to the length of time that NetBackup keeps the non-snapshot SAP backup images. The retention period for an automatic schedule controls how long NetBackup keeps records of when automatic backups occurred. For example, if your database is backed up once every Sunday morning, you should select a retention period of at least 2 weeks.
Effect of type of schedule on retention period	The type of schedule you select affects the retention period as follows: ■ Frequency-based scheduling Set a retention period that is longer than the frequency setting for the schedule. For example, if the frequency setting is set to one week, set the retention period to be more than one week. The NetBackup scheduler compares the latest record of the automatic backup schedule to the frequency of that automatic backup schedule. This comparison is done to determine whether a backup is due. So if you set the retention period to expire the record too early, the scheduled backup frequency is unpredictable. However, if you set the retention period to be longer than necessary, the NetBackup catalog accumulates unnecessary records. ■ Calendar-based scheduling The retention period setting is not significant for calendar-based scheduling. **Note:** If using Snapshot Client, the retention on the automatic schedule also determines how long the SAP phase 1 backup of the database files is retained.
Multiple copies	If you want to specify multiple copies of a backup for the policy, configure **Multiple copies** on the application backup schedule. If using Snapshot Client, also specify **Multiple copies** on the automatic schedule.

Add clients to a policy

The client list contains a list of the clients on which your scripts are run during an automatic backup. This list determines the clients that can send backup requests to the application schedule. A NetBackup client must be in at least one policy but can be in more than one.

For a NetBackup for SAP policy, clients you want to add must have the following items installed or available:

- SAP

- NetBackup client or server
- The backup or restore scripts

To add clients to a policy

1. Open the policy and select the **Clients** tab.
2. Select **Add**.
3. Type the name of the client and select the hardware and operating system of the client.

   If SAP is installed in a cluster, specify the virtual name of the SAP server as the client name.

   Note: If you installed NetBackup on more than one node in the SAP cluster, you must perform additional configuration.

   See "Reviewing the auto-discovered mappings" on page 72.

4. Select **Add**.

Adding NetBackup for SAP scripts to the backup selections list

The backup selections list in a database policy is different from the list in non-database policies. For example, in a Standard or MS-Windows policy, the list contains files and directories to be backed up. In a database policy, you specify scripts to be run.

Add scripts to the backup selections list only if you want to create a policy for automatic backups. In that case, add the scripts to a policy that has automatic backup schedules. NetBackup runs the scripts in the order that the scripts appear in the backup selections list.

To add scripts to the backup selections list

1. Ensure that the script resides on the client.

   See "Registering authorized locations used by a NetBackup database script-based policy" on page 252.

2. Open the policy and select the **Backup selections** tab.
3. Select **Add**.

4 In the **Pathname or directive** box, type the full path name of a script on the client.

For example:

It is recommended that you copy the script to a new file name or location so that it is not overwritten during upgrades.

5 Select **Add**.

Browse for scripts to add to the backup selections list (Windows)

On Windows you browse for a script that is located on a client and add that script to the backup selections list.

To browse for scripts to add to the backup selections list (Windows)

1 Ensure that the script resides on the client.

See "Registering authorized locations used by a NetBackup database script-based policy" on page 252.

2 Open the policy and click the **Backup Selections** tab.

3 Click **Add**.

4 Enter the name or full path to the script.

5 Click **Add**.

Rules for NetBackup for SAP scripts

Observe the following when you use scripts:

- To ensure that scripts run successfully on all clients, ensure that:
  - The scripts reside on each client in the client list and in the same location on each client
  - The script location is registered.
    See "Registering authorized locations used by a NetBackup database script-based policy" on page 252.
  - That NetBackup can access the location.
  - If you use NetBackup for SAP in a cluster, that the scripts reside in a location that is available after a failover.
- NetBackup installs sample scripts when you install the software; you can modify these scripts for your own use. Be sure to rename the scripts or copy the scripts to a location outside of the original installation location. This action ensures that

future NetBackup upgrades do not overwrite the SAP instance-specific modification that you are required to make before using the scripts.

See "NetBackup for SAP backup scripts" on page 53.

NetBackup for SAP backup scripts

If you configured a policy for automatic scheduling, you also need to specify backup scripts to be run automatically. You can either create your own scripts or you can modify one or more of the scripts that are included in the following locations:

Windows:

install_path\NetBackup\DbExt\SAP\samples

UNIX or Linux:

/usr/openv/netbackup/ext/db_ext/sap/scripts/sap_oracle
/usr/openv/netbackup/ext/db_ext/sap/scripts/sap_maxdb

Table 5-7 shows the sample scripts.

Table 5-7 NetBackup for SAP scripts

Script name	Used by (underlying database)	Comments
sap_offline_backup sap_offline_backup.cmd	Oracle without RMAN	Issues a brbackup command and performs a whole offline database backup. When it is initiated from an automatic schedule, the NetBackup bphdb process sets these variables in the environment, for use by the script or SAP backint interface. They identify the primary server and policy that caused the script to run. UNIX/Linux: $SAP_SERVER $SAP_POLICY Windows: %SAP_SERVER% %SAP_POLICY% If the script is initiated from the client host, these variables can be set into the environment manually. On UNIX/Linux, you should export the variables after they are set.
sap_online_backup.cmd sap_online_backup	Oracle without RMAN	Issues the brbackup and brarchive commands and performs a whole online database backup and then an offline redo log backup.

Table 5-7 NetBackup for SAP scripts *(continued)*

Script name	Used by (underlying database)	Comments

`sap_redo_log_backup`
 	
`sap_redo_log_backup.cmd`
	Oracle without RMAN	Issues the `brarchive` command and backs up the offline redo log files.

`sap_rman_backup`
 	
`sap_rman_backup.cmd`
	Oracle with RMAN	Issues the `brbackup` command to perform either an RMAN incremental level 0 or incremental level 1 cumulative backup. The level depends on whether the script was initiated from an Automatic Full or Automatic Cumulative schedule.
		
		If this script is to be initiated from the client, add logic to set the `SAP_FULL` or `SAP_CINC` variable appropriately. Otherwise the script exits with status 0 without taking a backup.

`sap_maxdb_backup`
 	
`sap_maxdb_backup.cmd`
	MaxDB	Issues a `dbmcli` command and uses the specified backup medium to perform either a full or cumulative incremental backup. The type depends on whether the script was initiated from an Automatic Full or Automatic Cumulative schedule.
		
		If this script is to be initiated from the client, add logic to set the `SAP_FULL` or `SAP_CINC` variable appropriately. Otherwise the script exits with status 0 without taking a backup.

If you decide to use these scripts, customize them for your environment.

Note: Always specify the correct script when configuring automatic backups or when starting operations through NetBackup. NetBackup for SAP does not generate an error if a restore script is used for a backup operation or a backup script is used for a restore operation.

Note: All scripts must be stored and run locally. One recommendation is that scripts should not be world-writable. Scripts are not allowed to be run from network or remote locations. Any script that is created and saved in the NetBackup `db_ext` (UNIX) or `dbext` (Windows) location needs to be protected during a NetBackup uninstall.

See "Registering authorized locations used by a NetBackup database script-based policy" on page 252.

Modifying the NetBackup for SAP backup scripts

This topic describes how to modify the backup scripts for NetBackup for SAP.

To modify the backup scripts

1. Copy the example script(s) from one of the following director(ies) to a different location on your client.

   UNIX or Linux:

   ```
   /usr/openv/netbackup/ext/db_ext/sap/scripts/sap_oracle
   /usr/openv/netbackup/ext/db_ext/sap/scripts/sap_maxdb
   ```

   Windows:

   ```
   install_path\NetBackup\DbExt\SAP\samples
   ```

   The SAP scripts must reside in an authorized location on the client. Do not store your scripts in the sample directory because your modifications are lost if you upgrade or reinstall. Always relocate your scripts to a safe location. If the client is clustered, this location must be available after a failover.

   See "Registering authorized locations used by a NetBackup database script-based policy" on page 252.

2. Enable proper permissions on the script files so NetBackup can run the scripts.

   On UNIX or Linux, set the access permissions of these scripts to `775`.

3. Modify the script according to the instructions in the file.

   On UNIX or Linux, remember to edit the lines that starts with `su`. If you do not include `su - user` in your script, it does not run with the proper account and environment variables. (*user* is the SAP administrator account.) This situation can lead to problems with database backups and restores.

   See "Example NetBackup for SAP backup script (Windows)" on page 57.

   See "Example NetBackup for SAP backup script (UNIX)" on page 59.

4 In case of sort_backup_type custom, custom sort file can be generated by `nbsapcustomsort` utility.

- Linux:

```
/usr/openv/netbackup/bin/nbsapcustomsort -u <oracle_username>
 -p <custom_sort_file_path>
```

- Windows:

```
install_path\NetBackup\bin\nbsapcustomsort.exe
-p <path_for_custom_sort_file_path>
```

Note: Ensure that path for `-p` is allowed listed and it should not contain any spaces.

See "Example NetBackup for nbsapcustomsort utility" on page 56.

5 Test the scripts you created.

See "Perform a manual backup" on page 77.

Example NetBackup for nbsapcustomsort utility

`nbsapcustomsort` utility is used to generate the `custom_sort_file` with name `sap_custom_sort_file` at:

UNIX or Linux:

```
<custom_sort_file_path>/sap_custom_sort_file
```

Windows:

```
<custom_sort_file_path>/sap_custom_sort_file
```

`nbsapcustomsort` utility is available at the following locations:

Linux:

```
/usr/openv/netbackup/bin/nbsapcustomsort
```

Windows:

```
install_path\NetBackup\bin\nbsapcustomsort.exe
```

The following are the parameters for `nbsapcustomsort`:

1. Required parameter for Linux: **-p <custom_sort_file_path> -u <username>**.

2. Required parameter for Windows: **-p <custom_sort_file_path>**.

   Note: User is the SAP administrator account.

3. Optional parameters: -t <total_streams> -v <streams_per_volume>
   - `-t` is the total number of streams for entire SAP Instance backup.
   - `-v` is the maximum number of streams per SAP device.

Example NetBackup for SAP backup script (Windows)

This topic includes an example backup script that was created for Windows.

For example, the `sap_offline_backup.cmd` script contains the following lines:

```
@REM
@REM $Header: sap_offline_backup.cmd,v 1.2 2002/11/20 00:47:59 $
@REM
@echo off

REM This environment variable are created by NetBackup (bphdb)

echo SAP_SCHEDULED = %SAP_SCHEDULED%
echo SAP_USER_INITIATED = %SAP_USER_INITIATED%
echo SAP_SERVER = %SAP_SERVER%
echo SAP_POLICY = %SAP_POLICY%

REM ------------------------------------------------------------------------
REM Replace cer below with the Oracle SID of the target database.
REM ------------------------------------------------------------------------

set ORACLE_SID=cer

REM ------------------------------------------------------------------------
REM Replace c:\oracle below with the Oracle home path.
REM ------------------------------------------------------------------------

set ORACLE_HOME=c:\oracle

REM ------------------------------------------------------------------------
REM Replace C:\oracle\CER with SAPData Home Path
REM ------------------------------------------------------------------------
```

```
set SAPDATA_HOME=C:\oracle\CER

REM ------------------------------------------------------------------------
REM Replace path with the correct sap archive path.
REM ------------------------------------------------------------------------

set SAPARCH=%SAPDATA_HOME%\saparch

REM ------------------------------------------------------------------------
REM Replace path with the correct sap backup path.
REM ------------------------------------------------------------------------

set SAPBACKUP=%SAPDATA_HOME%\sapbackup

REM ------------------------------------------------------------------------
REM Replace path with the correct sap reorg path.
REM ------------------------------------------------------------------------

set SAPREORG=%SAPDATA_HOME%\sapreorg

REM ------------------------------------------------------------------------
REM Replace path with the correct Path to  Brtools
REM ------------------------------------------------------------------------

set SAPEXE=C:\usr\sap\%ORACLE_SID%\sys\exe\run

REM ------------------------------------------------------------------------
REM Replace path with the correct BRBACKUP path.
REM ------------------------------------------------------------------------

set BRBACKUP=c:\usr\sap\%ORACLE_SID%\sys\exe\run\brbackup

REM Full offline backup

set CMD_LINE=%BRBACKUP% -u internal/ -c -d util_file -t offline -m all
%CMD_LINE%

REM ------------------------------------------------------------------------
REM To communicate with NetBackup's job monitor for an automatic schedule
REM a "STATUS_FILE" variable is created by NetBackup (bphdb) that contains
REM a path to a file.  This file is check by NetBackup to determine if the
REM automatic schedule was successful.  It is up to the script to write
REM a 0 (passed) or 1 (failure) to the status file to communicate to NetBackup
```

```
REM the status of the execution of the script.  The following code echo a 0
REM to %STATUS_FILE% if succcessful and echo a 1 to %STATUS_FILE% for a
REM failure.
REM ------------------------------------------------------------------------

if errorlevel 1 goto errormsg
echo BRBACKUP successful
if "%STATUS_FILE%" == "" goto end
if exist "%STATUS_FILE%" echo 0 > "%STATUS_FILE%"
goto end

:errormsg
echo Execution of BRBACKUP command failed - exiting
if "%STATUS_FILE%" == "" goto end
if exist "%STATUS_FILE%" echo 1 > "%STATUS_FILE%"

:end
```

Example NetBackup for SAP backup script (UNIX)

For example, the `sap_offline_backup` script contains the following lines:

```
#!/bin/sh
#
#NOTE:IF your SAP user (in this script orasap) runs in C shell, environmental
#variables can not be exported. In that case, you should modify this script to
#work in your environment. For example:
#   SAP_SERVER=$SAP_SERVER; export SAP_SERVER; (Correct for Bourne and Korn shells)
#   can change into
#   setenv SAP_SERVER $SAP_SERVER; (Correct for C shell)
#

#
#This environment variable are created by NetBackup (bphdb)
#

echo "SAP_SCHEDULED = $SAP_SCHEDULED"
echo "SAP_USER_INITIATED = $SAP_USER_INITIATED"
echo "SAP_SERVER = $SAP_SERVER"
echo "SAP_POLICY = $SAP_POLICY"

RETURN_STATUS=0
```

```
SAP_ENV=""
#
# If SAP_SERVER exists then export it to make it available to backint
#
if [ -n "$SAP_SERVER" ]
then
        SAP_ENV="$SAP_ENV SAP_SERVER=$SAP_SERVER; export SAP_SERVER;"
#if Oracle DBA account( orasap user) uses C Shell, comment the above line and
#uncomment next line
#       SAP_ENV="$SAP_ENV setenv SAP_SERVER $SAP_SERVER;"
fi

#
# If SAP_POLICY exists then export it to make it available to backint
#
if [ -n "$SAP_POLICY" ]
then
        SAP_ENV="$SAP_ENV SAP_POLICY=$SAP_POLICY;export SAP_POLICY;"
#if Oracle DBA account( orasap user) uses C Shell, comment the above line and
#uncomment next line
#       SAP_ENV="$SAP_ENV setenv SAP_POLICY $SAP_POLICY;"
fi

#
# Full offline backup
#

CMD_LINE="$SAP_ENV brbackup -c -d util_file -t offline -m all"

#
# The username on the "su" command needs to be replaced with the correct
# user name.
#
echo "Execute $CMD_LINE"
su - orasap -c "$CMD_LINE"

RETURN_STATUS=$?

exit $RETURN_STATUS
```

NetBackup for SAP backup script parameters

The SAP scripts that enable the SAP utilities to perform backups and restores use parameters defined in the following sources:

- Environment variables
- Configuration files. For example, `init`*`SID`*`.utl`, where *SID* is the instance.
- (UNIX or Linux) NetBackup configuration file (`bp.conf`)
- (Windows) NetBackup web UI configuration

You can use the different parameter sources to create SAP scripts to perform different database backup and restore tasks.

For example, you can define `%SAP_POLICY%` (`$SAP_POLICY`) in an SAP script to perform different types of backups.

Table 5-8 describes the environment variables that are created locally when NetBackup's automatic scheduler runs an SAP script. Their values can be used to direct the flow of actions within the script. Some affect the SAP `backint` interface, and can be changed within the script.

Table 5-8 Environment variables

Environment variable	Purpose
`$SAP_POLICY` `%SAP_POLICY%`	Names the NetBackup for SAP policy, if initiated from an automatic backup schedule. If set, specifies the policy to which the application backup jobs will be sent.
`$SAP_SCHEDULED` `%SAP_SCHEDULED%`	Is set to `1` if this operation was initiated from an automatic schedule. If not set to `1`, the script can then set values for the other variables during client-initiated operations.
`$SAP_SERVER` `%SAP_SERVER%`	Names the NetBackup server if initiated from an automatic backup schedule. If set, specifies the primary server to which the application backup jobs will be sent.
`$SAP_USER_INITIATED` `%SAP_USER_INITIATED%`	The opposite of SAP_SCHEDULED. Is set to `0` if this operation was initiated from an automatic schedule.
`SAP_SNC_SCHED`	Names the NetBackup for SAP automatic schedule, if initiated from an automatic schedule.
`SAP_FULL/SAP_CINC/SAP_INCR`	If the backup initiated from an automatic schedule, the appropriate one will be set to `1`, the others to `0`.

Table 5-8 Environment variables *(continued)*

Environment variable	Purpose
SAP_CLIENT	Names the NetBackup for SAP client from the policy, if initiated from an automatic schedule. If set, specifies the client name which the application backup job will use

Configuring the logon account for the NetBackup Client Service for NetBackup for SAP

To configure the logon account for the NetBackup Client Service for NetBackup for SAP

1. In the Windows Services application, open the **NetBackup Client Service** entry.
2. On the **Log On** tab, provide the following:
   - Type the password.
3. Stop and start the NetBackup Client Service.

About SAP configuration files

Table 5-9 lists the files that you need to modify depending upon whether your underlying database is an Oracle database or a MaxDB database.

Table 5-9 SAP configuration files

File	Underlying database	Comments
init*SID*.utl	Oracle and MaxDB	SAP parameter file. This file specifies policy, schedule, client, and other information to the backint interface.
init*SID*.sap	Oracle	SAP profile file. This file contains the information that the SAP tools use to interact with the backintand RMAN interfaces.
bsi.env	MaxDB	MaxDB configuration file. This file contains the information that the SAP tools use to interact with the backint interface.

See "About configuring NetBackup for SAP" on page 43.

See "Relationships between SAP database configuration files" on page 63.

See "NetBackup for SAP backup types" on page 46.

Relationships between SAP database configuration files

The SAP tools pass information in these configuration files to the `backint` interface.

When `backint` runs, the following occurs:

- In NetBackup for SAP on Oracle databases, you specify the actual `init`*`SID`*`.utl` file name as the argument to the `util_par_file` parameter in the `init`*`SID`*`.sap` file.
- In NetBackup for SAP on MaxDB databases, you specify the actual `init`*`SID`*`.utl` file name as the argument to the `PARAMETERFILE` parameter in the `bsi.env` file.
- The `init`*`SID`*`.utl` file name becomes the argument to the `backint` command's `-p par_file` parameter.

See "About SAP configuration files" on page 62.

See "NetBackup for SAP on MaxDB databases" on page 20.

See "NetBackup for SAP backup types" on page 46.

Modifying the initSID.utl file for NetBackup for SAP

The `init`*`SID`*`.utl` file is the SAP parameter file. Edit this file as part of configuring NetBackup for SAP.

To modify the `initSID.utl` **file**

1. Check for an existing parameter file.

   If an `init`*`SID`*`.utl` file already exists in the `$ORACLE_HOME/dbs` or `%ORACLE_HOME%\database` directory, copy it to a backup file.

2. Create a new parameter file.

   Copy the parameter file that is included in the NetBackup for SAP software distribution to the appropriate directory:

   Windows:

   ```
   %ORACLE_HOME%\database
   ```

   UNIX or Linux:

   ```
   $ORACLE_HOME/dbs
   ```

   For example, if the Oracle instance is `SAP`, copy the NetBackup example `.utl` file to `initSAP.utl`, as follows:

   Windows:

```
copy install_path\NetBackup\dbext\sap\scripts\initSAP.utl
%ORACLE_HOME%\database\initSAP.utl
```

   UNIX or Linux:

```
cp /usr/openv/netbackup/ext/db_ext/sap/scripts/sap_oracle/initSAP.utl \
$ORACLE_HOME/dbs/initSAP.utl
```

3. Modify `init`*`SID`*`.utl` file with a text editor to set the parameters to values that are appropriate to your site.

   The following parameters are required: `client`, `switch_list`, `switch_log`, `switch_sem`.

   You can set the following additional parameters depending on whether or not you use RMAN:

   - If you use RMAN, set `policy`, `schedule`, and `server`.
   - If you do not use RMAN, set `drives`, `policy`, `schedule`, and `server`.

Example NetBackup for SAP initSID.utl file

The following shows example `init`*`SID`*`.utl` files with all the parameters required for NetBackup for SAP on Oracle databases:

UNIX or Linux:

```
policy sap_policy_1
schedule Default-Application-Backup
client puffin
server puffin
drives 2
switch_list /oracle/sap/sapbackup/.switch.lis
switch_sem /oracle/sap/sapbackup/.switch.sem
switch_log /oracle/sap/sapbackup/.switch.log
```

Windows:

```
policy sap_policy_1
schedule Default-Application-Backup
client puffin
server puffin
drives 2
switch_list F:\oracle\SID\SAPBackup\.switch.lis
switch_sem F:\oracle\SID\SAPBackup\.switch.sem
switch_log F:\oracle\SID\SAPBackup\.switch.log
```

Note: The question mark (`?`) and ampersand (`&`) Oracle substitution characters and the `$ORACLE_HOME` (`%ORACLE_HOME%`) environment variable are not allowed in the `init`*`SID`*`.utl` file. They are not allowed in the *par_file* that the `backint -p` *`parfile`* option identifies.

Modifying the initSID.sap file (NetBackup for SAP Oracle databases only)

The `init`*`SID`*`.sap` file is the SAP profile file. This file signals to the SAP tools that you use the `backint` interface and conveys information about NetBackup for SAP to the SAP tools. Edit this file as part of configuring NetBackup for SAP when the underlying database is Oracle.

To modify the initSID.sap file

1. Check for an existing profile file.

   If an `init`*`SID`*`.sap` file already exists in the `$ORACLE_HOME/dbs` (`%ORACLE_HOME%\database`) directory, copy it to a backup file.

2. Create a new profile file.

   Copy the profile file that is included in the NetBackup for SAP software distribution to the `$ORACLE_HOME/dbs` (`%ORACLE_HOME%\database`) directory.

   For example, if the Oracle instance is `SAP`, copy the NetBackup example `.sap` file to `initSAP.sap`, as follows:

   Unix or Linux:

```
cp /usr/openv/netbackup/ext/db_ext/sap/scripts/sap_oracle/initSAP.sap \
$ORACLE_HOME/dbs/initSAP.sap
```

   Windows:

```
copy install_path\NetBackup\dbext\sap\scripts\initSAP.sap
%ORACLE_HOME%\database\initSAP.sap
```

3. Modify the `init`*`SID`*`.sap` file with a text editor to set the `backup_dev_type` and `util_par_file` parameters to values that are appropriate to your site.

4. (Conditional) Specify the `rman_parms` parameter and set the `NB_ORA_SAP` environment variable.

   Perform this step only if you use NetBackup for SAP on an Oracle database with RMAN. For example:

   ```
   rman_parms = "ENV=(NB_ORA_SAP=file)"
   ```

   where *file* specifies the full path to the `initSID.utl` file. Do not use substitution in the *file* specification.

   For example, this parameter might look like the following:

   Unix or Linux:

   ```
   rman_parms = "ENV=(NB_ORA_SAP=/apps/oracle/dbs/initCER.utl)"
   ```

   Windows:

```
rman_parms = "ENV=(NB_ORA_SAP=C:\apps\oracle\database\initCER.utl)"
```

   Depending on your site practices, you might want to specify the following additional NetBackup environment variables and their values as arguments to the `rman_parms` parameter: `NB_ORA_POLICY`, `NB_ORA_SCHED`, `NB_ORA_CLIENT`, and `NB_ORA_SERV`. You can also specify values for these environment variables in the `initSID.utl` file. If you specify differing values, the ones in `initSID.sap` override those in `initSID.utl`.

   For example, consider that you want to use one policy and schedule for the first part of a backup (data files). You want a different policy and schedule for the second part of a backup (the control file backup). You can set the `NB_ORA_POLICY` and `NB_ORA_SCHED` variables in `initSID.sap` to pertain to the first part of the backup. Then, you can set these variables to the different values that pertain to the second part of the backup in `initSID.utl`. NetBackup for SAP checks only `initSID.utl` when backing up the control files; it does not check `initSID.sap`.

   You can include other RMAN parameters in this file. For more information on other parameters, see your SAP documentation.

5. Save and close the `initSID.sap` file.

Example initSID.sap file (NetBackup for SAP on Oracle databases only)

For example, assume that `initSAP.sap` is your profile file and `initSAP.utl` is your parameter file. You can use the following procedure to set parameters.

To edit the `initSID.sap` file (NetBackup for SAP on Oracle databases only)

1. Use a text editor to open the `initSAP.sap` file.
2. Find the following line:

```
backup_dev_type = tape
```

3. Copy and paste this line under the original.

```
backup_dev_type = tape
backup_dev_type = tape
```

4. Comment out the original line.

```
#backup_dev_type = tape
backup_dev_type = tape
```

5. Change tape to util_file.

```
#backup_dev_type = tape
backup_dev_type = util_file
```

6. Find the following line:

```
util_par_file = file_path
```

7. Copy and paste this line under the original.

```
util_par_file = file_path
util_par_file = file_path
```

8 Comment out the original line.

```
#util_par_file = file_path
util_par_file = file_path
```

9 Change *file_path* to the path to the `initSID.utl` parameter file.

UNIX or Linux:

```
#util_par_file = file_path
util_par_file = ?/dbs/init@.utl
```

Windows:

```
#util_par_file = file_path
util_par_file = ?\database\init@.utl
```

This example file uses environment variable settings and Oracle substitution characters. When the SAP tools interpret the `initSID.sap` profile file, they replace the question mark (`?`) and at sign (`@`) characters.

These characters are replaced with the values assigned to the `$ORACLE_HOME` (`%ORACLE_HOME%`) and `$ORACLE_SID` (`%ORACLE_SID%`) environment variables, respectively.

Modifying the bsi.env file (NetBackup for SAP on MaxDB databases)

The `bsi.env` file specifies the following:

- The `backint` for MaxDB program's location, which is NetBackup for SAP on MaxDB.
- How to communicate with `backint`.

The database manager uses this information when it performs backups and restores.

`bsi.env` is the default name for this configuration file. If you want to change this name, you can specify a different name in the `BSI_ENV` environment variable in the MaxDB user environment.

Table 5-10 lists the parameters that you can include in the `bsi.env` file.

Table 5-10 bsi.env file parameters

Parameter	Description

`BACKINT /usr/openv/netbackup/bin/backint`
 	
`BACKINT <install_path>\Veritas\NetBackup \bin\backint.exe`
	Specifies the absolute path and file name of the `backint` for MaxDB program.

`INPUT absolute_path/file_name`
 	
`INPUT absolute_path\file_name`
 	
`OUTPUT absolute_path/file_name`
 	
`OUTPUT absolute_path\file_name`
 	
`ERROROUTPUT absolute_path/file_name`
 	
`ERROROUTPUT absolute_path\file_name`
	The `INPUT` parameter specifies the absolute path and file name of the standard input file. The `OUTPUT` parameter specifies the absolute path and file name of the standard output file. The `ERROROUTPUT` file specifies the absolute path and file name of the standard error output file.
	
	The database manager creates these files temporarily and uses them for standard input, output, and error output for `backint` for MaxDB.

`PARAMETERFILE absolute_path/file_name`
 	
`PARAMETERFILE absolute_path\file_name`
	Specifies the absolute path and file name of the configuration file for the `backint` for MaxDB program, `init`*`SAP`*`.utl`.

`TIMEOUT_SUCCESS seconds`	The number of seconds after which the backup tool exits and after the database kernel has copied all data successfully. Default is 300. This value is the period of time after which the database manager terminates the backup tool if all data has been copied to the pipes. Note that the backup tool always exits after this defined period. This exit occurs whether all the data is copied from the database kernel is backed up by the backup tool or not.
`TIMEOUT_FAILURE seconds`	This parameter is effective only if the database kernel has stopped the backup due to an error. In such a case, *seconds* specifies the number of seconds after which the backup tool exits. Default is 300. Allow enough time for the backup tool to save all the data that the database kernel copies. The backup tool always exits when this timeout is reached, even if not all of the copied data has been saved yet. However, do not specify too much time. You risk a situation in which a backup that is stopped by the database kernel blocks other backups, such as automatic nightly backups, from being started.

Table 5-10 bsi.env file parameters *(continued)*

Parameter	Description
ORIGINAL_RUNDIRECTORY *absolute_path*	Specifies the absolute path of the run directory of the source database. The default value is the run directory of the current database instance. This value is case-sensitive. Specify this parameter only for a migration. The following prerequisite must be met for migration from one computer to another: The version of `backint` for MaxDB that you use is capable of backing up from the source computer available on the target computer.

To create a MaxDB configuration file

1. Use a text editor to create the MaxDB configuration file, `bsi.env`.

   You can give the MaxDB configuration file a name other than `bsi.env`. If so, edit the `BSI_ENV` environment variable and include the full path to this file.

2. Use the parameters to specify backup characteristics.

   You must include the following parameters in `bsi.env: BACKINT`, `INPUT`, `OUTPUT`, `ERROROUTPUT`, and `PARAMETERFILE`. All other parameters are optional.

   The following is an example of a file with the minimum configuration specified:

   UNIX or Linux:

   ```
   BACKINT /usr/openv/netbackup/bin/backint
   INPUT /export/home/sapdb/logs/backint.in
   OUTPUT /export/home/sapdb/logs/backint.out
   ERROROUTPUT /export/home/sapdb/logs/backint.err
   PARAMETERFILE /export/home/sapdb/indep_data/wrk/initSAP.utl
   ```

   Windows:

   ```
   BACKINT install_path\NetBackup\bin\backint
   INPUT C:\sapdb\logs\backint.in
   OUTPUT C:\sapdb\logs\backint.out
   ERROROUTPUT C:\sapdb\logs\backint.err
   PARAMETERFILE C:\sapdb\indep_data\wrk\initSAP.utl
   ```

Backing up files mounted with LOFS with NetBackup for SAP (UNIX or Linux)

The NetBackup for SAP Agent for Unix/Linux supports backups of files from a loopback virtual file system (LOFS). An LOFS file system lets you create a virtual file system that provides access to existing files through the use of alternate pathname. Consider a loopback mount of the `/oracle` file system onto the `/database` file system. This loopback mount allows the `/oracle` file system to also appear under the `/database` file system. All files in `/oracle` are then accessible either from a pathname relative to `/oracle` or relative to `/database`. For example, `/database/sapdata1/system/system.dbf`.

If you have a local file system mounted as an LOFS, you do not need to select **Follow NFS**.

Select the “Follow NFS” option for backups of an LOFS if the actual file system (for example, `/oracle`) is either of the following:

- An NFS mount
- A Solaris global zone file system and mounted with LOFS onto a non-global zone.

About permissions for NetBackup for SAP log files (UNIX)

NetBackup uses the `/usr/openv/netbackup/logs` directory tree not only for the recording of troubleshooting information, but for progress and communication updates to users and other NetBackup applications. Restrictive permissions on these directories can not only disable the collection of troubleshooting data, but also prevent the application itself from functioning correctly.

See “Enable the debug logs manually (SAP) (UNIX/Linux) ” on page 196.

Reviewing the auto-discovered mappings

In certain scenarios, a NetBackup host shares a particular name with other hosts or has a name that is associated with a cluster. To successfully perform backups and restores with NetBackup for SAP, you must approve each valid auto-discovered mapping that NetBackup discovers in your environment. Or, manually add the mappings.

See the section called “Approve the auto-discovered mappings for a cluster” on page 74.

See the section called “Manually map host names” on page 75.

Examples of the configurations that have multiple host names include:

- A host is associated with its fully qualified domain name (FQDN) and its short name or its IP address.
- If the SAP server is clustered, the host is associated with its node name and the virtual name of the cluster.

These mappings are configured in the **Security > Host mappings** node in the NetBackup web UI. You can also use the `nbhostmgmt` command to manage the mappings. See the NetBackup Security and Encryption Guide and NetBackup Web UI Administrator's Guide for more details.

Auto-discovered mappings for a cluster

In a SAP cluster environment, you must map the node names to the virtual name of the cluster if the following apply:

- If the backup policy includes the cluster name (or virtual name)
  This mapping is not required for SAP HANA. That type of policy includes the node names in the cluster and not the virtual name of the cluster.
- If the NetBackup client is installed on more than one node in the cluster, the virtual name must be mapped to each node.
  If the NetBackup Client is only installed on one node, then no mapping is necessary.

Note: This procedure only applies to SAP, not to SAP HANA.

Approve the auto-discovered mappings for a cluster

To approve the auto-discovered mappings for a cluster

1. In the NetBackup web UI, expand **Security > Host mappings**.
2. Click the **Mappings to approve** tab.

   The list displays the hosts in your environment and the mappings or additional host names that NetBackup discovered for those hosts. A host has one entry for each mapping or name that is associated with it.

   For example, for a cluster with hosts `client01.lab04.com` and `client02.lab04.com`, you may see the following entries:

Host	Auto-discovered mapping
client01.lab04.com	client01
client01.lab04.com	clustername
client01.lab04.com	clustername.lab04.com
client02.lab04.com	client02
client02.lab04.com	clustername
client02.lab04.com	clustername.lab04.com

3. Click the name of the host.
4. Review the mappings for the host and click **Approve** if you want to use the discovered mappings.

   For example, if the following mappings are valid for `client01.lab04.com`, then you approve them.

Auto-discovered mapping	Valid name for
client01	The short name of the client
clustername	The virtual name of the cluster
clustername.lab04.com	The FQDN of the virtual name of the cluster

5 When you finish approving the valid mappings for the hosts, click on the **Hosts** tab.

For hosts `client01.lab04.com` and `client02.lab04.com`, you see entries for **Mapped host or IP address** that are similar to the following:

Host	Mapped host names/IP addresses
client01.lab04.com	client01.lab04.com, client01, clustername, clustername.lab04.com
client02.lab04.com	client02.lab04.com, client02, clustername, clustername.lab04.com

6 If you need to add a mapping that NetBackup did not automatically discover, you can add it manually.

Table 5-11 Example mapped host names for a SAP cluster environment

Environment	Host	Mapped host names
Cluster with two nodes	Physical name of *Node 1*	Virtual name of SAP server
	Physical name of *Node 2*	Virtual name of SAP server

Manually map host names

If you need to add a mapping that NetBackup did not automatically discover, you can add it manually.

To manually map host names

1 In the NetBackup web UI, expand **Security > Host mappings**.

2 Click on the **Hosts** tab.

3 Click **Add shared or cluster mappings**.

For example, type the name of the virtual name of the cluster. Then click **Add** to choose the hosts to which you want to map that virtual name.

Configuring the Maximum jobs per client

The **Maximum jobs per client** specifies the maximum number of concurrent backups that are allowed per client.

Note: Enter a large enough value for the **Maximum jobs per client** attribute to meet the number of jobs that SAP runs. You may need to experiment with different values at your site.

To configure the maximum jobs per client

1. On the left, select **Hosts > Host properties**.
2. Select the primary server.
3. If necessary, select **Connect**. Then select **Edit primary server**.
4. Select **Global attributes**.
5. Select the appropriate value for **Maximum jobs per client**.

   The default is 1.

You can use the following formula to calculate a smaller value for the Maximum jobs per client setting:

The following points help you to determine an appropriate setting:

- Determine which policies can be active concurrently based on scheduling and duration of backup.
- For each policy determine the number of jobs that can be sent concurrently be the client.
  - For SAP with RMAN backups, how many channels are specified in the `initSID.sap` file or the default instance configuration?
  - For SAP without RMAN, this is based on the keywords in the `initSID.utl` file; `sort_backup_type` and drives.
  - For MaxDB, this is the count of backup medium that is used for the backup. Typically one unless a parallel medium group is used.
  - Check if the policy specifies a lower concurrent job limit.
- Check if the settings for the storage units or storage servers that are used by the policies specify lower maximum numbers of concurrent jobs.
- For tape storage units, the number of tape drives may provide a lower limit. Check the MPX settings in the storage unit and policy; the lower MPX setting, times the number of tape drives, limits the number of jobs.

Perform a manual backup

After you configure the servers and assets in your environment, you can test the configuration settings with a manual backup. Perform a manual backup (or backups) from a policy with the automatic backup schedules that you created.

To perform a manual backup from a policy

1. On the left, select **Protection > Policies**.
2. Select the policy you want to test.
3. Select **Manual backup**.
4. Select the schedule that you want to use for the manual backup.
5. Select the clients that you want to include for the manual backup.

Testing multiple concurrent backup jobs for NetBackup for SAP for Oracle (without RMAN)

The following steps help you to test multiple concurrent jobs:

To test multiple drives

1. Complete the configuration of NetBackup, NetBackup for SAP, and SAP tools as described previously in this chapter.
2. Change the `drives` parameter in the following file:

   Windows:

   `%ORACLE_HOME%\database\init`*`SAP`*`.utl`

   UNIX or Linux:

   `$ORACLE_HOME/dbs/init`*`SAP`*`.utl`

   Change the value of this parameter to the number jobs the storage unit can accept concurrently.

   See “About parameters used in initSID.utl” on page 227.

   When you perform a backup, you should see two or more backups running.

   The NetBackup for SAP `backint` interface waits for all backups to complete before it reports success or failure to `brbackup`.

Chapter 6

Configuring RBAC for SAP HANA administrators

This chapter includes the following topics:

RBAC roles for the SAP HANA administrator

NetBackup enables control over which users can access which SAP HANA assets using Role Based Access Control (RBAC). You can grant RBAC access globally (to all SAP HANA assets) or to specific databases or instances.

The Default SAP HANA Administrator role has access to all SAP HANA assets (global). With this role the administrator can also manage credentials for instances. (These credentials are managed on the **Instances** tab in **Workloads > SAP**.) In addition, you may need other custom roles to give additional access to your SAP HANA administrators.

- A role that is restricted to an individual instance or database.
- A role that gives specific privileges to manage the host properties for SAP HANA clients.

Note the following:

- To create an RBAC role, you must have the RBAC Administrator role or the permissions to create roles.
- Contact your NetBackup administrator for assistance with creating roles and adding users to roles.

For information on the RBAC permissions and default roles, see the NetBackup API documentation.

Create a custom role for a non-SAP HANA administrator

A custom role can allow a non-SAP HANA administrator to sign into the NetBackup web UI with limited access. Use this role if you do not want an administrator (who is not SAP HANA administrator) to have the Default SAP HANA Administrator role. With this custom role, this type of administrator can only view SAP HANA jobs and does not have access to SAP HANA assets or their information.

To create a custom role for a non-SAP HANA administrator

1. On the left, select **Security > RBAC** and click **Add**.
2. Select **Custom role** and click **Next**.
3. Provide a Role name and a description. For example, include a description that the role allows a non-SAP HANA administrator to view SAP HANA jobs.
4. Under **Permissions**, select **Assign**.
5. On the **Assets** tab, go to **SAP HANA assets**, then select **View jobs**.
6. Select **Assign**.
7. Under **Workloads**, select **Assign**.
8. Select the option **Apply permissions to all existing and future SAP HANA assets**.
9. Select **Assign**.
10. Under **Users**, select **Assign**. Then add the users that you want to have this RBAC role.
11. Select **Assign**.
12. When you are done configuring the role, select **Add role**.

Chapter 7

Configuring NetBackup for SAP HANA clients for use with scripts

This chapter includes the following topics:

About configuring the primary and the media servers

You must configure the primary and the media servers before you configure the NetBackup client. The procedures for configuring primary and media servers are very similar. However, Cohesity recommends that you configure all server information such as storage devices and volumes from the primary server. Following this configuration sequence helps ensure that the primary server properly administers the media servers.

For initial NetBackup server configuration, in the NetBackup web UI and select **Settings > Getting Started**. Configure the storage devices, volumes, and the catalog backup and then create a backup policy.

Note: Before beginning the NetBackup configuration for each HANA appliance, check that `libelf.so.0` is linked to the `libelf.so.0.8.12` file in each HANA box operating system. The location of the library is `/usr/lib64/`.

If the `libelf.so.0.8.12` file is not present at the `/usr/lib64/` location then copy the `libelf.so.0.8.12` file from install of `SUSE SLES 11 SP1 to /usr/lib64/`.

Run the following commands:

- `cd /usr/lib64/`
- `ln -s libelf.so.0.8.12 libelf.so.0`

For more details about the `libelf.so.0.8.12` file error, see the following article:

https://support.cohesity.com/s/article/article-100009886

See "Configuring the media server for NetBackup SAP HANA" on page 81.

See "Configuring the primary server for NetBackup SAP HANA" on page 81.

Configuring the media server for NetBackup SAP HANA

For SAP HANA, for debugging and troubleshooting purpose you need to set the logging level to 5 (maximum).

On the left select **Hosts > Host properties**. Select the media server and then select **Edit media server**. Select **Logging** and edit the **Global logging level**.

Configuring the primary server for NetBackup SAP HANA

When you configure the primary server for SAP HANA, you must make the following additional changes:

- You must add the media server to the list of servers that the client accepts connections from. To add the media server:
  - On the left select **Hosts > Host properties**.
    Select the client and then select **Edit client**.
    Select **Servers > Additional Servers**. Then select **Add**.
- Configure the maximum number of concurrent jobs in the selected storage unit. The maximum concurrent jobs setting on the storage unit specifies the maximum number of jobs that NetBackup can send to a disk storage unit at one time. The

default setting is one job and the job count can range from 0 to 256. For an SAP policy type that is configured for HANA, set this value to a value greater than 64.

For an SAP policy type (specifically for HANA) when the storage unit is the AdvancedDisk type, or MSDP, set the following values:

- **Maximum jobs per client** - Set to a value in accordance to the value of the `Maximum Streams Per Drive * number of drives`.
- Set the **Maximum jobs per policy** - Set to a value higher than the value of `Maximum Streams Per Drive * number of drives`. For example, if you have two drives and **Maximum streams per drive** is set to 3, the total expected jobs to that storage unit is 6.

- Set the logging level.
  See "Setting the logging level" on page 82.
- Set the value of the maximum jobs per client.
  See "Setting maximum jobs per client" on page 82.
- Create an SAP HANA-specific backup policy.
  See "Creating a backup policy for SAP HANA" on page 83.

Setting the logging level

For SAP HANA, for to debug or troubleshoot issues you need to set the logging level to**5 (maximum)**.

To set the logging level:

1. Open the NetBackup web UI.
2. Open the NetBackup web UI.
3. On the left, click **Hosts > Host properties**.
4. Select the primary server. If necessary, click **Connect**. Then click **Edit primary server**.
5. Click **Logging**.
6. For the **Global logging level**, set the logging level to **5 (Maximum)**.
7. After the jobs run correctly, change the logging level to **0 (Minimum)**.

Setting maximum jobs per client

The **Maximum jobs per client** property specifies the maximum number of backup and archive jobs that the NetBackup clients can perform concurrently. The default is one job.

For the SAP HANA software, set this value to 100.

To change the timeout period

1. Open the NetBackup web UI.
2. On the left, click **Hosts > Host properties**.
3. Select the primary server. If necessary, click **Connect**. Then click **Edit primary server**.
4. Click **Global attributes**. Set the **Maximum jobs per client** to a value of **100**.

   See the NetBackup Web UI Administrator's Guide for more information on this setting.
5. Click **Save**.

Creating a backup policy for SAP HANA

This topic describes how to create a new backup policy for the SAP HANA.

To create a backup policy for SAP HANA

1. Open the NetBackup web UI.
2. Sign in with an account that has the RBAC Administrator role or similar permissions.
3. On the left, click **Protection > Policies**.
4. Click **Add**. Then type a unique name for the **Policy name**.
5. Select the **SAP** in the **Policy type** list.
6. Select a disk-based storage unit from the **Policy storage** list.
7. On the **Attributes** tab, select the options **Follow NFS**, **Compress**, and **Encrypt**.

   Note: When you use MSDP or a virtual tape library as the storage unit, do not select the **Compress** and **Encrypt** options.
8. Complete the entries on the **Attributes** tab.

   See "About policy attributes" on page 46.
9. Add other policy information as follows:
   - NetBackup automatically creates a schedule for SAP HANA. Make any changes on the **Schedules** tab.
   - On the **Clients** tab, add all the SAP HANA nodes.
   - Add scripts to the backup selections list.

See “Adding NetBackup for SAP scripts to the backup selections list” on page 51.

10 After you complete adding the required clients and scripts, click **Create**.

Configuring redirected restores

In an SAP HANA multiple-host system, recover data of one host by another host with the following instructions.

To allow a host to redirect a restore

1 On the primary server, create the `altnames` file for the host that you want to have permissions to perform redirected restores.

2 For example, to give `HostB` permissions to restore from another host, create the following file:

- On Windows
  `install_path\NetBackup\db\altnames\HostB`
- On Linux
  `/usr/openv/netbackup/db/altnames/HostB`

3 In the `altnames` file, add the names of the clients whose files the requesting client wants to restore.

For example, if you want `HostB` to have permissions to redirect restores from `HostA`, add `HostA` to the `HostB` file.

Configuring the NetBackup client

The following steps help you configure the NetBackup client:

- Setting the log levels.
- Updating the `initSAP.utl` file.

See “Enable the debug logs manually (SAP) (UNIX/Linux) ” on page 196.

See “Set the debug level on a SAP client (UNIX) ” on page 199.

See “ Updating the initsap.utl file” on page 85.

If you need to un-install or remove the NetBackup client software, refer to the NetBackup Installation Guide.

Updating the initsap.utl file

A default `initSAP.utl` file is provided with NetBackup. The `initSAP.utl` is a text file that contains comments, parameters (required and optional), and parameter values. These parameters determine the backup and restore procedure between NetBackup and the SAP tools.

The `initsap.utl` file that is located at the following path,
`/usr/openv/netbackup/ext/db_ext/sap/scripts/sap_hana/initSAP.utl`

For the SAP HANA software, you must change the name of the policy in the configuration file. Uncomment the policy keyword and specify the correct policy name that is configured on NetBackup primary server to take an SAP HANA backup.

Configuring the SAP HANA Agent

To configure the SAP HANA agent, you must perform a node level configuration. You must also configure the backint file.

Node level configuration

The Database Instance installation creates a directory `/usr/sap/<SID>` containing links to the installation directory.

So the path `/usr/sap/<SID>/SYS/global/hdb/opt/hdbbackint` is the same on each HANA installation, independent of the hardware vendor.

- Manually create a `node_names.txt` file under `/usr/openv/netbackup/ext/db_ext/sap`.
  This file should contain all probable client names (appliance node), one name per line.
  For example:
  ```
  node1.test.lab
  node2.test.lab
  ```
  The SAP HANA admin should update this file whenever a new node is added or removed from the appliance.
  Using the root user, give `+r` permissions to the newly created file `node_names.txt` on every client (appliance node).
- Create a link from SAP HANA's `hdbbackint to hdbbackint_script` script instead of NetBackup's `hdbbackint` binary.
  ```
  ln -s /usr/openv/netbackup/bin/hdbbackint_script
  /usr/sap/<SID>/SYS/global/hdb/opt/hdbbackint
  ```

Note: If the `opt` directory is not present, at `/usr/sap/<SID>/SYS/global/hdb`location, you must create it and then create a `hdbbackint` link as mentioned above.

The HANA DB Master node triggers the backups on all the nodes.

SAP HANA backint configuration

The following configurations must be done from the SAP HANA studio:

- Create `hdbbackint` soft link from `/usr/sap/<SID>/SYS/global/hdb/opt/hdbbackint to /usr/openv/NetBackup/bin/hdbbackint_script` for every database instance.
- The parameter file (`initSAP.utl`) must be specified for data in the SAP HANA database instance configuration. To specify the parameter file, go to **Instance**->**Configuration**-> **global.ini** > **data_backup_parameter_file**.
- The parameter file (`initSAP.utl)` must be specified for logs in the SAP HANA database instance configuration. To specify backup using the parameter file, go to **Instance**->**Configuration**->**global.ini**-> **log_backup_parameter_file**. To specify backup using the `backint` file, go to **Instance**->**Configuration**->**global.ini**-> **log_backup_using_backint**->**SYSTEM = true**.

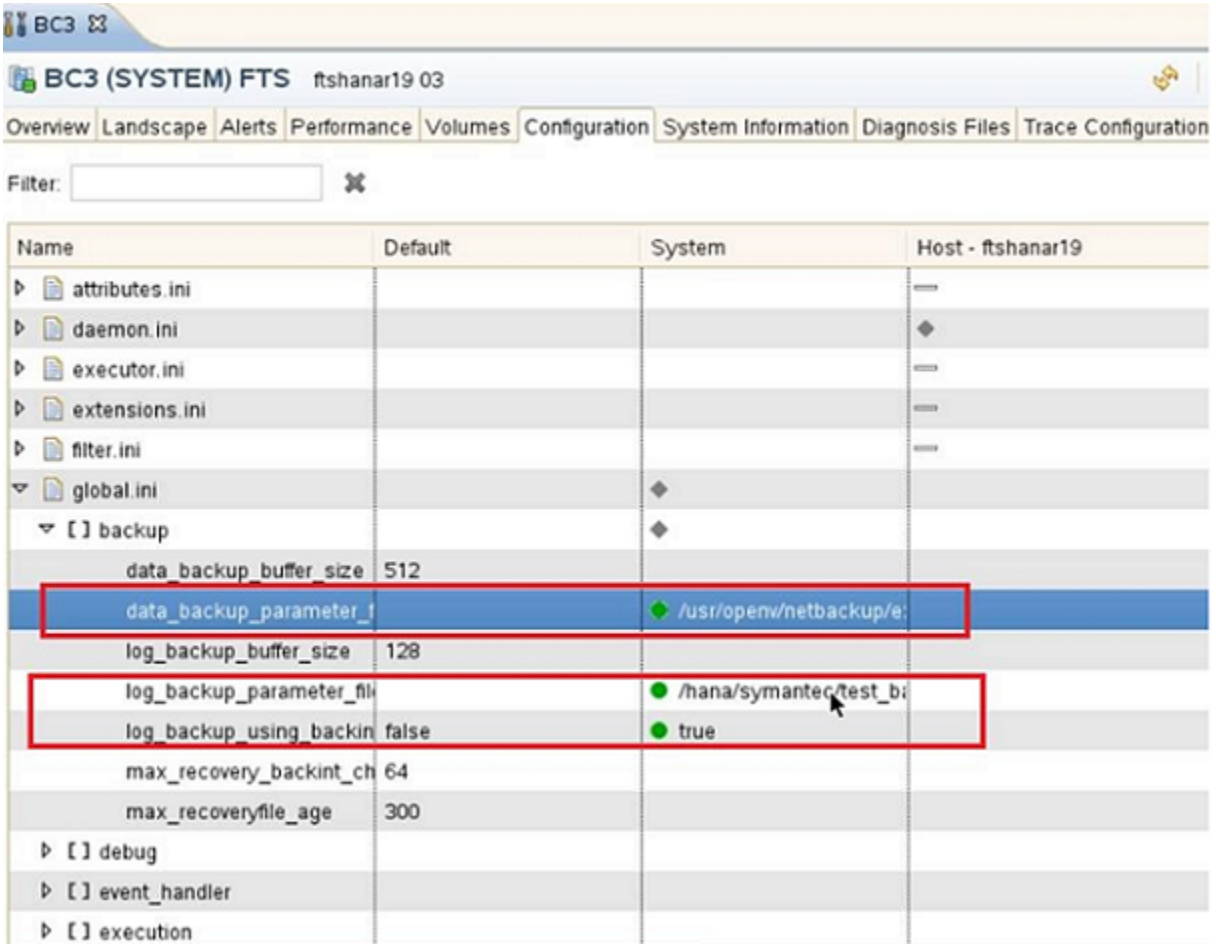

Note: While restoring from the HANA studio, locate log backups if the log backups were written to the file system and subsequently moved. You need to specify their current location. If you do not specify an alternative location for the log backups, the system uses the location where the log backups were first saved. The directory that is specified is searched recursively.

Configuring SAP HANA instance for redirected restore

You can redirect a restore of an SAP HANA instance to another SAP HANA instance.

Note: For additional information about SAP HANA Multi-Tenant Database Containers, see the SAP documentation.

To configure redirected restores of an SAP HANA instance

1. Add all destination instance of SAP HANA hosts in the source instance NetBackup policy.

   Ensure that the backup policy name is the same in the source and the target `initSAP.utl` file. More details are available on how to add the destination host client from the Clients tab.

   See "Creating a backup policy for SAP HANA" on page 83.

2. Add the current source SAP HANA master node name in the destination SAP HANA Master node file. Open the node file from the following location:

   `/usr/openv/netbackup/ext/db_ext/sap/node_names.txt`

   For more details on node-level cofiguration, see See "Configuring the SAP HANA Agent" on page 85.

3. Follow the steps in the Copying a Database Using Third-Party Backup Tools topic in the *SAP HANA Administration Guide*. This topic describes the requirements to use the $(SAPSYSTEMNAME) in the SAP HANA configuration.

Chapter 8

Managing SAP HANA assets and their credentials

This chapter includes the following topics:

About the Workloads > SAP utility in SAP HANA context

NetBackup displays the instances, databases, and availability groups that it discovers in the **Workloads > SAP** node of the NetBackup web UI. It also displays any auto-discovered and manually-added instances and databases that were discovered.

The web UI indicates the intelligent policy name in the **Protected by** column when a backup using the intelligent policy is performed on an individual database.

The **SAP** node contains the following tabs:

- **Instances**
  Contains all the SAP HANA instances that NetBackup discovers or that you manually added.
- **Databases**
  Contains all the SAP HANA databases that NetBackup discovers.

About discovery of SAP HANA objects

NetBackup discovery runs regularly and gathers information for instances in your environment. The data expires after one hour. The NetBackup Discovery Service (`nbdisco`) runs “shallow” discovery every 8 hours for instances on the clients for that primary server. The NetBackup Agent Request Service (NBARS) polls the primary server every 5 minutes for any non-expired data.

By default, this service reports to the primary server when it finds SAP HANA instances. However, the user can turn off discovery for a specific client, with the `bpsetconfig` utility. See the `REPORT_CLIENT_DISCOVERIES` option in the NetBackup Administrator’s Guide, Volume I.

Confirmation messages in the web UI

A message `Starting the discovery of databases...` displays after you click **Discover databases**. This message only indicates that a request was made to start the discovery process. However, database discovery can fail for different reasons. For example, if the instance is not associated with valid credentials or the host cannot be reached. You can consider the deep discovery is successful when the message displays: `Successfully started the discovery of databases. Click Refresh to update the list.`

Discover instances on demand

You can manually start the NetBackup discovery process if you want to immediately discover new SAP HANA instances in your environment.

To discover new SAP HANA instances

1. Click the **Instances** tab.
2. Click **Discover instances**.
3. Select the hosts that contain the instances you want to discover.
4. Click **Discover**.

Discover databases on demand

You have to manually start the NetBackup discovery process to discover databases in your environment.

Note that after you recover a database that was deleted from the SAP HANA environment NetBackup does not immediately reflect the updated database state. To refresh the database state you must run discovery manually.

To discover databases

1. Select the **Databases** tab.
2. Select **Discover databases**.
3. Select the host and the instance that is associated with the databases.

   Note that only registered instances are shown in this list.
4. Select **Discover**.

About registering SAP HANA instances

To allow for full discovery of SAP HANA assets and to protect those assets, you must register instances with a credential.

Instances can be registered by manually registering individual instances.

Registering instances when SAP HANA hosts are clustered or use multiple NICs

When NetBackup discovers a SAP HANA cluster, it adds a single entry on the **Instances** tab. This instance represents all nodes in the cluster. The host name is the virtual name of the SAP HANA cluster. When you add credentials for this

instance NetBackup validates the credentials on the active node. The credentials must be valid for all nodes in the cluster.

When NetBackup discovers a SAP HANA host that uses multiple NICs, it adds an entry using the NetBackup client name on the **Instances** tab. If you installed the NetBackup client using the public interface name, you must configure the NetBackup client name as the private interface name. Then add credentials to the instance with its private interface name. For a SAP HANA cluster that uses multiple NICs, add credentials to the instance with the private virtual name of the SAP HANA cluster.

Validation of credentials

After you register credentials, NetBackup validates the credentials and starts the database discovery.

For a SAP HANA cluster, NetBackup validates the credentials on the active node. The credentials must be valid for all nodes in the cluster. Note that the registered date reflects the date and time the credential was added or updated. It does not indicate if the credentials are valid.

Authentication options for SAP HANA credentials

The NetBackup web UI supports SAP HANA SYSTEMDB database users: SYSTEM, Backup administrator and Backup operator are granted with appropriate privileges.

Register SAP HANA instance with a new credential

To add a new credential for SAP HANA instance

1. Review the recommendations and requirements for the type of credentials that you want to use for authentication.
2. Select the **Instances** tab.
3. Select the check box for each instance that you want to register. Then select **Manage credentials**.
4. Select **Add credentials** and select **Next**.
5. Select the authentication option to register SAP HANA instances.

Add a credential for SAP HANA

1. On the left, select **Credential management**.
2. On the **Named credentials** tab, select the **Add** button.
3. From the **Credential store** options, select **NetBackup**.
4. Select the **Start** button.
5. Provide the following properties:
   - Credential name (for example: *server_credential1*)
   - Tag (for example: *workload name*)
   - Description (for example: This credential is used to access *workload name*)
6. Select the **Next** button.
7. From the **Category** list, select **SAP** .
8. Provide the authentication details that are needed to connect to the SAP HANA.
9. Select the **Next** button.
10. Add one or more RBAC roles that you want to have access to the credential.
    - Select the **Add** button.
    - Select the role name.
    - Select the credential permissions that you want the role to have.
11. Select the **Next** button and follow the prompts to complete the wizard.

View the credential name that is applied to an asset

You can view the named credential that is configured for an asset type. If the credentials are not configured for a particular asset, this field is blank.

Edit or delete a named credential

You can edit the properties for a named credential or delete a named credential NetBackup from the **Credential management**.

Edit a named credential

You can edit a named credential to change the following: credential tag, description, category, authentication details, or permissions. You cannot change the credential name.

To edit a named credential

1. On the left, select **Credential management**.
2. On the **Named credentials** tab, locate and select the check box for the credential that you want to edit.
3. Select **Edit** and update the credential as needed.
4. Review the changes and select **Finish**.

Delete a named credential

You can delete a named credential that you no longer need to use with NetBackup. Be sure to apply another credential to any assets that use the credential you want to delete. Otherwise, backups and restores may fail for those assets.

To delete a named credential

1. On the left, select **Credential management**.
2. On the **Named credentials** tab, locate and select the check box for the credential that you want to delete.
3. Select **Delete > Delete**.

Remove SAP HANA instances

Use this procedure to remove the instances that no longer exist in your environment.

To remove SAP HANA instance

1. Click the **Instances** tab.
2. Locate and select the check box for the instance.
3. Click **Remove**.

Note: If you remove an instance, all assets that are associated with the deleted instance are no longer protected. You can still recover existing backup images, but backups of the instance fail.

Remove SAP HANA databases

Use this procedure to remove the databases that no longer exist in your environment.

To remove a SAP HANA database

1. Click the **Databases** tab.
2. Locate and select the check box for the database.
3. Click **Remove**.

Note: If you remove a database, you can still recover existing backup images. However, backups of the database fail.

Manually add SAP HANA instance

Newly discovered SAP HANA instances are automatically displayed. However, you may not want to wait for the discovery service to discover a new instance. In this case you can add an instance manually.

To manually add a SAP HANA instance

1. On the left, click **Workloads > SAP**, then click the **Instances** tab.
2. Click the **Instances** tab.
3. Click **Add**.
4. Provide the **Host** name where the instance resides and the **Instance name**.
   - For a SAP HANA cluster, the host name is the virtual name of the SAP HANA cluster. You do not need to add each node in the cluster.
   - For a multi-NIC environment, the host name is the private interface name of the SAP HANA host or of the virtual SAP HANA.
   - For a failover cluster instance, enter the instance cluster name. NetBackup enumerates the FCI under the physical node names and the instance cluster name.
5. Click **Next**.
6. Review the roles that have access to the instance. Click **Add** to give additional roles access to the instance.

7 Click **Manage credentials** to add the credentials for this instance.

You may omit credentials at this time. The instance is marked as unregistered and the **Registered** column is empty.

8 Click **Finish**.

Chapter 9

Configuring backups with SAP HANA Server Intelligent Policy

This chapter includes the following topics:

About SAP HANA Intelligent Policies

SAP HANA Intelligent Policy lets you create a single policy to protect multiple SAP HANA instances or the databases in an instance. These instances can be spread over multiple clients. You can select SAP HANA instances for a policy from a list of instances that are automatically discovered in the NetBackup environment.

The SAP HANA Intelligent Policy includes the following criteria:

- Storage unit and media to use

- Policy attributes
- Backup schedules: Full, differential-incremental, cumulative incremental
- The SAP HANA Options: Backint response timeout, Enable multri-stream backups and Maximum parallel backup jobs

Create SAP HANA Intelligent Policy

This topic describes how to create an intelligent policy to protect SAP HANA instances or databases. Other topics cover how to protect clusters.

To create a SAP HANA Intelligent Policy

1. Before you configure an intelligent policy ensure that you have registered the SAP HANA instances that you want to protect.
2. Sign in to the primary server as a user that has the RBAC Administrator role or a role that can manage policies.
3. On the left, select **Protection > Policies**.
4. Select the **Add** button.
5. Type a unique name for the new policy.
6. In the **Policy type** list, select SAP policy.
7. Complete the entries on the **Attributes** tab.
8. Add other policy information as follows:
   - On the **Protect databases** tab, choose how you want to protect SAP HANA. If you choose the **Protect databases** option, you can select individual databases.
   - Add schedules.
   - (Optional) Make changes to any tuning parameters.
9. When you have completed the policy configuration, select **Create**.

   In the **Workloads > SAP** utility, the properties for an instance or database indicate the name of any intelligent policies that protect those objects.

Schedule properties for SAP HANA Intelligent Policies

This topic describes how to configure certain schedule properties for SAP HANA Intelligent Policies. Other schedule properties vary according to your specific backup strategy and system configuration.

Table 9-1 describes how the schedule properties affect a SAP HANA Intelligent Policy.

Table 9-1 Description of schedule properties

Property	Description
Type of backup	Specifies the type of backup that this schedule can control. The selection list shows only the backup types that apply to the policy you want to configure.
Schedule type	You can schedule a backup in one of the following ways: ■ **Frequency** **Frequency** specifies the period of time that can elapse until the next backup operation begins on this schedule. For example, assume that the frequency is 7 days and a successful backup occurs on Wednesday. The next full backup does not occur until the following Wednesday. Typically, incremental backups have a shorter frequency than full backups. The frequency can be hours, days, or weeks. For transaction log backups, the frequency can also be minutes. ■ **Calendar** The **Calendar** option lets you schedule the backup operations that are based on specific dates, recurring week days, or recurring days of the month.
Retention	Specifies a retention period to keep backup copies before they are deleted. The retention period for a schedule controls how long NetBackup keeps records of when scheduled backups occurred. Set the time period to retain at least two full backups of your database. In this way, if one full backup is lost, you have another full backup to restore. The type of schedule you select affects the retention period as follows: ■ Frequency-based scheduling Set a retention period that is longer than the frequency setting for the schedule. For example, if the frequency setting is set to one week, set the retention period to be more than one week. When NetBackup expires a backup image it does not notify SAP HANA. Use SAP HANA to periodically delete expired backup sets from the SAP HANA repository. ■ Calendar-based scheduling The retention period setting is not significant for calendar-based scheduling.

Table 9-1 Description of schedule properties *(continued)*

Property	Description
Media multiplexing	Multiplexing is useful if you have many simultaneous backups using the same tape drive. However, it can interfere with SAP HANA recovery due to how SAP HANA requests streams during restore. In most cases, Cohesity does not recommend multiplexing multiple SAP HANA streams from the same backup to a single tape.

Schedule backup types for SAP HANA Intelligent Policies

The **Type of backup** attribute specifies the type of backup that the schedule controls.

Refer to the following guidelines when you configure schedules:

- The backup operation is skipped for a specific database if the database recovery model is not supported for the selected backup type.
- If a differential backup runs and a full backup do not already exist for the database or filegroup, NetBackup can convert the backup to a full backup. Similarly, NetBackup can convert transaction log backups if a full backup does not already exist for the database.

Table 9-2 Schedule backup types for SAP HANA Intelligent Policies

Backup type	Description
Full backup	A complete backup of the database that contains all of the data files and the log file.
Cumulative incremental backup	Includes all changes since the full backup, regardless of any previous incremental backups
Differential incremental backup	The data changed since the last full data backup (complete data backup or data snapshot).
Log backup	Backs up the transactions that have occurred since the last log backup.

Add databases to a policy

This topic describes how to add databases to a policy when you choose the **Protect databases** option.

To add databases to a policy

1. In the policy, select the **Databases** tab.
2. Select **Protect databases**.
3. Select **Add**.

   All instances that you registered are displayed.
4. Select the check box next to each database that you want to add to the list.

   When you select individual databases, you must manually add any new databases in your environment to a policy. In this case, NetBackup does not dynamically create a list of databases at run-time.

   For the databases that are hosted on a SAP HANA cluster, the **Host name** represents the virtual name of the SAP HANA.

Performance tuning and configuration options

The **SAP** tab in SAP HANA backup policy contains the tuning parameters that can improve the performance of your backups. The settings that affect performance are discussed in this topic.

Table 9-3 Tuning parameters for SAP HANA backups

Field	Description
Backint response timeout	This SAP HANA option will wait up to 10 minutes for the backup agent to respond before timing out. The default value is 600 seconds (10 minutes).
Enable multiple-streams backups	This option divides the backup operation into multiple concurrent streams. A stream corresponds to a job in the activity monitor. For example, if the value is **3**, each database is backed up using three jobs. This configuration applies in any situation in which SAP HANA dumps data faster than your tape drive is capable of writing. The number of backup channels is 1. The range is 1–32. The default value for the threshold to enable parallel backups is 128 GB.
Maximum parallel backup jobs	This option is the number of backup operations to start simultaneously, per database instance. The range is 1–32. The default is 1. You may need to configure other options when you configure two or more parallel backup operations.

Perform a manual backup of SAP HANA instance

After you configure the servers and assets in your environment, you can test the configuration settings with a manual backup. Perform a manual backup (or backups) from a policy with the automatic backup schedules that you created.

Alternatively, you can use Backup now to perform a manual backup of an asset in the **Workloads** node.

To perform a manual backup from a policy

1. On the left, select **Protection > Policies**.
2. Select the policy you want to test.
3. Select **Manual backup**.
4. Select the schedule that you want to use for the manual backup.
5. For SAP HANA Intelligent Policies, select the databases or instances that you want to include for the manual backup.

Backing up SAP HANA instances and databases using web UI

Use the NetBackup web UI to back up SAP HANA instances and databases.

To discover SAP HANA instances

1. In the NetBackup web UI, on the left, select **SAP**.
2. In the **Instances** tab, click **Discover instances**.
3. On the **Discover instances** page, select the required client and click **Discover**.
4. Refresh the **Instances** tab to see the new instances that are added for the selected client.
5. Select the instance and click **Manage credentials**.
6. On the **Manage credentials** page, click **Add credentials**. Enter the following information:
   - Credential name - Name for the credential, for example: SAP credential
   - Tag - Tag for the credential
   - Description - Information on credential
   - Credentials for SAP database - SAP user name and SAP user password
7. Click **Next**. The credentials are validated.

8 Click **Close**.

9 In the **Databases** tab, click **Discover databases**.

10 On the **Discover databases** page, select the required databases > **Hosts and Instances**, and click **Discover**.

11 Refresh the **Databases** tab to see the new databases that are added.

12 On the left, select **Protection > Policies** and click **Add**.

13 On the **Create policy** page, on the **Attributes** tab, specify the policy information.

14 Select **SAP** as the policy type.

Instead of the **Clients** tab, the **Protect databases** tab is displayed.

The **Backup selections** tab is disabled and a new tab called **SAP** is displayed.

15 In the **Schedules** tab, specify the details of the backup schedule for the SAP data backup.

See "Schedule properties for SAP HANA Intelligent Policies" on page 99.

16 In the **Protect databases** tab, click the **Protect databases** option and click **Add**.

17 On the **Select database** page, select the databases that you want to protect and click **Select**.

18 Click **Save**.

19 Select the SAP policy that you have created and click **Manual backup**.

20 On the **Manual backup** page, select the required schedule and the client, and click Backup.

21 On the left, select **Activity monitor** to see the SAP backup job details.

Chapter 10

Performing backups and restores of SAP

This chapter includes the following topics:

Using NetBackup for SAP to perform a backup

You can use the following methods to perform a backup with NetBackup for SAP:

- Manual SAP backups
  You can use the NetBackup server software to manually run an automatic backup schedule for the SAP policy.
  See "Perform a manual backup" on page 77.
- Automatic SAP backups

The most convenient way to back up your database is to set up schedules for automatic backups. When the NetBackup scheduler invokes a schedule for an automatic backup, the SAP scripts start the database backup.

The SAP scripts run as follows:

- In the same order as they appear in the file list
- On all clients that have them (that is, matching path names)

There are instructions that explain how to create and configure a policy that uses the SAP scripts to back up your database automatically according to the schedules you specify.

- Client-initiated SAP backup
  The SAP tools on the SAP host can, either manually or via a non-NetBackupscheduler, initiate a backup to the NetBackup backint or libsapora/orasbt APIs.

Using BRTools to start an SAP backup (for Oracle database only)

Note: This information applies to NetBackup for SAP on Oracle databases only. A different configuration is required for NetBackup for SAP on a MaxDB database.

See "About NetBackup for SAP on MaxDB databases" on page 188.

You can start a backup in one of the following ways.

- Through the `SAPDBA` utility's menus
- Through the `brbackup` command line

When you use the SAPDBA utility or the brbackup command to start a backup, the following events occur:

- `brbackup` status messages appear on the console. These messages report when the database server is started or stopped. They also report when the backup mode of the tables is changed.
- The `brbackup` command starts the NetBackup for SAP `backint` interface, and it submits files to be backed up.
- The NetBackup for SAP `backint` interface processes input files and calls the XBSA interface. During the file-online mode, each database file is backed up, one at a time. The NetBackup for SAP `backint` interface coordinates with brbackup using a semaphore file.

- After all files are backed up, the full file list is displayed. The file list is displayed in the format that the NetBackup for SAP `backint` interface specification requires as to success or failure. This format includes a backup ID (BID) to be used for later restores. The SAP tools maintain a log of the backup session to keep track of the BID. This allows cross-referencing by `brrestore`.

Performing an SAP offline backup with SAPDBA

The following procedure shows you how to perform an offline backup with SAPDBA.

To perform an offline backup with SAPDBA

1. Verify that you have completely configured NetBackup, NetBackup for SAP, and the SAP tools.

   See "About configuring NetBackup for SAP" on page 43.

2. Verify that you are user `sapadm`.
3. Run the `stopsap R3` command to stop SAP.
4. Start `sapdba`.
5. Type `h` to select the `Backup database` menu item.
6. Type `d` to select `Objects for backup`.
7. Type `g` and type the name of the object to back up.

   For example: `PSAPUSER1D`.

8. Type `q` to go back to the previous menu.

   The information that is returned indicates the following:

   - The backup device type is `external backup tool (backint)`.
   - The backup type is `offline_force`.
   - Tablespace is `PSAPUSER1D`.

9. Type `S` to start the backup.

Performing brbackup online SAP backups

You can use `brbackup` instead of `sapdba` to perform SAP database backups. This section shows an online backup with dynamic switching of files into and out of backup mode. You can change the `backup_type` and `backup_dev_type` by changing the `initSAP.sap` parameter file or by specifying `-t online` and and `-d util_file_online` on the `brbackup` command. Here is what these changes look like in the `initSAP.sap` file:

```
backup_type = online

backup_dev_type = util_file_online
```

This backup method allows `sapdba` or `brbackup` to use the switch files with the NetBackup for SAP `backint` interface. This method provides a better online backup for very large files because only the necessary tablespaces are placed in backup mode. When NetBackup is ready to process another file, it notifies `brbackup` to take the completed file out of backup mode and put the next file into backup mode.

To perform brbackup online SAP backups

1. Verify that you have completely configured NetBackup, NetBackup for SAP, and the SAP tools.

   See "About configuring NetBackup for SAP" on page 43.

2. Verify that you are user `sapadm`.

3. Enter the brbackup command.

   - To initiate an online backup, enter the following command:

     ```
     brbackup -d util_file_online -t online -m all
     ```

   - To initiate a full backup when you use RMAN, enter the following command:

     ```
     brbackup -d rman_util -t online -m full -c
     ```

   - To initiate an incremental backup when you use RMAN, enter the following command:

     ```
     brbackup -d rman_util -t online -m incr -c
     ```

Performing an SAP archive

An archive is performed in a fashion similar to a backup, but using `brarchive` instead of `brbackup`, or alternatively using `brbackup -a`. You can use the NetBackup multiple copies feature or `brarchive` command to create multiple copies that are stored on separate media or media servers. Depending on how you use NetBackup for SAP, create archive log backups copies with the method that is easiest for you.

The following examples show how to backup and restore an SAP archive:

Example 1 Issue the following command to create the SAP archive:

```
brarchive -d util_file -s
```

Example 2

To restore the first available copy of an archive log file, issue the following command:

```
brrestore -d util_file -a log_number
```

where *log_number* is the number of the archive log you want to restore.

Using the NetBackup multiple copies feature to create an SAP archive

If the option for multiple copies is enabled in the backup schedule used for the job, NetBackup automatically makes up to four copies of a backup. For more information on the multiple copies feature, see the NetBackup Administration Guide, Volume 1.

Using two brarchive commands in sequence to create an SAP archive

You can also create two copies of the archive log by issuing two `brarchive` commands in sequence.

For the first copy, call `brarchive` with the `-s` option. For the second copy, call `brarchive` with the `-sc` option. Both commands automatically determine that the archive log file needs to be backed up and perform the backup.

Be sure that each brarchive command uses either a separate utl file or separate environment variable values where either the policy or schedule name differ. The policy or schedule should identify the unique media or storage units to use to ensure that the second copy is not written to the same media as the first copy (and subject to loss at the same time). For example, you can write each backup copy to two different volume pools, to disk and tape, or to two different DSUs.

The following example shows how two `brarchive` commands are used in sequence to create two copies of an SAP archive. This example assumes that only one archive log is to be backed up, archive log `77`.

To create the backup copies	To create the first backup copy, issue the following commands:
	`SAP_SCHED=To-DSU ; export SAP_SCHED`
	`brarchive -d util_file -s`
	To create the second backup copy, issue the following commands:
	`SAP_SCHED=To-Pool1 ; export SAP_SCHED`
	`brarchive -d util_file -sc`
	On Windows, the syntax to set the environment variables is:
	`SET SAP_SCHED=To-DSU`

To restore the archived log files	To restore from the first backup copy, issue the following command:
	`brrestore -a 77`
	To restore from the second backup copy, issue the following command:
	`brrestore -a2 77`

Performing SAP restores using the SAPDBA utility (SAP on Oracle databases only)

To restore a partial database or full Oracle database, you must have a list of valid restores. Use the `sapdba` system to maintain the list of restores for specific tablespace or complete database restores.

When the user restores either individual tablespaces or full databases, the user is prompted before an existing copy of the target file is deleted. `sapdba` then invokes the `brrestore` command.

`brrestore` submits the BID and file name list to the NetBackup for SAP `backint` interface. The `backint` interface cross-references the exact date and time to when the backup was made and uses NetBackup to recover the file. The `backint` interface monitors the progress of the restore and reports status back to `brrestore`.

Upon completion, the `backint` interface saves a copy of the NetBackup restore logs for auditing purposes. `sapdba` then provides required database recovery, such as media recovery, and restarts the database server.

To use brrestore to restore database files

1. In Oracle-based environments without RMAN, use the command with the following options:

```
brrestore -d util_file -b last -m full -c force
```

2. In Oracle-based environments with RMAN, use the command with the following options:

```
brrestore -d rman_util -b last -m full -c force
```

Redirecting a restore to a different client (SAP on Oracle databases only)

You have the option to restore an SAP database to a client other than the one that originally supplied the backup. This process to restore data to a different client is called a redirected restore.

Note: In some situations, you can restore from lower to higher release levels of Oracle. For more information, see your Oracle documentation.

To perform a redirected restore, the following conditions must be present:

- The source client and destination client must have identical computer system architectures, OS versions, and bit levels.
- The source client and destination client must have identical Oracle RDBMS levels.

The redirected restore has the following process:

- Configure the NetBackup server.
- Copy files from the source client to the destination client.
- Configure the NetBackup for SAP destination client.
- Perform the redirected restore.

The following procedures explain each part of the process. For more information on redirected restores, see the NetBackup Administrator's Guide, Volume 1.

To configure the NetBackup server

1. Verify that you are logged on as the NetBackup administrator.
2. Configure the NetBackup server to allow redirected restores.

- Allow restores from any client.
  Continue with step 3.
- Allow restores from only selected clients.
  Continue with step 5.

3 To allow any client to restore the backups of any other client, create the following file on the NetBackup primary server:

  Windows:

  `install_path\NetBackup\db\altnames\No.Restrictions`

  UNIX or Linux:

  `/usr/openv/netbackup/db/altnames/No.Restrictions`

4 Continue with step 7.

5 To allow only the destination client to restore from all other clients, create the following file:

  Windows:

  `install_path\NetBackup\db\altnames\client_name`

  UNIX or Linux:

  `/usr/openv/netbackup/db/altnames/client_name`

  where *client_name* is the name of the destination client that is allowed to perform the redirected restore.

6 If you performed step 5, and want to restrict the destination client to only restore from the source client, add the name of the NetBackup for SAP source client to the *client_name* file.

7 On the destination client, back up the existing control files to preserve the original configuration.

  These files are usually in the `$ORACLE_HOME/dbs` (`%ORACLE_HOME%\database`) directory.

To configure the NetBackup for SAP destination client

1 Change the client parameter on the destination client.

  Do one of the following:

  - Change the client parameter in the `initSID.utl` file on the destination client to the name of the source client.

- Set the `SAP_CLIENT` environment variable on the destination client to the name of the source client. For example:

  ```
  SAP_CLIENT=source_client
  export SAP_CLIENT
  ```

  The `SAP_CLIENT` variable has higher priority than the value in the `initSID.utl` file.

2 (Windows) On the destination client, create a file path like the one on the source client.

For example:

- Source client actual path is:

  ```
  C:\Oracle\SAP
  ```

- Destination client actual path is:

  ```
  C:\Oracle\SAP
  ```

3 (UNIX or Linux) Create a symbolic link on the destination client that resembles the actual path of the source client. Point it to the directory that points to the SAP files.

To copy the configuration and the control files (UNIX or Linux)

- Copy configuration and control files from the NetBackup for SAP source client to the NetBackup for SAP destination client.
  - Copy the following files from the source client, usually `$ORACLE_HOME/dbs`, to the same location on the destination client:

    ```
    initSID.sap
    initSID.utl
    initSID.ora
    initSID.dba
    ```

  - Copy the `$SAPDATA_HOME/sapbackup` directory from the source client to the same location on the destination client.

To copy the configuration and the control files (Windows)

- Copy configuration and control files from the NetBackup for SAP source client to the NetBackup for SAP destination client.
  - Copy the following files from the source client, usually `%ORACLE_HOME%\database`, to the same location on the destination client:

```
initSID.sap
initSID.utl
initSID.ora
initSID.dba
```

- Copy the `%SAPDATA_HOME%\sapbackup` directory from the source client to the same location on the destination client.

To perform a NetBackup for SAP redirected restore

1. Log onto the NetBackup for SAP destination client.
2. Enter the following command:

   ```
   brrestore -d util_file -b last -m full
   ```

   Do not use the `-m dest_dir` option with `brrestore`. This option restores the files to the directory you specify, rather than to the original directories.

SAP redirected restore sample environment (UNIX)

This sample environment shows how a redirected restore uses the `restore_filter` script.

This example assumes the following:

- Source client
  - Actual path is `/home_db/oracle/sap`
  - Includes soft link `/oracle/SAP` (`SAP` is the system ID), which points to `/home_db/oracle/sap`.
- Destination client
  - Actual path is `/home2/sap`
  - Also has soft link `/oracle/SAP`, which points to `/home2/sap`.
  - Make soft link `/home_db/oracle/sap`, which points to `/home2/sap`

In the `restore_filter` script, substitute the path that was provided in the input file list of the destination client with the actual path of the source client. The resulting script sample might look like the following:

```
#!/bin/sh
# this shell is used to change some logically linked files
```

```
# during a restore
sed -e 's#/oracle/SAP/#/home_db/oracle/sap/#' $1 > $2
```

When the backup starts, the file list contains the file path with the following soft link:

```
/oracle/SAP/sapdata1/btabd_1/btabd.data1
```

However, the file is backed up with the following actual path:

```
/home_db/oracle/sap/sapdata1/btab_d/btabd.data
```

When the request for a restore is issued, the input file list contains file paths with soft links. These are converted to the actual path of the destination client. Because this path is different from the source client path, the restore would fail in the inquire phase. To prevent this kind of failure, use the `restore_filter` script.

Disaster recovery using BRRECOVER

When you run disaster recovery using BRRECOVER, if the `init`*`SID`*`.sap` file is not present, BRRECOVER calls BACKINT without the “-p init<SID>.utl” parameter.

If NetBackup for SAP is called without the `util` file parameter, it checks for the SAP_RECOVERY environment variable. If the environment variable is set to “disaster”, the restore is allowed without `util` file parameter. NetBackup for SAP then restores using the default SAP policy.

Restarting failed NetBackup for SAP backups and restores

A backup or a restore can fail or terminate because of system error or resource contention. For example, a backup can fail because of the following reasons: the tape library ran out of tapes, the storage unit ran out of space, or a network connection was lost. In the case of large environments with more than a terabyte of data, it might not be feasible to repeat the backup all over again.

If the policy is enabled for checkpoint restart, SAP tools can resume the backup from a point shortly before the failure. If enabled, the default checkpoint interval is 15 minutes unless adjusted on the backup policy’s attributes tab. For a failed job, NetBackup retains partial image and catalog entries in the NetBackup catalog for the files that have been already backed up.

SAP tools can also restart a restore of just the files that were not previously restored. This ability is independent of the checkpoint restart option in the backup policy.

Caution: Attempts to restart a backup or restore from the Activity Monitor will fail.

Note: SAP tools use `backint` to restart failed backups and restores of Oracle databases.

Sap tools do not support the use of RMAN to restart failed backups and restores of Oracle databases. MaxDB also cannot restart failed backups or restores.

Restarting NetBackup for SAP backups

The `brbackup` command initiates the backup for SAP on Oracle databases. The `brbackup` command calls the NetBackup for SAP agent (`backint`) with appropriate options and gives a list of files for backup. Backups can be either successful or unsuccessful.

The results are as follows:

- For a successful backup:
  For each file that is backed up successfully, `backint` sends the following message to `brbackup`:

  ```
  #SAVED BID filename
  ```

  For each `SAVED ...` message that `brbackup` receives, `brbackup` writes the following information to its log file:
  UNIX or Linux:

  ```
  #FILE..... /home1/orasap/inp/file1

  #SAVED.... VXF1134574401
  ```

  Windows:

  ```
  #FILE..... C:\oracle\SUD\sapdata1\file1.dbf

  #SAVED.... VXF1134574401
  ```

- For an unsuccessful backup:
  For each file that is not backed up successfully, `backint` sends the following message to `brbackup`:

  ```
  #ERROR filename
  ```

  For each `ERROR ...` message that `brbackup` receives, `brbackup` writes the following message to its log file:
  UNIX or Linux:

```
BR0233E Backup utility has reported an error while saving

file /home/orasap/inp/filename
```

Windows:

```
BR0233E Backup utility has reported an error while saving

file C:\oracle\SUD\sapdata1\file1.dbf
```

A backup can fail for several reasons that are unrelated to NetBackup operations:

- The tape library ran out of tapes
- A network connection problem occurred
- The user killed the backup but wants to restart it sometime later

For situations like these, NetBackup for SAP saves the partial image for the files that have been already backed up.

Catalog entries in the NetBackup catalog are also saved. It returns the following messages:

- `#SUCCESS BID filename` - For the files that were backed up
- `#ERROR filename` - For the files that were not backed up

If a backup has terminated because of any problems that are not related to NetBackup, you do not have to repeat a complete backup again. Use the `brbackup` command and the `-f logfile` parameter. The `brbackup` command automatically determines the files that still need to be backed up. The next section describes how to use the `-f` parameter to restart a backup.

Configuring the checkpoint restart option in a NetBackup for SAP policy

To restart backups from the point of failure, you must enable checkpoint in your NetBackup for SAP policy. To enable this feature, in the attributes for the policy, select **Take Checkpoints Every**. By default, this option is disabled. When this option is not enabled, a failed backup that is based on this policy is restarted from the beginning of the job.

When **Take Checkpoints Every** is enabled, by default the NetBackup for SAP agent takes a checkpoint every 15 minutes. You can configure this time interval for a different duration.

Restarting a NetBackup for SAP backup with the brbackup -f command

You can use the `brbackup` command's `-f` option to back up only the files that failed to be backed up. You do not have to specify that all files be backed up again.

The `brbackup` command examines the specified file. It determines the files that were backed up successfully and the files that were not backed up successfully. `brbackup` sends `backint` the list of files that still need to be backed up. `brbackup` starts a new job only for the files that still need to be backed up.

Example 1. The following command specifies a log file:

```
brbackup -d util_file -m all -c -f  bdprcayp.aff
```

Example 2. The following command specifies the last backup job:

```
brbackup -d util_file -m all -c -f last
```

To restart a NetBackup for SAP backup with the brbackup -f command

- Run the `brbackup` command with the `-f` *`logfile`* parameter.

  For log file, specify one of the following:

  - The log file name of the failed job. When the log file is specified, `brbackup` checks it for the files that were not backed up successfully.
  - The keyword `last`. When this keyword is specified, `brbackup` checks the status of the last backup job for the files that were not backed up successfully.

Example: Restarting a failed NetBackup for SAP backup job (UNIX or Linux)

This example shows how to restart a failed backup job on UNIX or Linux. In this example, assume the following:

- The backup job was deliberately cancelled from the activity manager. In a real user case, a job might fail for different reasons.
- You enabled checkpoint (the **Take Checkpoints Every** option in the policy).
- You ran the backup according to the following `brbackup` command:

  ```
  brbackup -d util_file -t offline -m all -c force
  ```

The following is an excerpt from the output log:

```
BR0280I BRBACKUP time stamp: 2006-02-28 21.07.59
BR0057I Backup of database: SUD
BR0058I BRBACKUP action ID: bdsbecnl
BR0059I BRBACKUP function ID: aff
BR0110I Backup mode: ALL
BR0077I Database files for backup:

/oracle/SUD/origlogA/log1_m1.dbf
/oracle/SUD/origlogB/log2_m1.dbf
/oracle/SUD/origlogA/log3_m1.dbf
/oracle/SUD/origlogB/log4_m1.dbf
/oracle/SUD/saparch/cntrl/cntrlSUD.dbf

BR0061I 29 files found for backup, total size 853.604 MB
```

The previous output log shows the 29 files for backup. `brbackup` calls the NetBackup SAP agent for backup.

In this example, we have configured the checkpoint frequency for 5 minutes. By the time first checkpoint was taken (that is, in the first 5 minutes), 13 files were backed up. The user killed the job from the NetBackup activity monitor after NetBackup took the first checkpoint. `backint` returned `#SUCCESS` messages to `brbackup` for 13 files that were backed up. For the files that were yet to be backed up, `backint` retuned `#ERROR`. The following are the output messages from the `brbackup` log:

```
BR0280I BRBACKUP time stamp: 2006-02-28 21.08.04
BR0229I Calling backup utility with function 'backup'...

BR0278I Command output of '/usr/sap/SUD/SYS/exe/run/backint -u SUD -f backup -I /oracle/
SUD/sapbackup/.bdsbecnl.lst -t file -p /home1/orasap/ora920/dbs/initSUD.utl
-c':

*********************************************************************

Program:                      /usr/sap/SUD/SYS/exe/run/backint 6.5
Input File:                   /oracle/SUD/sapbackup/.bdsbecnl.lst
Profile:                      /home1/orasap/ora920/dbs/initSUD.utl

Program:                      backint 6.5
Input File:                   C:\oracle\SUD\sapbackup\.bdsbedhj.lst
Profile:                      C:\oracle\ora920\dbs\initSUD.utl

Function:                        BACKUP
Backup Type:                  BACKUP_FILE
```

```
**************************************************************
BR0280I BRBACKUP time stamp: 2006-02-28 21.15.50

#FILE..... /oracle/SUD/sapdata1/btabd_1/btabd.data1

#SAVED.... VXF1141141105

BR0280I BRBACKUP time stamp: 2006-02-28 21.15.50

#FILE..... /oracle/SUD/sapdata2/btabi_1/btabi.data1

#SAVED.... VXF1141141105

….
….
BR0233E Backup utility has reported an error while saving file

/oracle/SUD/saparch/cntrl/cntrlSUD.dbf

**************************************************************
BR0280I BRBACKUP time stamp: 2006-02-28 21.15.50

BR0279E Return code from '/usr/sap/SUD/SYS/exe/run/backint -u SUD -f backup -I /oracle/
SUD/sapbackup/.bdsbecnl.lst -t file -p /home1/orasap/ora920/dbs/initSUD.utl
-c': 2

BR0232E 13 of 29 files saved by backup utility
BR0280I BRBACKUP time stamp: 2006-02-28 21.15.50
BR0231E Backup utility call failed

BR0280I BRBACKUP time stamp: 2006-02-28 21.15.52
BR0304I Starting and opening database instance SUD ...

BR0280I BRBACKUP time stamp: 2006-02-28 21.16.00
BR0305I Start and open of database instance SUD successful

BR0056I End of database backup: bdsbecnl.aff 2006-02-28 21.15.50
BR0280I BRBACKUP time stamp: 2006-02-28 21.16.00
BR0054I BRBACKUP terminated with errors
```

You can then restart the preceding job with following command:

```
brbackup -d util_file -t offline -m all -f last -c force
```

The `brbackup` command examined the last backup log and found that the backup failed for some files. It writes following messages in the beginning of output log for this session:

```
BR0051I BRBACKUP 6.40 (22)
BR0055I Start of database backup: bdsbedhj.aff 2006-02-28 21.16.35

BR0453W Error message found in /oracle/SUD/sapbackup/bdsbecnl.aff:

'BR0233E Backup utility has reported an error while saving file
/oracle/SUD/sapraw/rawDev4'
BR0453W Error message found in /oracle/SUD/sapbackup/bdsbecnl.aff:
'BR0233E Backup utility has reported an error while saving file
/oracle/SUD/sapraw/rawDev5'

….
….

BR0453W Error message found in /oracle/SUD/sapbackup/bdsbecnl.aff:
'BR0279E Return code from '/usr/sap/SUD/SYS/exe/run/backint -u SUD -f backup -I
/oracle/SUD/sapbackup/.bdsbecnl.lst -t file -p /home1/orasap/ora920/dbs/initSUD.utl
-c': 2'

BR0453W Error message found in /oracle/SUD/sapbackup/bdsbecnl.aff:

'BR0232E 13 of 29 files saved by backup utility'

BR0453W Error message found in /oracle/SUD/sapbackup/bdsbecnl.aff:

'BR0231E Backup utility call failed'

BR0459W Backup /oracle/SUD/sapbackup/bdsbecnl.aff terminated with errors

….
….
BR0280I BRBACKUP time stamp: 2006-02-28 21.16.36
BR0057I Backup of database: SUD
BR0058I BRBACKUP action ID: bdsbedhj
BR0059I BRBACKUP function ID: aff
BR0110I Backup mode: ALL
BR0077I Database files for backup:

/oracle/SUD/origlogA/log1_m1.dbf
/oracle/SUD/origlogB/log2_m1.dbf
/oracle/SUD/origlogA/log3_m1.dbf
/oracle/SUD/origlogB/log4_m1.dbf
/oracle/SUD/saparch/cntrl/cntrlSUD.dbf

BR0061I 16 files found for backup, total size 283.502 MB
BR0091I 13 files were already saved in: bdsbecnl.aff
```

As the preceding output log shows, 13 files were already backed up. The following is the output of a successful `brbackup` restart job. This job is the one that backed up the remaining 16 files:

```
BR0280I BRBACKUP time stamp: 2006-02-28 21.16.41
BR0229I Calling backup utility with function 'backup'...

BR0278I Command output of '/usr/sap/SUD/SYS/exe/run/backint -u SUD -f backup -I
/oracle/SUD/sapbackup/.bdsbedhj.lst -t file -p /home1/orasap/ora920/dbs/initSUD.utl

-c':
*********************************************************************

Program:                      /usr/sap/SUD/SYS/exe/run/backint 6.5
Input File:                   /oracle/SUD/sapbackup/.bdsbedhj.lst
Profile:                      home1/orasap/ora920/dbs/initSUD.utl

Function:                     BACKUP
Backup Type:                  BACKUP_FILE
*********************************************************************

BR0280I BRBACKUP time stamp: 2006-02-28 21.21.38

#FILE..... /oracle/SUD/sapraw/rawDev4

#SAVED.... VXF1141141607

….
….
BR0280I BRBACKUP time stamp: 2006-02-28 21.21.38

#FILE..... /oracle/SUD/saparch/cntrl/cntrlSUD.dbf

#SAVED.... VXF1141141607
*********************************************************************

BR0280I BRBACKUP time stamp: 2006-02-28 21.21.38
BR0232I 16 of 16 files saved by backup utility
BR0230I Backup utility called successfully
```

Example: Restarting a failed NetBackup for SAP backup job (Windows)

This example shows how to restart a failed backup job. In this example, assume the following:

- The backup job was deliberately cancelled from the activity manager. In a real user case, a job might fail for different reasons.
- You enabled checkpoint (the **Take Checkpoints Every** option in the policy).
- You ran the backup according to the following `brbackup` command:

```
brbackup -d util_file -t offline -m all -c force
```

The following is an excerpt from the output log:

```
BR0280I BRBACKUP time stamp: 2006-02-28 21.07.59
BR0057I Backup of database: SUD
BR0058I BRBACKUP action ID: bdsbecnl
BR0059I BRBACKUP function ID: aff
BR0110I Backup mode: ALL
BR0077I Database files for backup:

C:\oracle\SUD\origlogA\log1_m1.dbf
C:\oracle\SUD\origlogB\log2_m1.dbf
C:\oracle\SUD\origlogA\log3_m1.dbf
C:\oracle\SUD\origlogB\log4_m1.dbf
C:\oracle\SUD\saparch\cntrl\cntrlSUD.dbf

BR0061I 29 files found for backup, total size 853.604 MB
```

The previous output log shows the 29 files for backup. `brbackup` calls the NetBackup SAP agent for backup.

In this example, we have configured the checkpoint frequency for 5 minutes. By the time first checkpoint was taken (that is, in the first 5 minutes), 13 files were backed up. The user killed the job from the NetBackup activity monitor after NetBackup took the first checkpoint. `backint` returned `#SUCCESS` messages to `brbackup` for 13 files that were backed up. For the files that were yet to be backed up, `backint` retuned `#ERROR`. The following are the output messages from the `brbackup` log:

```
BR0280I BRBACKUP time stamp: 2006-02-28 21.08.04
BR0229I Calling backup utility with function 'backup'...

BR0278I Command output of 'backint.exe -u SUD -f backup -I
C:\oracle\SUD\sapbackup\.bdsbecnl.lst -t file -p C:\oracle\ora920\dbs\initSUD.utl
-c':

*********************************************************************
```

```
Program:                      backint 6.5
Input File:                   C:\oracle\SUD\sapbackup\.bdsbedhj.lst
Profile:                      C:\oracle\ora920\dbs\initSUD.utl

Function:                        BACKUP
Backup Type:                BACKUP_FILE
*****************************************************************
BR0280I BRBACKUP time stamp: 2006-02-28 21.15.50

#FILE..... C:\oracle\SUD\sapdata1\btabd_1\btabd.data1

#SAVED.... VXF1141141105

BR0280I BRBACKUP time stamp: 2006-02-28 21.15.50

#FILE..... C:\oracle\SUD\sapdata2\btabi_1\btabi.data1

#SAVED.... VXF1141141105

….
….
BR0233E Backup utility has reported an error while saving file

C:\oracle\SUD\saparch\cntrl\cntrlSUD.dbf

*****************************************************************
BR0280I BRBACKUP time stamp: 2006-02-28 21.15.50

BR0279E Return code from 'backint -u SUD -f backup -I
C:\oracle\SUD\sapbackup\.bdsbecnl.lst -t file -p C:\oracle\ora920\dbs\initSUD.utl
-c': 2

BR0232E 13 of 29 files saved by backup utility
BR0280I BRBACKUP time stamp: 2006-02-28 21.15.50
BR0231E Backup utility call failed

BR0280I BRBACKUP time stamp: 2006-02-28 21.15.52
BR0304I Starting and opening database instance SUD ...

BR0280I BRBACKUP time stamp: 2006-02-28 21.16.00
BR0305I Start and open of database instance SUD successful

BR0056I End of database backup: bdsbecnl.aff 2006-02-28 21.15.50
BR0280I BRBACKUP time stamp: 2006-02-28 21.16.00
BR0054I BRBACKUP terminated with errors
```

You can then restart the preceding job with following command:

```
brbackup -d util_file -t offline -m all -f last -c force
```

The `brbackup` command examined the last backup log and found that the backup failed for some files. It writes following messages in the beginning of output log for this session:

```
BR0051I BRBACKUP 6.40 (22)
BR0055I Start of database backup: bdsbedhj.aff 2006-02-28 21.16.35

BR0453W Error message found in C:\oracle\SUD\sapbackup\bdsbecnl.aff:

'BR0233E Backup utility has reported an error while saving file
C:\oracle\SUD\sapdata5\datafile.4'
BR0453W Error message found in C:\oracle\SUD\sapbackup\bdsbecnl.aff:
'BR0233E Backup utility has reported an error while saving file
C:\oracle\SUD\sapdata5\datafile.5'

….
….

BR0453W Error message found in C:\oracle\SUD\sapbackup\bdsbecnl.aff:
'BR0279E Return code from 'backint -u SUD -f backup -I
C:\oracle\SUD\sapbackup\.bdsbecnl.lst -t file -p C:\oracle\ora920\dbs\initSUD.utl
-c': 2'

BR0453W Error message found in C:\oracle\SUD\sapbackup\bdsbecnl.aff:

'BR0232E 13 of 29 files saved by backup utility'

BR0453W Error message found in C:\oracle\SUD\sapbackup\bdsbecnl.aff:

'BR0231E Backup utility call failed'

BR0459W Backup C:\oracle\SUD\sapbackup\bdsbecnl.aff terminated with errors

….
….
BR0280I BRBACKUP time stamp: 2006-02-28 21.16.36
BR0057I Backup of database: SUD
BR0058I BRBACKUP action ID: bdsbedhj
BR0059I BRBACKUP function ID: aff
BR0110I Backup mode: ALL
BR0077I Database files for backup:

C:\oracle\SUD\origlogA\log1_m1.dbf
C:\oracle\SUD\origlogB\log2_m1.dbf
```

```
C:\oracle\SUD\origlogA\log3_m1.dbf
C:\oracle\SUD\origlogB\log4_m1.dbf
C:\oracle\SUD\saparch\cntrl\cntrlSUD.dbf

BR0061I 16 files found for backup, total size 283.502 MB
BR0091I 13 files were already saved in: bdsbecnl.aff
```

As the preceding output log shows, 13 files were already backed up. The following is the output of a successful `brbackup` restart job. This job is the one that backed up the remaining 16 files:

```
BR0280I BRBACKUP time stamp: 2006-02-28 21.16.41
BR0229I Calling backup utility with function 'backup'...

BR0278I Command output of 'backint -u SUD -f backup -I
C:\oracle\SUD\sapbackup\.bdsbedhj.lst -t file -p C:\oracle\ora920\dbs\initSUD.utl

-c':
*********************************************************************

Program:                      backint 6.5
Input File:                   C:\oracle\SUD\sapbackup\.bdsbedhj.lst
Profile:                      C:\oracle\ora920\dbs\initSUD.utl

Function:                     BACKUP
Backup Type:                  BACKUP_FILE
*********************************************************************

BR0280I BRBACKUP time stamp: 2006-02-28 21.21.38

#FILE..... C:\oracle\SUD\sapdata5\datafile.4

#SAVED.... VXF1141141607

….
….
BR0280I BRBACKUP time stamp: 2006-02-28 21.21.38

#FILE..... C:\oracle\SUD\saparch\cntrl\cntrlSUD.dbf

#SAVED.... VXF1141141607
*********************************************************************

BR0280I BRBACKUP time stamp: 2006-02-28 21.21.38
BR0232I 16 of 16 files saved by backup utility
BR0230I Backup utility called successfully
```

Restarting SAP database restores on Oracle

You can restore an SAP database that is based on Oracle by using the `brrestore` command. The `brrestore` command writes a record into the input file for each file to be restored. This record is named *BID filename*. `brrestore` then calls the NetBackup for SAP agent (`backint`) with the appropriate options.

Restores can be either successful or unsuccessful.

For each record in the input file, the brrestore command expects one of the following messages from backint:

- For a successful restore:

  ```
  #FILE..... filename
  #RESTORED BID
  ```

- For an unsuccessful restore:

  ```
  #ERROR filename
  ```

  If the restore is successful, `backint` reports `#RESTORED BID filename`. Then, `brrestore` writes the following:
  UNIX or Linux:

  ```
  #RESTORED. VXF1147974254
  ```

  Windows:

  ```
  #RESTORED. VXF1147974254
  ```

  `brrestore` does not write any messages if `backint` reports `#ERROR filename`. If the restore job fails after restoring some files, the NetBackup for SAP agent (`backint`) reports `#RESTORED` for the files it restored. It reports `#ERROR` for the files that it did not restore.

Restarting a restore with the brrestore -f option

You can use the `brrestore` command's `-f` option to restore only the files that failed to be restored. You do not have to specify that all files be restored again.

To restart a restore

- Run the `brrestore` command with the `-f logfile` parameter.

  For `logfile`, specify one of the following:

  - The log file name of the failed job. When the log file is specified, `brrestore` checks it for the files that were not restored successfully.

- The keyword `last`. When this keyword is specified, `brbackup` checks the status of the last restore job for the files that were not restored successfully.

The `brrestore` command examines the specified file. It determines the files that were restored successfully and the files that were not restored successfully. `brrestore` sends `backint` the list of files that still need to be restored. `brrestore` starts a new job only for the files that still need to be restored.

The following command specifies a log file:

```
brrestore -d util_file -m all -f rdsqcxdf.rsb
```

The following command specifies the last restore job:

```
brrestore -d util_file -m all -f last
```

Example: Restarting a NetBackup for SAP restore job (UNIX or Linux)

This section shows how to restart a failed restore job. Assume that the you deliberately cancelled the job from the activity monitor. In reality, the job might fail due to different reasons.

This example shows the log messages that `brrestore` created while restoring the files for the `CER` database, using the following command:

```
brrestore -d util_file -b last -m all
```

The following are excerpts from the corresponding `brrestore` log file:

```
BR0280I BRRESTORE time stamp: 2006-05-18 23.18.15
BR0407I Restore of database: CER
BR0408I BRRESTORE action ID: rdsqcxdf
BR0409I BRRESTORE function ID: rsb
BR0449I Restore mode: ALL
BR0419I Files will be restored from backup: bdsqcwtl.aff 2006-05-18 23.14.01
BR0416I 19 files found to restore, total size 645.148 MB
BR0421I Restore device type: util_file
BR0134I Unattended mode with 'force' active - no operator confirmation allowed

BR0280I BRRESTORE time stamp: 2006-05-18 23.18.16
BR0229I Calling backup utility with function 'restore'...

BR0278I Command output of '/usr/sap/CER/SYS/exe/run/backint -u CER -f restore -I
/oracle/CER/sapbackup/.rdsqcxdf.lst -t file -p /home/orasap/Ora920/dbs/initCER.utl -c':

*******************************************************************************
```

```
Program: /usr/sap/CER/SYS/exe/run/backint 6.5
Input File: /oracle/CER/sapbackup/.rdsqcxdf.lst
Profile: /home/orasap/Ora920/dbs/initCER.utl

Function: RESTORE
*******************************************************************************
BR0280I BRRESTORE time stamp: 2006-05-18 23.18.57

#FILE..... /oracle/CER/sapdata1/btabd_1/btabd.data1

#RESTORED. VXF1147974254

….
….
BR0280I BRRESTORE time stamp: 2006-05-18 23.18.57

#FILE..... /oracle/CER/sapdata1/el46bi_1/el46bi.data1

#RESTORED. VXF1147974254
*******************************************************************************

BR0374E 7 of 19 files restored by backup utility
BR0280I BRRESTORE time stamp: 2006-05-18 23.18.57
BR0231E Backup utility call failed
```

After you start the `brrestore` command, you cancel the restore job from the activity monitor. As the preceding log indicates, `brrestore` restored 7 out of the 19 files. You can restore the remaining files as part of a new `brrestore` session. In a production situation, you should identify and correct the cause of failure before restarting `brrestore`.

Enter the following command with the `-f` option to start `brrestore`:

```
brrestore -d util_file -m all -f last
```

The following output is an excerpt from the corresponding `brrestore` log file:

```
BR0453W Error message found in /oracle/CER/sapbackup/rdsqcxdf.rsb:

'BR0374E 7 of 19 files restored by backup utility'

BR0453W Error message found in /oracle/CER/sapbackup/rdsqcxdf.rsb:

'BR0231E Backup utility call failed'

BR0471W Restore /oracle/CER/sapbackup/rdsqcxdf.rsb terminated with errors

BR0428W File /oracle/CER/sapdata6/es46bd_1/es46bd.data1 will be overwritten
```

```
BR0280I BRRESTORE time stamp: 2006-05-18 23.19.20
BR0407I Restore of database: CER
BR0408I BRRESTORE action ID: rdsqcxfs
BR0409I BRRESTORE function ID: rsb
BR0449I Restore mode: ALL
BR0419I Files will be restored from backup: bdsqcwtl.aff 2006-05-18 23.14.01
BR0416I 12 files found to restore, total size 520.094 MB
BR0445I 7 files were already restored in the following run: rdsqcxdf.rsb
BR0421I Restore device type: util_file
```

`brrestore` checks the previous log file and finds that 12 files need to be restored and that 7 files are already restored. `brrestore` directs `backint` to restore the 12 files. The log file is as follows:

```
BR0134I Unattended mode with 'force' active - no operator confirmation allowed

BR0280I BRRESTORE time stamp: 2006-05-18 23.19.20
BR0229I Calling backup utility with function 'restore'...

BR0278I Command output of '/usr/sap/CER/SYS/exe/run/backint -u CER -f restore -I /oracle/
CER/sapbackup/.rdsqcxfs.lst -t file -p /home/orasap/Ora920/dbs/initCER.utl -c':

*******************************************************************

Program:/usr/sap/CER/SYS/exe/run/backint 6.5
Input File:/oracle/CER/sapbackup/.rdsqcxfs.lst
Profile:/home/orasap/Ora920/dbs/initCER.utl

Function:RESTORE
*******************************************************************
BR0280I BRRESTORE time stamp: 2006-05-18 23.20.09

#FILE..... /oracle/CER/sapdata6/es46bd_1/es46bd.data1

#RESTORED. VXF1147974254

….
….
BR0280I BRRESTORE time stamp: 2006-05-18 23.20.09

#FILE..... /oracle/CER/sapdata1/system_1/system.data1

#RESTORED. VXF1147974254
*******************************************************************

BR0280I BRRESTORE time stamp: 2006-05-18 23.20.09
```

```
BR0374I 12 of 12 files restored by backup utility
BR0230I Backup utility called successfully
```

As the preceding log shows, the remaining 12 files are restored successfully.

Example: Restarting a NetBackup for SAP restore job (Windows)

This section shows how to restart a failed restore job. Assume that the you deliberately cancelled the job from the activity monitor. In reality, the job might fail due to different reasons.

This example shows the log messages that `brrestore` created while restoring the files for the `CER` database, using the following command:

```
brrestore -d util_file -b last -m all
```

The following are excerpts from the corresponding `brrestore` log file:

```
BR0280I BRRESTORE time stamp: 2006-05-18 23.18.15
BR0407I Restore of database: CER
BR0408I BRRESTORE action ID: rdsqcxdf
BR0409I BRRESTORE function ID: rsb
BR0449I Restore mode: ALL
BR0419I Files will be restored from backup: bdsqcwtl.aff 2006-05-18 23.14.01
BR0416I 19 files found to restore, total size 645.148 MB
BR0421I Restore device type: util_file
BR0134I Unattended mode with 'force' active - no operator confirmation allowed

BR0280I BRRESTORE time stamp: 2006-05-18 23.18.16
BR0229I Calling backup utility with function 'restore'...

BR0278I Command output of 'backint -u CER -f restore -I
C:\oracle\SUD\sapbackup\.rdsqcxdf.lst -t file -p C:\oracle\Ora920\dbs\initCER.utl
-c':

*******************************************************************************

Program: backint 6.5
Input File: C:\oracle\SUD\sapbackup\.rdsqcxdf.lst
Profile: C:\oracle\Ora920\dbs\initCER.utl

Function: RESTORE
*******************************************************************************
BR0280I BRRESTORE time stamp: 2006-05-18 23.18.57

#FILE..... C:\oracle\SUD\sapdata1\btabd_1\btabd.data1
```

```
#RESTORED. VXF1147974254

….
….
BR0280I BRRESTORE time stamp: 2006-05-18 23.18.57

#FILE..... C:\oracle\SUD\sapdata1\el46bi_1\el46bi.data1

#RESTORED. VXF1147974254
**********************************************************************

BR0374E 7 of 19 files restored by backup utility
BR0280I BRRESTORE time stamp: 2006-05-18 23.18.57
BR0231E Backup utility call failed
```

After you start the `brrestore` command, you cancel the restore job from the activity monitor. As the preceding log indicates, `brrestore` restored 7 out of the 19 files. You can restore the remaining files as part of a new `brrestore` session. In a production situation, you should identify and correct the cause of failure before restarting `brrestore`.

Enter the following command with the `-f` option to start `brrestore`:

```
brrestore -d util_file -m all -f last
```

The following output is an excerpt from the corresponding `brrestore` log file:

```
BR0453W Error message found in C:\oracle\SUD\sapbackup\rdsqcxdf.rsb:

'BR0374E 7 of 19 files restored by backup utility'

BR0453W Error message found in C:\oracle\SUD\sapbackup\rdsqcxdf.rsb:

'BR0231E Backup utility call failed'

BR0471W Restore C:\oracle\SUD\sapbackup\rdsqcxdf.rsb terminated with errors

BR0428W File C:\oracle\SUD\sapdata6\es46bd_1\es46bd.data1 will be overwritten

BR0280I BRRESTORE time stamp: 2006-05-18 23.19.20
BR0407I Restore of database: CER
BR0408I BRRESTORE action ID: rdsqcxfs
BR0409I BRRESTORE function ID: rsb
BR0449I Restore mode: ALL
BR0419I Files will be restored from backup: bdsqcwtl.aff 2006-05-18 23.14.01
BR0416I 12 files found to restore, total size 520.094 MB
```

```
BR0445I 7 files were already restored in the following run: rdsqcxdf.rsb
BR0421I Restore device type: util_file
```

`brrestore` checks the previous log file and finds that 12 files need to be restored and that 7 files are already restored. `brrestore` directs `backint` to restore the 12 files. The log file is as follows:

```
BR0134I Unattended mode with 'force' active - no operator confirmation allowed

BR0280I BRRESTORE time stamp: 2006-05-18 23.19.20
BR0229I Calling backup utility with function 'restore'...

BR0278I Command output of '/usr/sap/CER/SYS/exe/run/backint -u CER -f restore -I
C:\oracle\SUD\sapbackup\.rdsqcxfs.lst -t file -p C:\oracle\Ora920\dbs\initCER.utl -c':

*********************************************************************

Program:/usr/sap/CER/SYS/exe/run/backint 6.5
Input File:C:\oracle\SUD\\sapbackup\.rdsqcxfs.lst
Profile:C:\oracle\Ora920\dbs\initCER.utl

Function:RESTORE
*********************************************************************
BR0280I BRRESTORE time stamp: 2006-05-18 23.20.09

#FILE..... C:\oracle\SUD\sapdata6\es46bd_1\es46bd.data1

#RESTORED. VXF1147974254

….
….
BR0280I BRRESTORE time stamp: 2006-05-18 23.20.09

#FILE..... C:\oracle\SUD\sapdata1\pooli_1\pooli.data1

#RESTORED. VXF1147974254
*********************************************************************

BR0280I BRRESTORE time stamp: 2006-05-18 23.20.09
BR0374I 12 of 12 files restored by backup utility
BR0230I Backup utility called successfully
```

As the preceding log shows, the remaining 12 files are restored successfully.

Chapter 11

Performing backups and restores of SAP HANA

This chapter includes the following topics:

About backups and restores of SAP HANA

You can run backups using SQL commands. You can also use Cron to schedule the backups. Check the SAP website for the SAP HANA Administration Guide.

Backups for SAP HANA can be initiated from NetBackup. SAP HANA data restore is initiated from the SAP HANA studio and not from NetBackup web UI. For more information, see the following tech note on the Cohesity Support website for full backup:

https://support.cohesity.com/s/article/article-100010190

See the following tech note on the Cohesity Support website for full, differential, and incremental backup:

https://support.cohesity.com/s/article/article-100015942

Using the SAP HANA studio for backup

Use the SAP HANA studio to backup the SAP HANA data. The following steps help you to set up the backup.

To set up the backup using the SAP HANA studio

1. On the SAP HANA studio interface, select the HANA instance that you want to backup. In this case <SID>.
2. Right-click on the instance and select the **Back Up**option as shown.

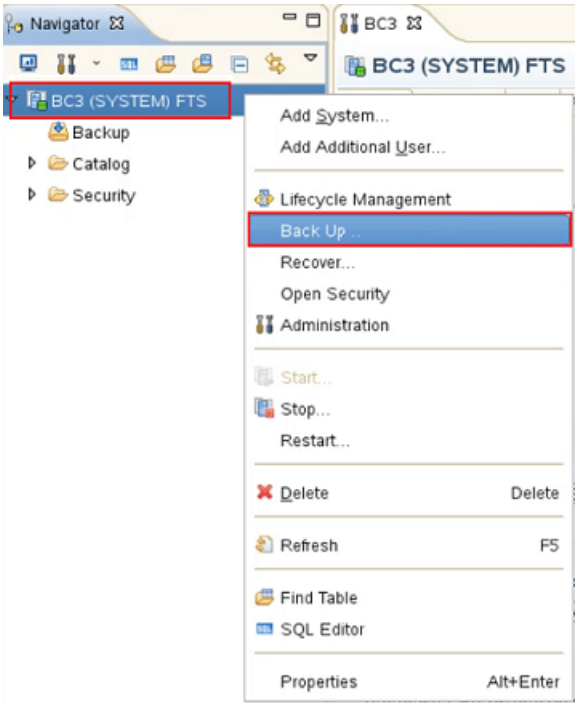

3. The Backup of System screen is displayed. Select the **Backint** option. Specify the **Backup Destination** and **Backup Prefix** as shown and click **Next**.

Note: If you use SAP HANA SPS 10, you can also select **Differential Data Backup** or **Incremental Data Backup** as the **Backup Type**. Enter a **Backup Prefix** that indicates if the backup type is incremental or differential. For example, for an incremental data backup, the backup prefix may be *INCRE_DATA_BACKUP*.

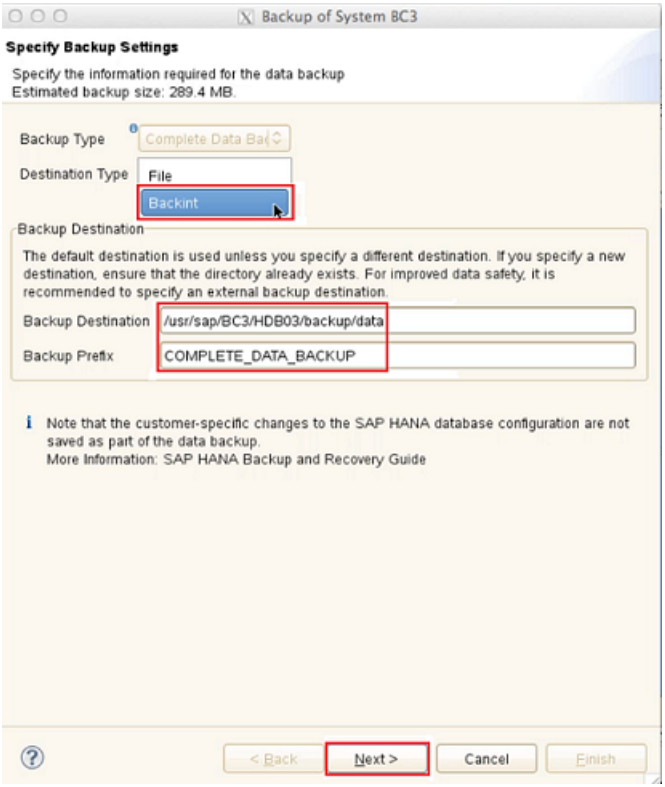

4 Confirm the settings under **Review Backup Settings** as shown and click **Finish**.

5 Open the Activity monitor to see the backup details.

For a successful backup, the following results are displayed in the SAP HANA interface:

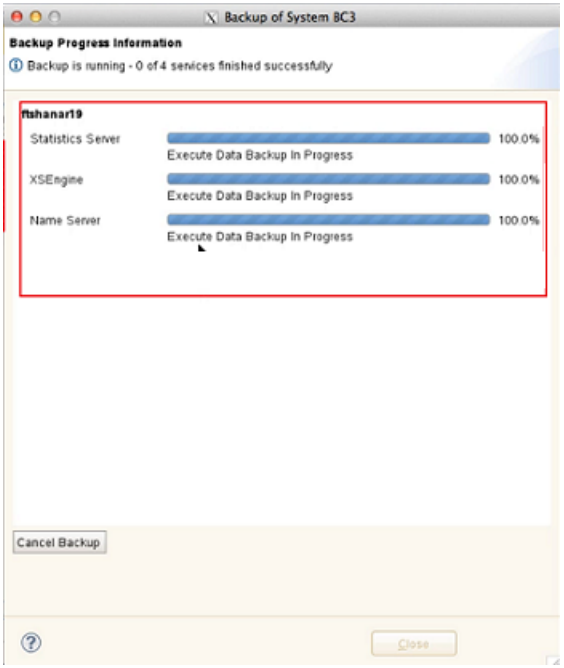

6 On the Backup of System screen, click **Close**.

Using the SAP HANA studio for restore

Use the SAP HANA studio to backup the SAP HANA data. The following steps help you to set up the restore.

Setting up the restore using the SAP HANA studio

1 On the SAP HANA studio interface, select the HANA instance that you want to recover. In this case <SID>.

2 Right-click on the instance and select the **Recover** option as shown.

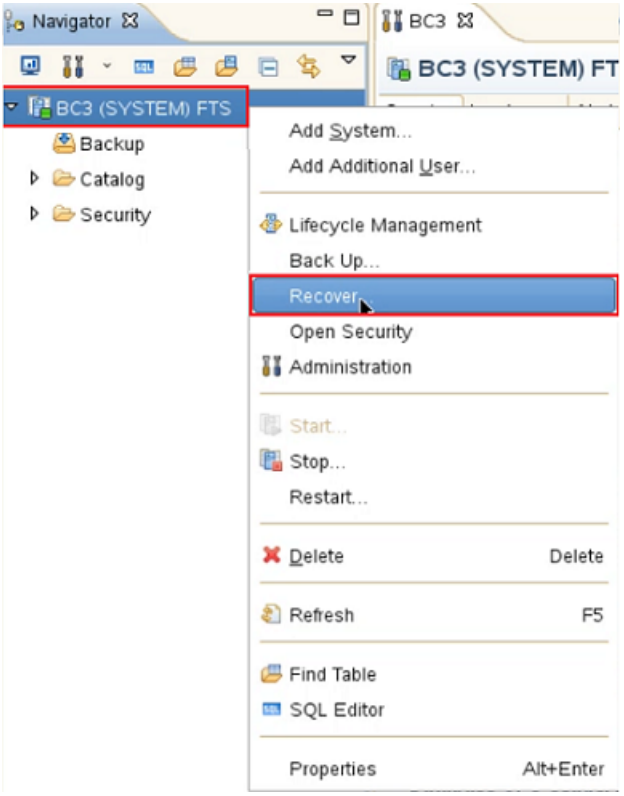

3 For a recovery, the system must be offline. Select **OK** on the System Shutdown screen as shown.

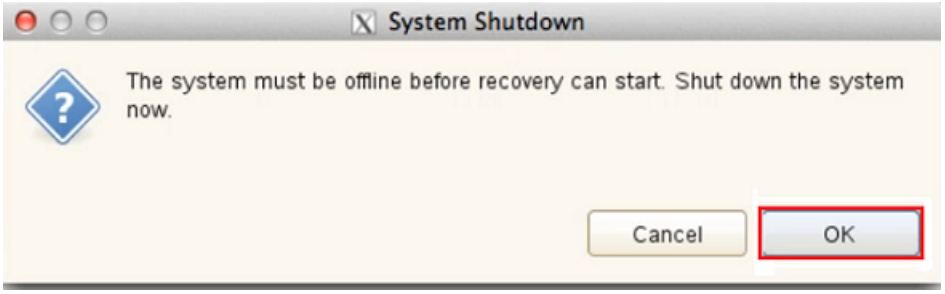

4 The **Logon to SAP System Host screen** displays. Specify the username and password as shown and click **OK**.

5 In the **Recovery of System** screen, select the **Recover the database to its most recent state** option and click **Next**.

6 Check the details under **Recovery of the Log Backups** and click **Next**.

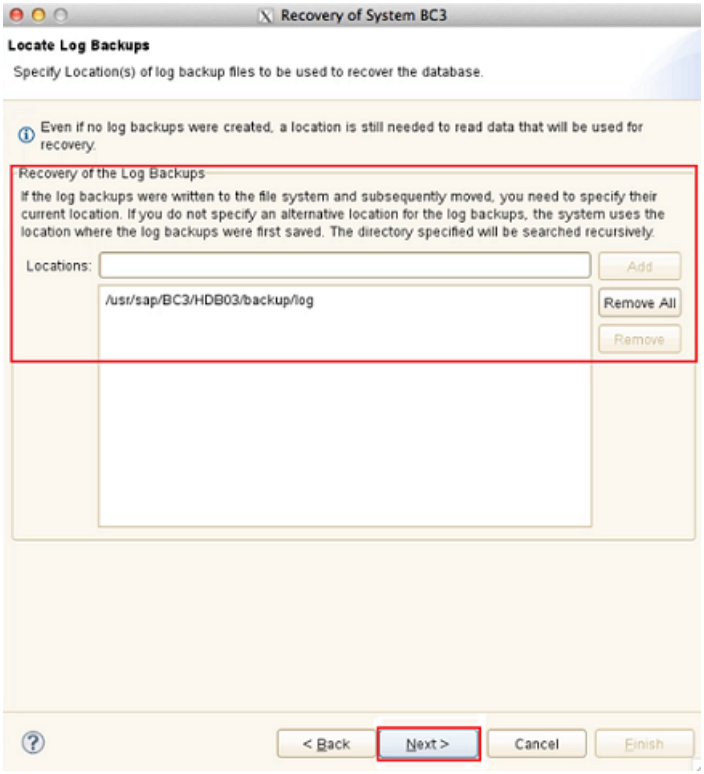

7 Under **Select Data Backup**, select the data backup to recover the SAP HANA database and click **Next**.

8 In the Other Settings screen, in the **Check Availability of Delta and Log Backups** section, select the **Third-Party Backup Tool (Backint)** check box.

Note: If you use SAP HANA SPS 10 or later, then ensure that in the **Use Delta Backups** section, the **Use Delta Backups** check box is selected.

---

9 Confirm the recovery details under **Review Recovery Settings** and click **Finish**.

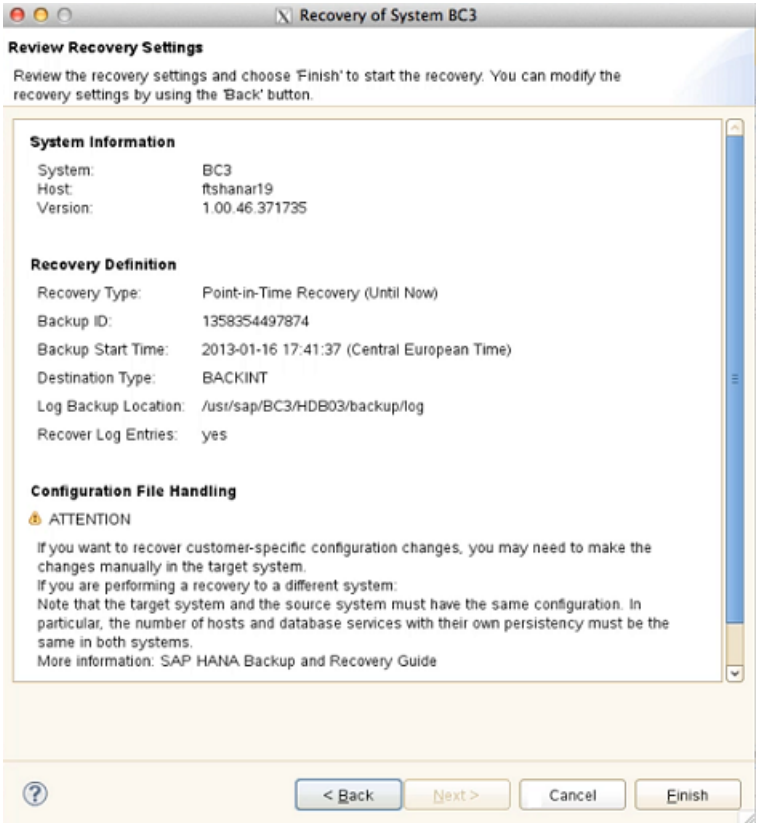

10 Open the Activity monitor to see the restore details.

For a successful restore, the following results are displayed in the SAP HANA interface:

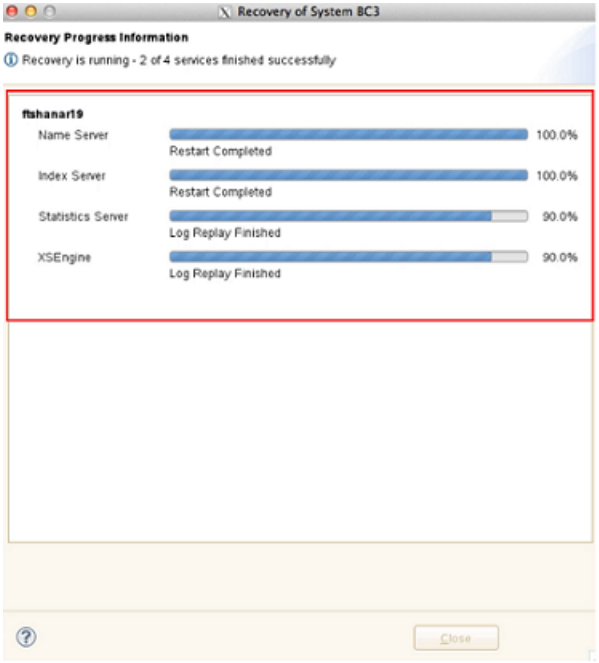

11 After the recovery, on the **Recovery of System** screen, click **Close**.

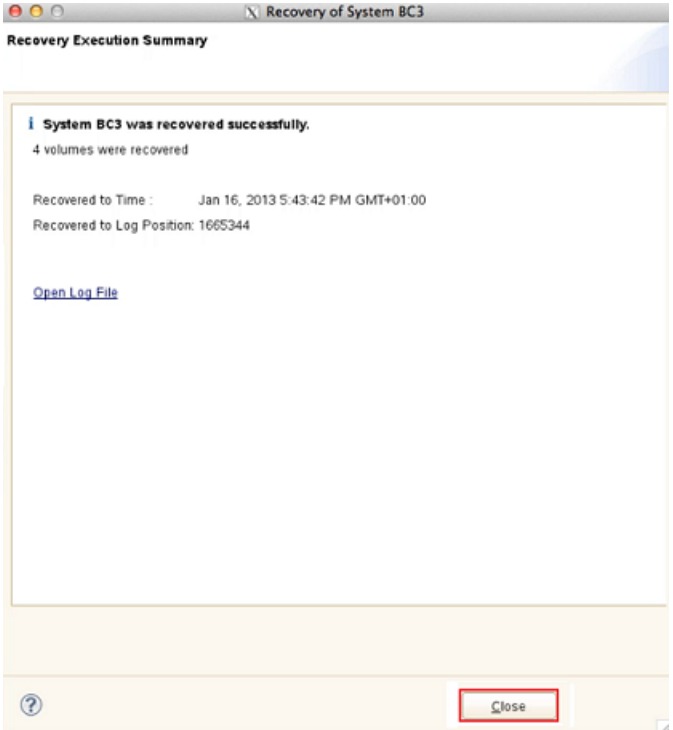

Using SAP HANA studio for SAP HANA instance redirected restore

Use the SAP HANA studio to backup the SAP HANA data. The following steps help you to set up the redirected restore.

To set up the SAP HANA instance redirected restore using the SAP HANA studio

1. On the SAP HANA studio interface, select the destination SAP HANA instance that you want to recover.
2. Right-click the SAP HANA instance, and select **Backup and Recovery** > **Recover System**.
3. On the **Shutdown System** screen, select **OK**.

   Before a recovery can start, the system must be offline.

4 In the **Recovery of System <*system name*>** screen, select one of the following options, and click **Next**.

- **Recover the database to its most recent state** - If you select this option, see step 5 for further steps.
- **Recover the database to the following point in time** - If you select this option, see step 5 for further steps.
- **Recover the database to a specific data backup or storage snapshot** – Select this option when you want to restore only the data backup and not the log backups. If you select this option, see step 10 for further steps.

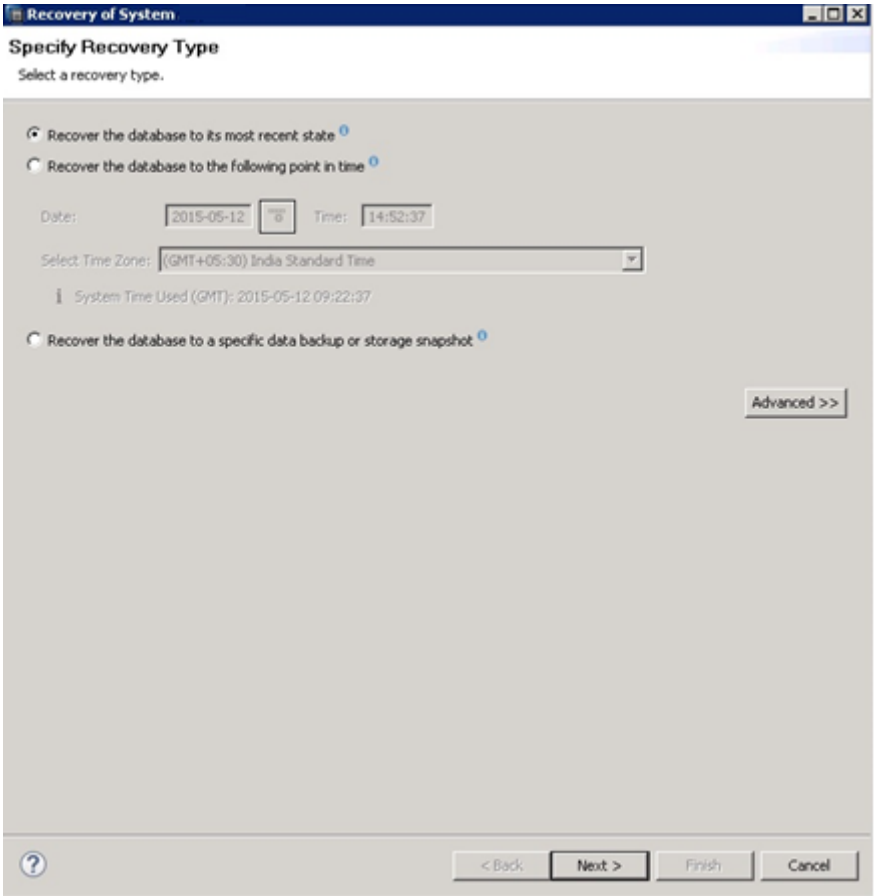

5 Specify the source instance name in the **Source System** field, and click **Next**.

Note: By default, the **Source System** field is prepopulated with the destination SAP HANA instance name.

In SAP HANA SPS 10 and later, you must select the **Backint System Copy** check box to enable the **Source System** field.

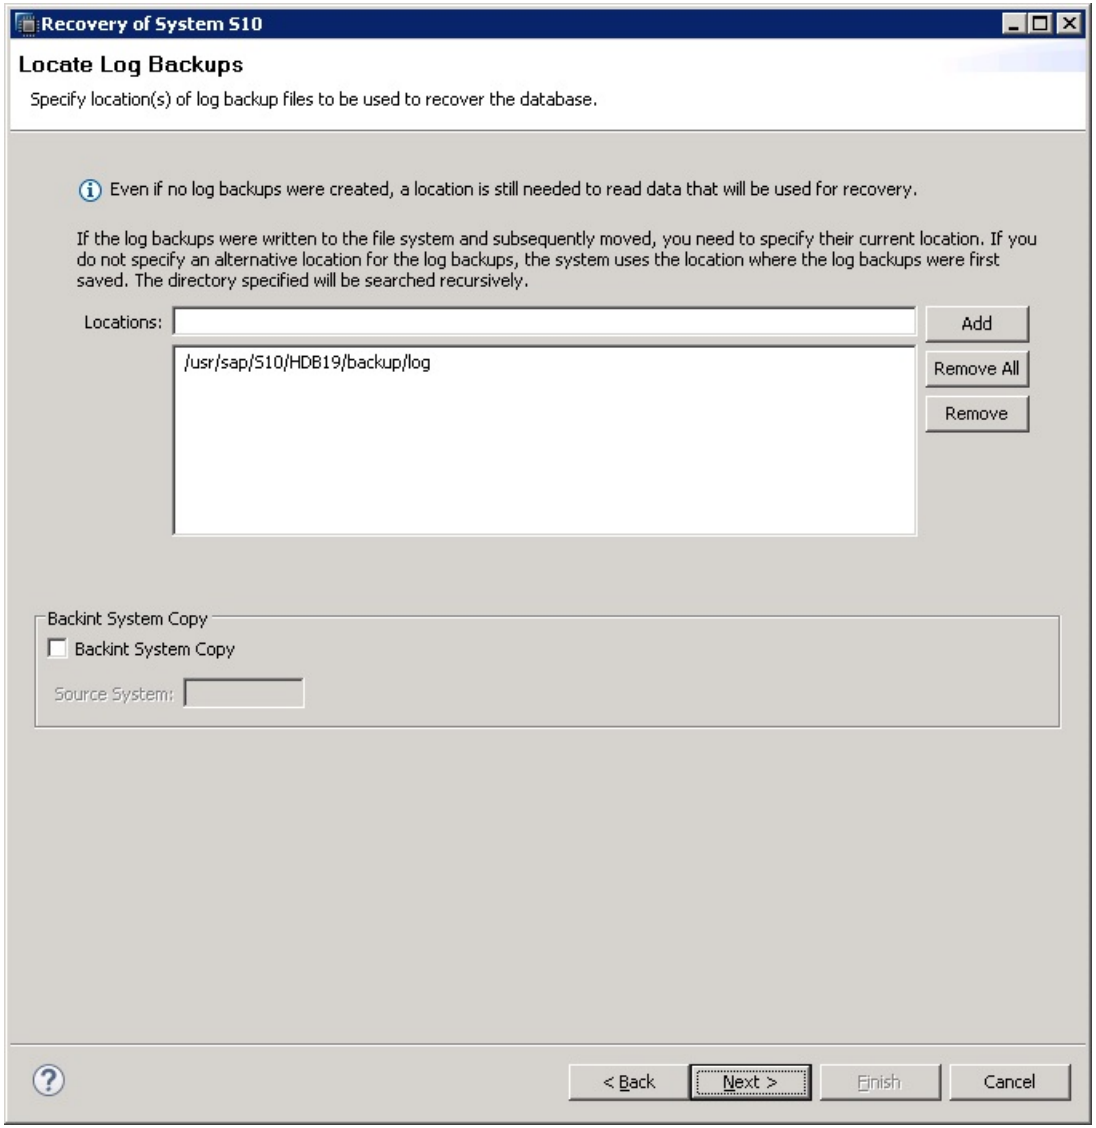
Recovery of System S10
Locate Log Backups
Specify location(s) of log backup files to be used to recover the database.
Even if no log backups were created, a location is still needed to read data that will be used for recovery.
If the log backups were written to the file system and subsequently moved, you need to specify their current location. If you do not specify an alternative location for the log backups, the system uses the location where the log backups were first saved. The directory specified will be searched recursively.
Locations:
Add
/usr/sap/S10/HDB19/backup/log
Remove All
Remove
Backint System Copy
Backint System Copy
Source System:
< Back
Next >
Finish
Cancel

6 In the **Select a Backup** screen, select a data backup to recover the SAP HANA database, and click **Next**.

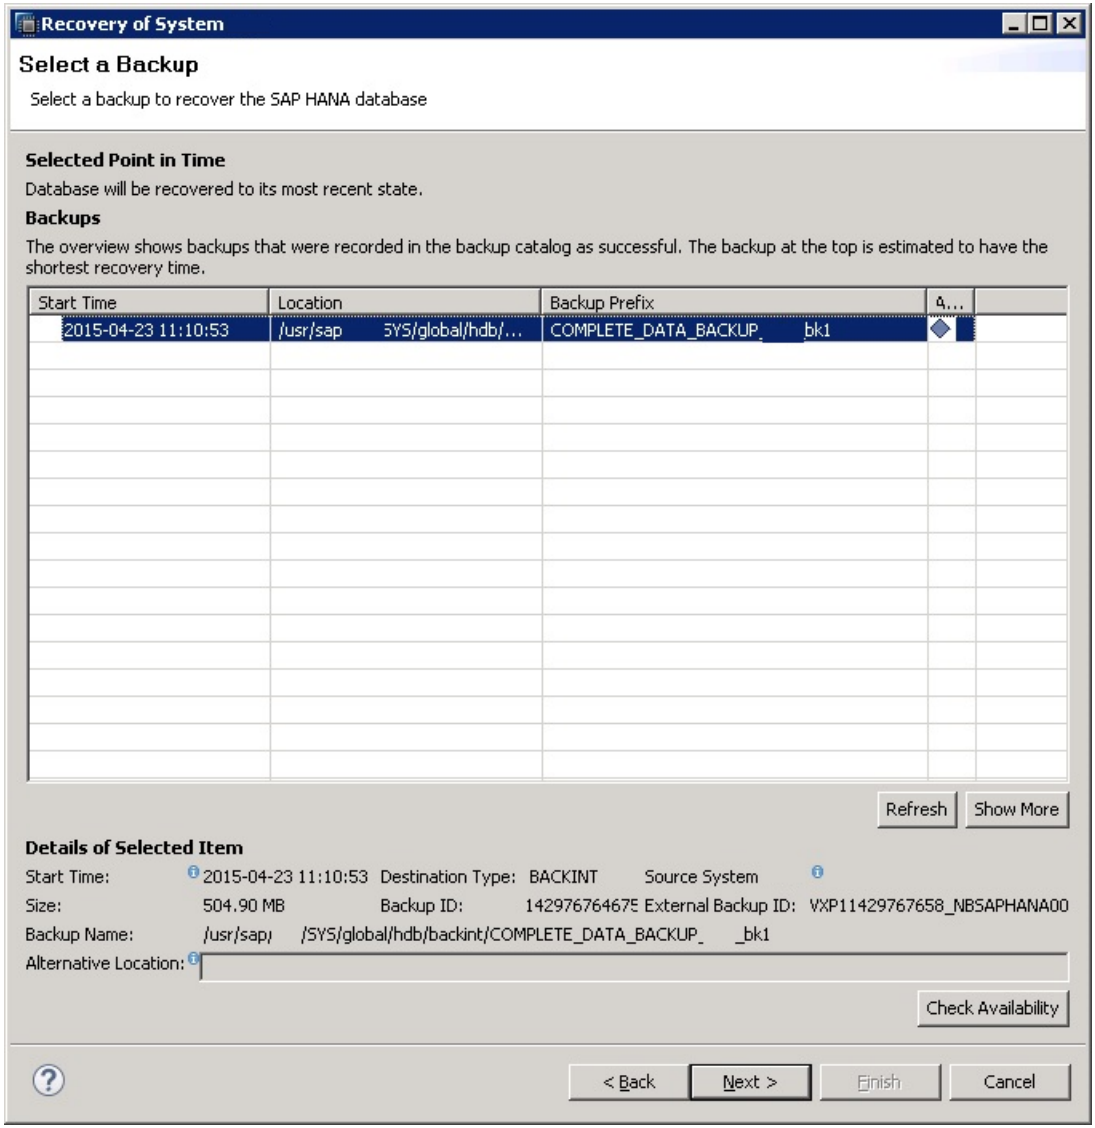

7 In the **Check Availability of Log Backups** section, clear the **File System** check box and select the **Third-Party Backup Tool (Backint)** check box. In the **Initialize Log Area** section, select the **Initialize Log Area** check box, and then click **Next**.

Note: If you use SAP HANA SPS 10 and later, ensure that the **Use Delta Backups** check box is selected.

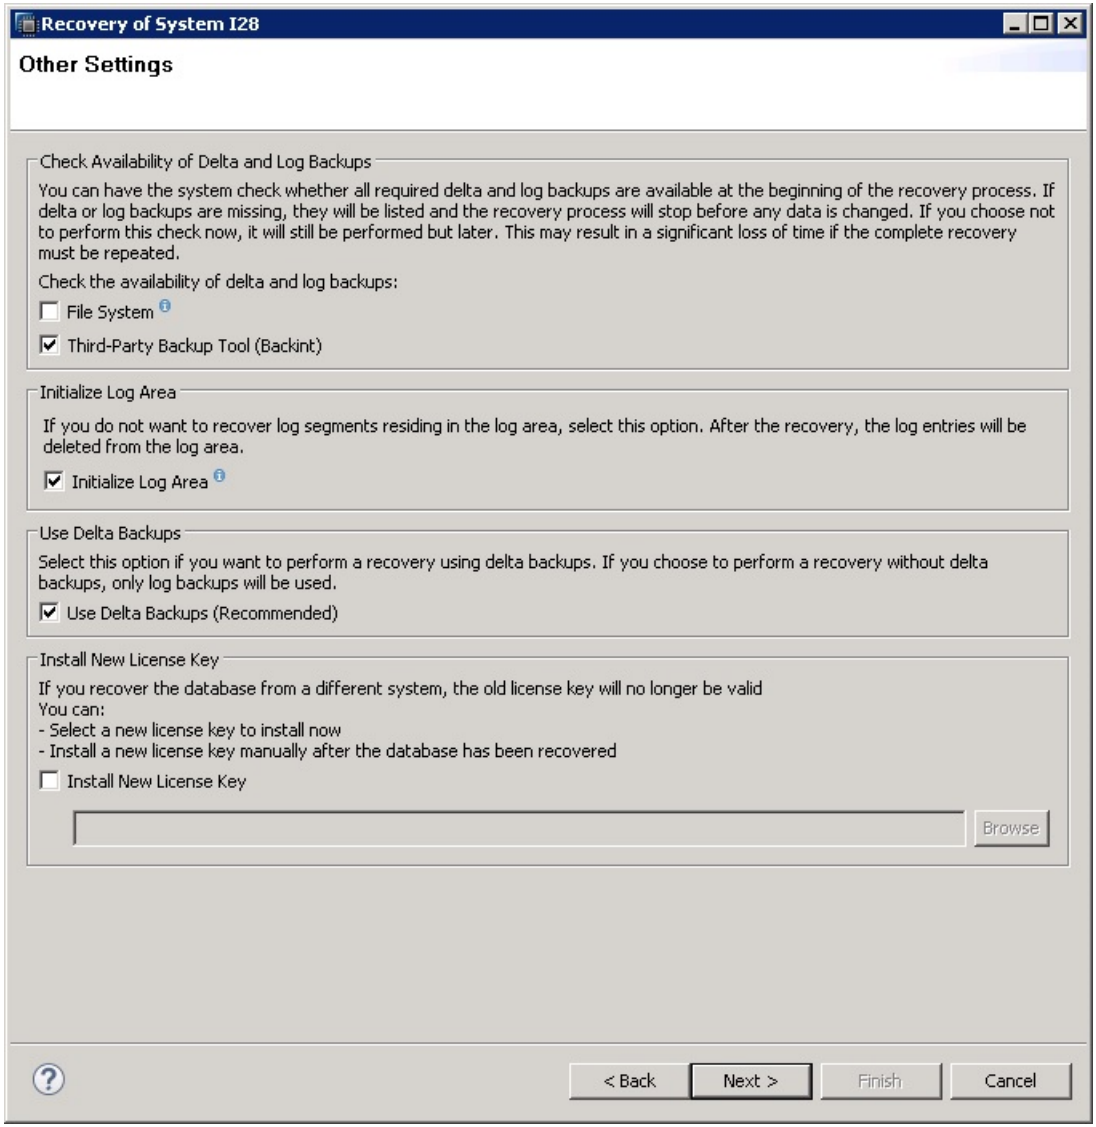
Recovery of System I28
Other Settings
Check Availability of Delta and Log Backups
You can have the system check whether all required delta and log backups are available at the beginning of the recovery process. If delta or log backups are missing, they will be listed and the recovery process will stop before any data is changed. If you choose not to perform this check now, it will still be performed but later. This may result in a significant loss of time if the complete recovery must be repeated.
Check the availability of delta and log backups:
File System
Third-Party Backup Tool (Backint)
Initialize Log Area
If you do not want to recover log segments residing in the log area, select this option. After the recovery, the log entries will be deleted from the log area.
Initialize Log Area
Use Delta Backups
Select this option if you want to perform a recovery using delta backups. If you choose to perform a recovery without delta backups, only log backups will be used.
Use Delta Backups (Recommended)
Install New License Key
If you recover the database from a different system, the old license key will no longer be valid
You can:
- Select a new license key to install now
- Install a new license key manually after the database has been recovered
Install New License Key
Browse
< Back
Next >
Finish
Cancel

8 In the **Review Recovery Settings** screen, confirm the recovery details, and click **Finish**.

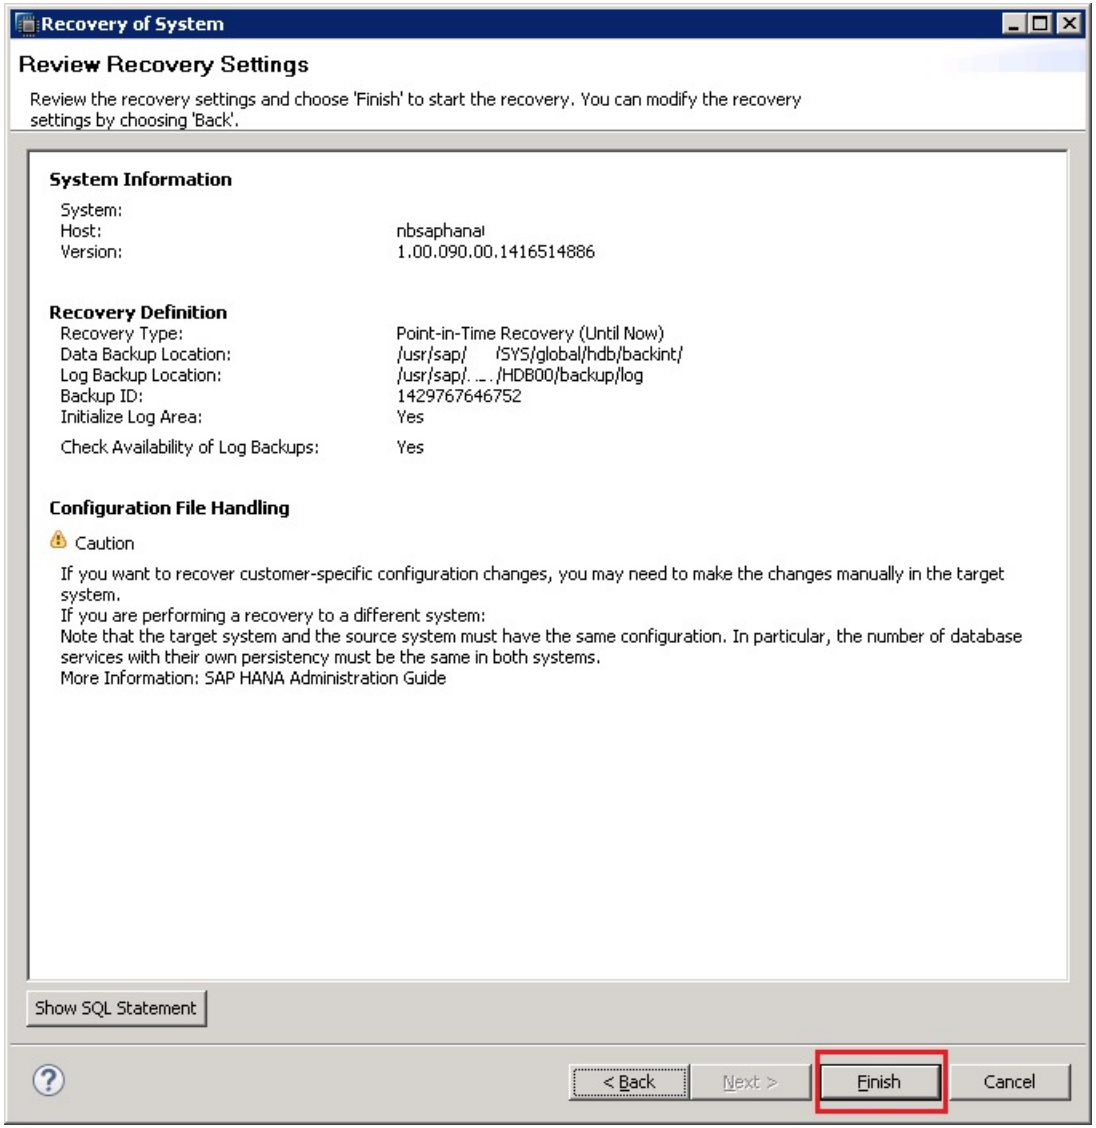

9 Open the Activity monitor to view the recovery information.

10 If you have selected **Recover the database to a specific data backup or storage snapshot** in step 3 and clicked **Next**, then in the **Specify Backup Location** screen, you can select one of the following options:

- **Search for the catalog in Backint only** – see Step 11 for further steps.
- **Specify backup without catalog** – see Step 16 for further steps.

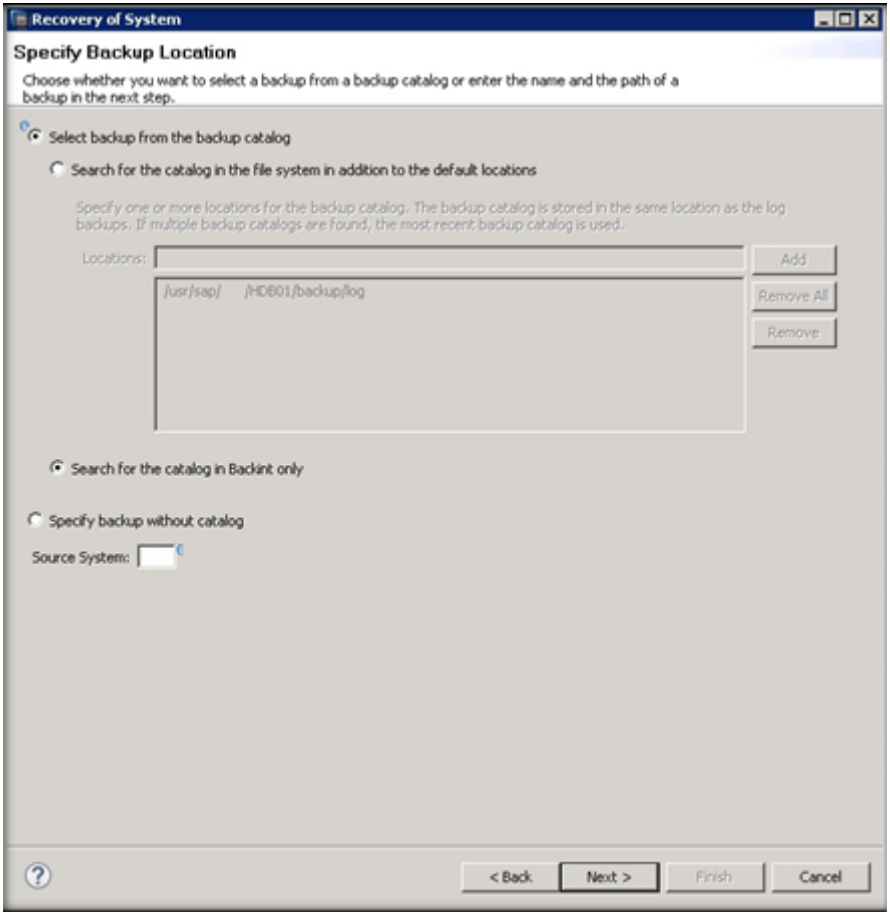

11 Specify the source instance name in the **Source System** field, and click **Next**.

Note: By default, the **Source System** field is prepopulated with the destination SAP HANA instance name.

If you use SAP HANA SPS 10 and later, you must select the **Backint System Copy** check box to enable the **Source System** field.

---

12 In the `Select a Backup` screen, select a data backup to recover the SAP HANA database, and click **Next**.

13 In the **Other Settings** screen, the **Initialize Log Area** check box is selected by default. Click **Next**.

Note: If you use SAP HANA SPS 10 and later, ensure that the **Use Delta Backups** check box is selected.

---

14 In the **Review Recovery Settings** screen, confirm the recovery details, and click **Finish**.

15 Check the Activity monitor to view the recovery information.

16 Specify the source instance name in the **Source System** field, and click Next.

Note: By default, the **Source System** field is prepopulated with the destination SAP HANA instance name.

If you use SAP HANA SPS 10 and later, you must select the **Backint System Copy** check box to enable the **Source System** field.

---

17 In the **Specify the Backup to Recover** screen, specify the **Destination Type** as **Backint** and in the **Backup Prefix** field, specify the backup prefix. Click **Next**.

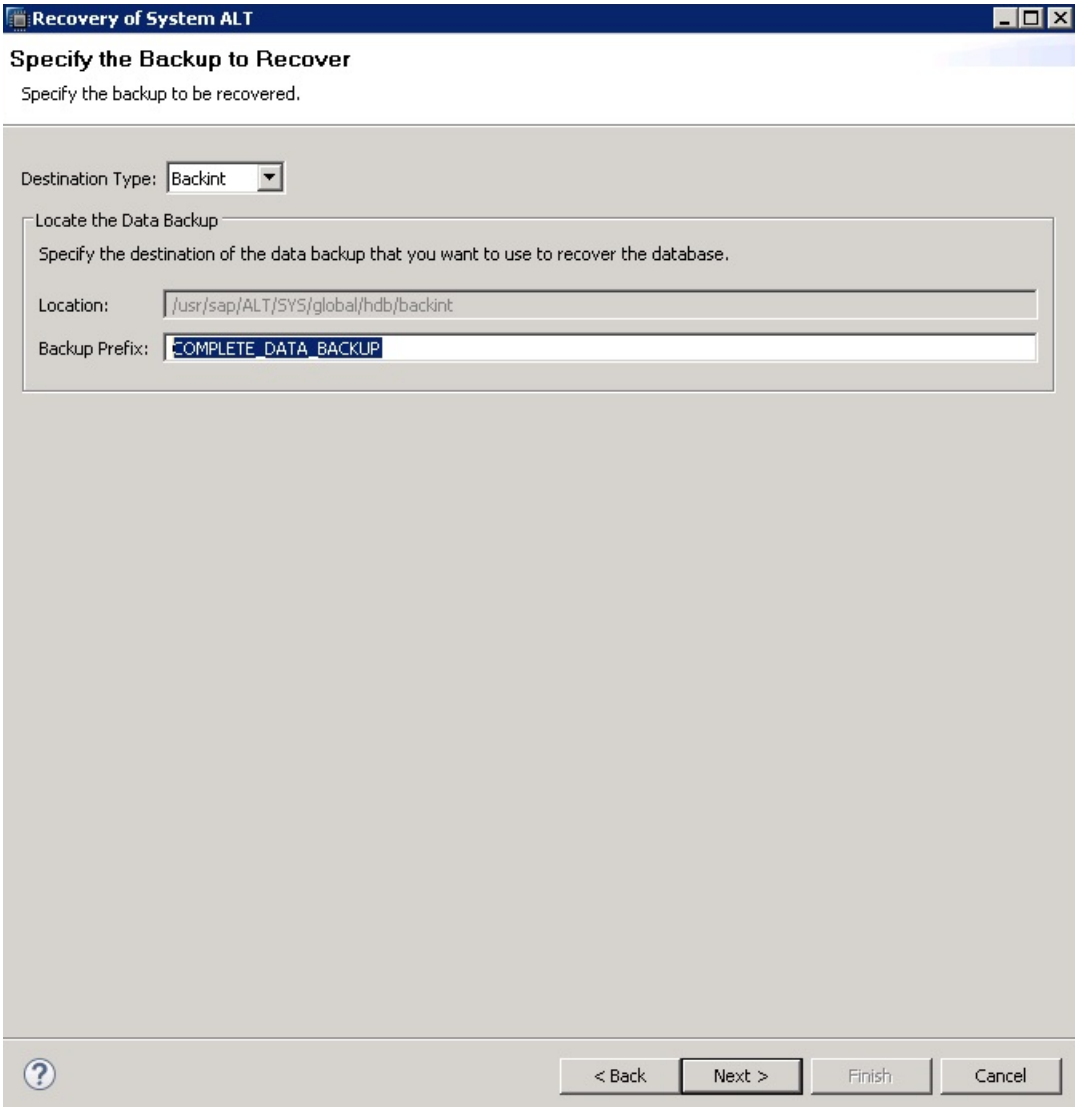

18 In the **Other Settings** screen, the **Initialize Log Area** check box is selected by default. Click **Next**.

Note: If you use SAP HANA SPS 10 and later, ensure that the **Use Delta Backups** check box is selected.

19 In the **Review Recovery Settings** screen, confirm the recovery details, and click **Finish**.

20 Open the Activity monitor in NetBackup to view the recovery information.

Configuring multistreaming with third-party backup tools

By default, multistreaming is disabled, which means that SAP HANA uses a single channel for data backups and delta backups. If required, you can configure SAP HANA to use additional channels. When multiple channels are used, SAP HANA distributes the data equally across the available channels. All the parts of a multistreamed backup are approximately the same size.

Multistreamed data backups are configured individually for each database.

Note: To create multistreamed data backups and delta backups of an SAP HANA database, the third-party backup tool must also be configured to use multiple channels. To get optimal performance, ensure that the third-party backup tool is configured correctly.

For more information about the configuration of the backup tool, consult the vendor documentation.

To enable multistreaming

1 To enable multistreaming for data backups and delta backups, configure the parameters `parallel_data_backup_backint_channels` and `parallel_data_backup_backint_size_threshold` for each database. The parameters are in the backup section of `global.ini`.

2 Specify the appropriate values for each service.

- `parallel_data_backup_backint_channels` - Specify the number of channels to be used for multistreaming.
  You can specify a value between 1 and 32. The maximum number of channels permitted for each service is 32.

The default value is 1.

A value of 1 means that data backups with third-party backup tools are created through a single channel. This disables multistreaming. To enable multistreaming, specify a value greater than 1.

Note: Each additional channel requires an additional IO buffer (by default, 512 MB). Ensure that increasing the number of channels does not have a negative impact on memory consumption.

If backup encryption is enabled, three additional IO buffers are needed for each channel.

- `parallel_data_backup_backint_size_threshold` - Specify the minimum data backup size for multistreaming to be enabled. The value is specified in GB.

  You can specify a value of 1 GB or more.

  The default value is 128 GB.

  The threshold applies to each service individually. If the data backup size for a service is below the threshold, that service will not be multistreamed, even if multistreaming is enabled (parallel_data_backup_backint_channels > 1).

3 (Optional) Ensure that all buffers fit into memory. You can optimize the configuration of the following parameters:

- Reduce the backup buffer size.
  Reconfigure data_backup_buffer_size.
- Reduce the number of channels.
  Reconfigure parallel_data_backup_backint_channels.

4 Save.

The changes take effect immediately.

Chapter 12

NetBackup for SAP with Snapshot Client

This chapter includes the following topics:

NetBackup for SAP with Snapshot Client overview

The NetBackup for SAP Snapshot Client software consolidates a variety of snapshot-based technologies into a single, easy-to-use backup solution. When

NetBackup for SAP is used with Snapshot Client, the environments that are based on Oracle are supported.

Note: Environments that are based on MaxDB databases are not supported.

The ability to restore your environment quickly depends on your ability to back up business-critical data quickly. Backups enable you to restore your environment in the event of logical database errors or physical errors such as hardware failures or disasters. Snapshot backups enable you to minimize the length of time that a database is in backup mode or offline. Cohesity recommends that customers perform online backups at least daily and offline backups once a week. Since periods of low SAP system activity are few and far between because customers tend to run long batch jobs on weekends or during night hours, snapshots make it possible to protect the data with minimal delay.

Table 12-1 describes the Snapshot Client features.

Table 12-1 NetBackup for SAP with Snapshot Client features

Snapshot Client feature	Description
Array and software snapshot integration	Supports a variety of array and software snapshots and provides a base for all Snapshot Client solutions.
Snapshot backups	A snapshot backup occurs when NetBackup creates a point-in-time copy of the disk volumes on which the database resides. When snapshot backup is used with Snapshot Client, NetBackup for SAP backs up Oracle objects by taking snapshot images of the component files on the local host. Later it backs up the snapshot to a storage unit. The snapshot process is nearly instantaneous; so user access to the database is not interrupted. Client operations and user access continue without interruption during the backup. The backup does not affect the performance or availability of the database. You can perform snapshot backup by using the `backint` command or by using RMAN proxy copy.

Table 12-1 NetBackup for SAP with Snapshot Client features *(continued)*

Snapshot Client feature	Description
Instant recovery backups	This feature makes backups available for instant recovery from disk. Instant recovery combines snapshot technology with the ability to do rapid disk-based restores. NetBackup creates the image without interrupting user access to data. The snapshot can optionally be backed up to tape or other storage long-term storage, but is retained for recovery in the short term. Instant recovery makes it possible to perform file promotion and rollback without accessing the long-term storage. Additionally for UNIX and Linux, instant recovery makes it possible to perform block-level restores.
Off-host backups	An off-host backup shifts the burden of the backup process onto a separate backup agent, such as an alternate client or data mover. An off-host backup reduces the effect on the client's resources that a local backup ordinarily causes. The backup agent reads the data from the client disk and writes it to storage. You can perform off-host backups by using the`backint` command or by using RMAN proxy.
Block-level incremental backup and recovery	This feature is available for UNIX and Linux. Less data leads to high performance data protection. A Block-Level Incremental (BLI) Backup uses the change tracking capabilities of the Veritas File System (VxFS) Storage Checkpoint feature. In a BLI backup, only the changed blocks of data are backed up, not the entire file or file system. A BLI backup saves time, decreases the amount of backup media that is required, and significantly reduces CPU and network overhead during backups.
RMAN proxy copy	Proxy copy is an extension to Oracle's media management API. A proxy copy is a special type of backup in which RMAN turns over control of the data transfer to the NetBackup for SAP Agent. The agent can then manage the entire data movement between the disks that contain the Oracle data files and the storage devices that NetBackup manages. With proxy copy, RMAN provides a list of files that require backup or restore to the NetBackup for SAP Agent. The NetBackup for SAP Agent determines how the data is moved.

Table 12-1 NetBackup for SAP with Snapshot Client features *(continued)*

Snapshot Client feature	Description
RMAN stream-based backups	RMAN stream-based operations are the standard way by which NetBackup for SAP implements conventional RMAN backups and restores. In a stream-based backup, NetBackup moves the data that the Oracle server process provides. NetBackup for SAP captures the data stream content that RMAN provides and stores it to media. If the user configures SAP for RMAN with multiple channels, the following occurs: ■ RMAN makes multiple backup requests to NetBackup in parallel ■ NetBackup backs up and catalogs the multiple backups as separate images
RMAN proxy and RMAN stream-based backups intermixed	RMAN Proxy backups can be used to perform off-host snapshot-based split mirror full backups. For example, you can use RMAN proxy for weekly full backups. RMAN stream-based backups can be used to perform on-host incremental backups. In this way you can perform daily backups. NetBackup for SAP offers you the flexibility to intermix backup methods. You can specify both RMAN proxy copy off-host snapshot backups and RMAN stream-based on-host incremental backups. When you perform both types of backups, you remove the backup load from the production host during full backups. You also reduce the amount of data for incremental backups.

Using NetBackup for SAP with Snapshot Client to back up large databases

Cohesity recommends that customers with production databases back up their environments on a daily basis. This daily backup often is not feasible for databases of a size between 100 GB to over 1 TB.

This issue exists due to the following concerns:

- Server performance. The backup process for large databases can cause severe performance problems because the process consumes the database server's resources. CPU time, the system bus, the I/O bus, hard disk controllers, and

volume controllers become saturated. As a result, online use of the SAP system is limited and system performance is unpredictable during the backup.

- System availability. Traditionally, backup activities were carried out when there was little or no system activity. This time window usually occurred at night. In today's production environments, which require little or no system downtime, this window is small, if one even exists.
- Network performance. Instability and further performance loss may be experienced with a backup of large databases from the production host and over the network.

NetBackup for SAP with Snapshot Client supports split mirror backups. Split mirror backups are the recommended backup method for large databases because these backups overcome the preceding concerns.

Figure 12-1 SAP split mirror backup scenario

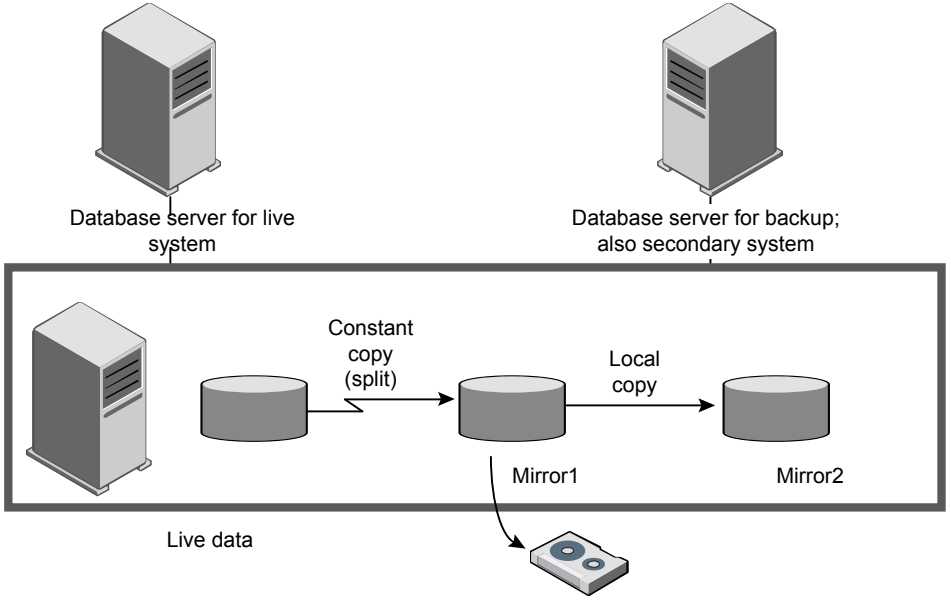

In SAP environments, the Snapshot Client technology supports the following major backup strategies:

- Off-host backup, which offers more performance. It offloads database backup activity and CPU cycles from the production host to the backup host. Thus, it improves the performance of the production environment.
- Snapshot backup, which requires no downtime of your production system. SAP supports both offline split mirror and online split mirror backups. In an online

split mirror backup, the production database remains available for user transactions. The need for backup windows is eliminated and 24/7 uptime functionality is provided for continuous business transactions.

- (UNIX or Linux) Block-Level Incremental (BLI) Backup. BLI backups decrease the amount of backup media that is required for incremental backups and to significantly reduce CPU and network overhead during backups. BLI, may not be used for incremental backups with SAP with RMAN.
- Split mirror backups. Because the mirrors are split from their standard devices and mounted on the backup server, the backup does not overload the network. The backup is run on the backup server without affecting the network.

How the NetBackup for SAP Snapshot Client works

A backup is initiated when:

- An automatic schedule runs the backup script that is specified in the Backup selections of the policy. The backup script contains `brtools` commands to start a backup.
- A user or external process on the client host invokes the `brtools` commands to start a backup.

The `brtools` include the `brbackup` command. The `brbackup` command initiates backup of the specified objects. The `brbackup` command provides a list of physical file names to NetBackup for SAP.

See "SAP database objects supported by advanced backup methods" on page 166.

The argument to the `brbackup` command `-d` parameter determines if the backup uses the backint or RMAN backup interfaces. Then either brbackup or RMAN provides the file list to NetBackup for SAP. NetBackup for SAP verifies that the policy it selected for the backup is configured with the appropriate Snapshot Client attributes, and then initiates file-based backups of the Oracle files, and performs the data movement.

Before NetBackup for SAP performs the backups, it requests that `brbackup` or RMAN put the data files being backed up into backup mode.

NetBackup then creates a snapshot of the files. When the snapshot is complete, NetBackup for SAP signals back to `brbackup` or RMAN to take the data files out of backup mode. The data files that are selected for backup are in backup mode for the time necessary to capture a snapshot of the data.

How NetBackup for SAP performs a split mirror backup using backint

Figure 12-2 lists the sequence of operations when you use `backint` to perform a backup.

Figure 12-2 SAP Oracle online and offline backups by using the `brbackup` command and `backint` command

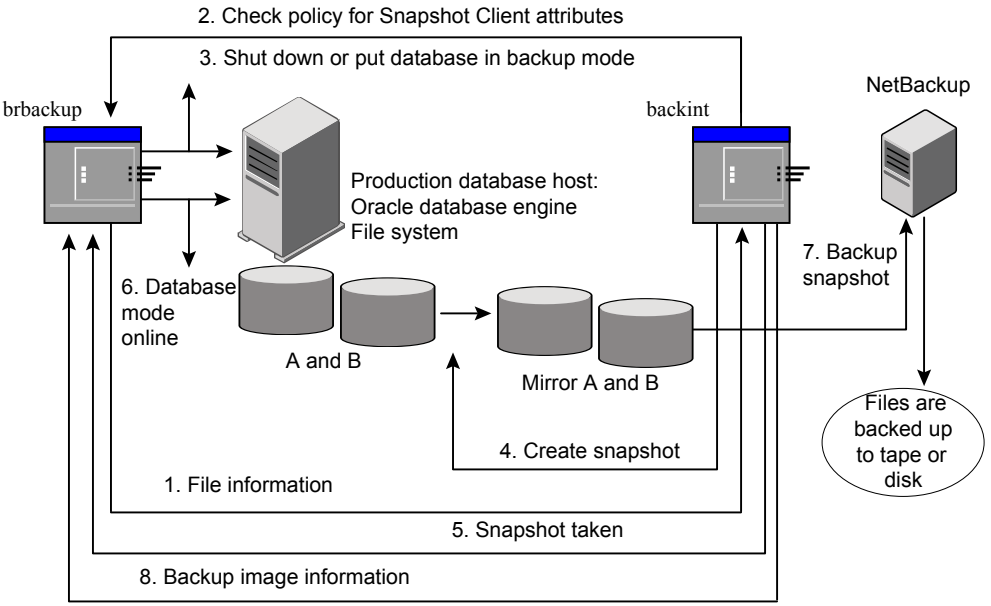

When NetBackup for SAP performs a backup by using the backint, the following occurs:

1 NetBackup for SAP receives a list of files to back up from `brbackup`.

2 NetBackup for SAP queries the policy to check whether the Snapshot Client policy attributes are specified.

When NetBackup for SAP is ready to back up the files, it sends a request to `brbackup` to put the data files into backup mode

Depending on the parameters with which the brbackup command was invoked, brbackup either shuts down the database or puts the database in backup mode.

- If `brbackup` is invoked with following command, `brbackup` shuts down the database:

  ```
  # brbackup -d util_file_online -t offline -m all -c force
  ```

- If `brbackup` is invoked with following command, `brbackup` puts the database or tablespace into backup mode:

  ```
  # brbackup -d util_file_online -t online -m all -c force
  ```

4 NetBackup for SAP quickly creates a snapshot of the files to back up. This operation typically takes a few seconds or minutes.

5 NetBackup for SAP notifies `brbackup` that the snapshot has been taken and that the database can be either started or taken out of backup mode.

6 `brbackup` either starts the database or takes the database out of backup mode.

7 NetBackup backs up the snapshot that it created.

8 NetBackup returns backup image information to `brbackup`.

See "The util_file parameter incompatible with Snapshot Client" on page 166.

See "How the NetBackup for SAP Snapshot Client works" on page 162.

How NetBackup for SAP performs a split mirror backup using RMAN

Figure 12-3 lists sequence of operations when you use RMAN proxy to perform a backup.

Figure 12-3 SAP Oracle online and offline backups by using RMAN proxy

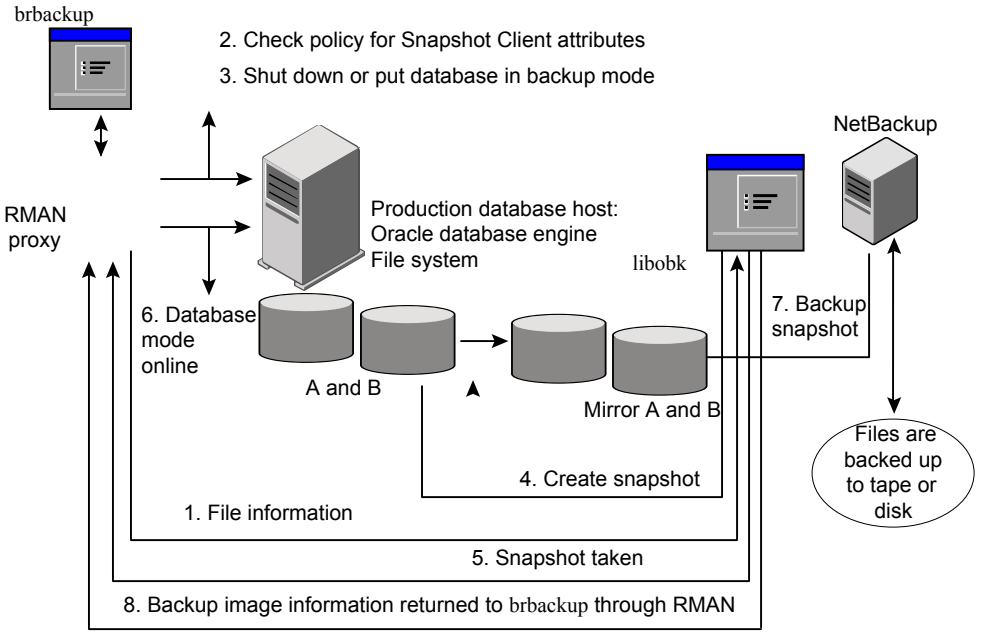

When NetBackup for SAP performs a backup by using RMAN, the following occurs:

1 NetBackup for SAP receives a list of files to back up from RMAN.

2 NetBackup for SAP queries the policy to check whether the Snapshot Client policy attributes are specified.

3 When NetBackup for SAP is ready to back up the files, it sends a request to RMAN. The RMAN puts the data files into backup mode or shuts down. NetBackup sends this request by using an SBT API.

4 Depending on the parameters with which the brbackup command was invoked, RMAN either shuts down the database or puts the database in backup mode.

- If `brbackup` is invoked with following command, RMAN shuts down the database:

  ```
  # brbackup -d rman_util -t offline -m full -c force
  ```

- If `brbackup` is invoked with following command, RMAN puts the database or tablespace into backup mode:

  ```
  # brbackup -d rman_util -t online -m full -c force
  ```

5 NetBackup for SAP quickly creates a snapshot of the files to back up. This operation typically takes a few seconds or minutes.

6 NetBackup for SAP notifies RMAN that the snapshot has been taken and that the database can be either started or taken out of backup mode.

7 RMAN either starts the database or takes the database out of backup mode.

8 NetBackup backs up the snapshot that it created.

9 NetBackup returns backup image information to `brbackup`.

The util_file parameter incompatible with Snapshot Client

NetBackup for SAP does not perform snapshot backups if the following parameter is specified on the `brbackup` command line:

```
-d util_file
```

When you specify this parameter, `brbackup` either shuts down the database or puts the database in backup mode. Then `brpbackup` calls the NetBackup for SAP Agent. The database remains in backup mode for the duration of entire backup.

Doing snapshot-based backups in this scenario is not useful. The snapshot technology adds no value because the database is either offline or in backup mode for several hours. If you initiate backups with this parameter in effect, NetBackup for SAP performs only standard backups directly to the storage unit. NetBackup for SAP ignores any snapshot-related configuration in the policy

About the NetBackup for SAP restore process

If you use the NetBackup for SAP Snapshot Client during a backup, you can use the typical NetBackup for SAP restore process. The Snapshot Client has no effect on the restore process.

Follow the existing NetBackup for SAP mechanism to perform restores of split mirror backups or block-level incremental backups (UNIX or Linux). If you use the instant recovery with rollback option, you must export the `SAP_RESTORE` environment variable.

SAP database objects supported by advanced backup methods

The `brbackup` command and Oracle RMAN control the kinds of database objects that can be backed up using Snapshot Client backup methods. Oracle allows proxy copy backups of databases, tablespaces, and data files.

For profile and configuration files, the NetBackup for SAP `backint` command performs standard backups directly to the configured storage unit. NetBackup for SAP uses advanced methods only when it moves Oracle data files.

The following is additional information on special file types for UNIX or Linux:

- Symbolic links and raw data files. NetBackup for SAP with Snapshot Client backs up and restores the data files that consist of symbolic links and regular files. NetBackup for SAP backs up and restores both the symbolic link and the file. If you select **Retain snapshots for Instant Recovery or SLP management**, the symbolic link must reside on the same file system as the data file. When you use instant recovery, if the symbolic link resides on a different file system than the data file it points to, the restore fails.

  NetBackup for SAP with Snapshot Client backs up and restores the data. Files that are created on raw partitions:

- Quick I/O data files. NetBackup for SAP with Snapshot Client backs up and restores Quick I/O Oracle data files. A Quick I/O file consists of two components: A hidden file with space allocated for it and a link that point to the Quick I/O interface of the hidden file.
  - On the backup, NetBackup for SAP with Snapshot Client follows the symbolic link and backs up both components of the Quick I/O file: The symbolic link and the hidden file.
  - On the restore, NetBackup for SAP with Snapshot Client restores both components from the backup image. If one or both of the components are missing, NetBackup for SAP with Snapshot Client creates the missing component(s).

Configuration requirements for snapshot backups with NetBackup for SAP

Each snapshot type has its own hardware requirements, software requirements, compatibility with certain features, and the snapshot methods that are supported. Special requirements apply for specific types of backups. See the NetBackup Hardware and Cloud Storage Compatibility List (HCL) for more information. Familiarize yourself with this information before you configure any snapshot backups.

The following list highlights some of the requirements that pertain to database agents:

- Snapshot Client backups do not back up all database objects. Your backup configuration must include schedules to perform snapshot and stream-based backups. This configuration ensures that the entire database can be restored successfully.

- On UNIX, the user identification and group identification numbers (UIDs and GIDs) associated with the files to be backed up must be available. The UID and GID must be available to both the primary client and the alternate backup client. The UID on the primary client and the alternate backup client must be the same. Similarly, the GID on the primary client and the alternate backup client must be the same.
- Ensure that the data files reside on a volume or a file system that does not contain archive logs, control files, or executables.
- One reason to have two different volumes is to separate the data files from the other files. If the logs are configured on the same volumes (or file systems) as the data files, the logs are temporarily frozen while NetBackup takes the snapshot. The process cannot access the logs when the database is active, so the database activity may freeze until the logs become accessible again.
  Another reason for writing the data files to their own repository is because it is required for an instant recovery point-in-time rollback. Only data files can exist on the volume that you want to restore.
- The hardware and software that is required for the appropriate snapshot method must be installed and configured correctly.
- NetBackup Snapshot Client must be installed and configured correctly, and the primary server must have a valid license for this option.
- To perform off-host backups, specify the off-host in the backup policy and ensure that host has the software and permissions to mount the snapshot.
- To perform Snapshot Client backups or restores, you need `brtools` version 6.40, patch level 36 or greater.

About configuring Snapshot Client with NetBackup for SAP

This topic explains how to configure snapshot and instant recovery backups for the SAP policy. For information on how a snapshot method is automatically selected and details on the types of backup methods, see the NetBackup Snapshot Manager for Data Center Administrator's Guide.

Snapshot backups do not back up all database objects. Your backup configuration must include one or more automatic schedules to perform snapshot backups and one or more application schedules to perform stream-based backups. This configuration ensures that the entire database can be restored successfully.

For snapshot or instant recovery backups, configure the following SAP policy and schedules as follows:

- On UNIX, snapshot methods for the file systems (raw partitions or logical volumes) on which the data objects (data files) reside.
- A backup method on the policy attributes dialog box.
- An Automatic Full Backup schedule to perform snapshot and off-host backups of the data files in `brbackup` phase 1.
- An Application Backup schedule to back up the profile files in `brbackup` phase 2. NetBackup for SAP does not support snapshot backups for `brbackup` phase 2.

Configuring a snapshot policy for NetBackup for SAP

The following procedure shows how to configure a snapshot policy with optional snapshot retention and off-host backup. For information on instant recovery policies, see the following topic.

See "Prerequisites for instant recovery backups (UNIX or Linux)" on page 172.

To configure a snapshot policy for NetBackup for SAP

1. Open the policy you want to configure.
2. Click on the **Attributes** tab.
3. Select the **SAP** policy type.
4. Select a policy storage unit from the **Policy storage** list.
5. Click **Perform snapshot backups**.
6. (Optional) Click **Snapshot options** to choose a snapshot method.

   By default NetBackup chooses a snapshot method for you. To choose a snapshot method, click **auto** (the default) or click one of the methods that are presented in the list.

   The snapshot method that you must use depends on your hardware environment and software environment. Only certain snapshot methods are supported in certain environments. See the NetBackup NAS Administrator's Guide and the NetBackup Snapshot Manager for Data Center Administrator's Guide. Also see the Hardware and Cloud Storage Compatibility List (HCL).

   You can configure only one snapshot method per policy. For example, assume that you want one snapshot method for clients a, b, and c, and a different method for clients d, e, and f. Then you need to create two policies for each group of clients and select one method for each policy.

7 (Optional) Select **Perform off-host backup**.

By default, the client that hosts the database performs the backup. If you want to reduce the I/O processing load on the client that hosts the database, specify an alternate client to perform the backup.

8 (Conditional) Select an off-host backup method.

The following off-host backup methods are available:

Use Alternate client (UNIX and Windows clients)	If you select **Alternate client**, also specify the name of the client to perform the backup. This option may require additional configuration. The alternate client must be a client that shares the disk array.
Use Data mover (UNIX clients only).	See "Configuring a NAS snapshot policy for NetBackup for SAP on UNIX/Linux" on page 175.

9 Click the **Schedules** tab.

10 Click **Add**.

11 Configure an Automatic schedule for the database files.

These files are backed up in `brbackup` phase 1.

12 (Conditional) Go to the **Instant Recovery** group, then select **Snapshots only**.

This setting suppresses NetBackup's default behavior, which is to copy the snapshot to a storage unit. When you select **Snapshots only**, NetBackup creates the on-disk snapshot copy of the database, but it does not copy the snapshot to a storage unit. The on-disk snapshot becomes the only backup copy. Note that the on-disk snapshot is not considered to be a replacement for a traditional backup.

13 Configure an Application backup schedule that backs up profile and configuration files, which are backed up in `brbackup` phase 2.

NetBackup for SAP copies the profile files or configuration files to the storage unit you select.

14 On the **Clients** tab, specify the clients to include in this policy.

15 On the **Backup selections** tab, specify a backup script.

See "About NetBackup for SAP backup types " on page 171.

16 Configure other attributes and add any additional schedules and backup selections.

About NetBackup for SAP backup types

The following backup type roles are available on the **Schedules** tab of the policy. You can use this information when you configure the NetBackup for SAP with Snapshot Client.

Table 12-2 Backup type roles

Backup Type	Description
Application Backup	The Application Backup schedule enables user-controlled NetBackup operations from the client. This schedule type allows operations the client initiates and operations an automatic schedule initiates from the NetBackup primary server. NetBackup uses the Application Backup schedule when the SAP user starts a backup manually. Configure at least one Application Backup schedule for each SAP policy. By default, an Application Backup is configured automatically. NetBackup for SAP uses the Application Backup schedule to back up the profile files that are backed up in `brbackup` phase 2.
Automatic Full Backup	An Automatic Full Backup schedule specifies the dates and times on which NetBackup is to start backups. Backups commence by running the SAP scripts automatically in the order they appear in the file list. If there is more than one client in the SAP policy, NetBackup for SAP runs the SAP scripts on each client. You must configure the Automatic Full Backup schedule in order for Snapshot Client methods to back up Oracle data files.
Automatic Differential incremental backup	In a differential incremental backup, NetBackup for SAP backs up all data blocks that changed since the most recent full or incremental backup. This type of backup can be used with the Snapshot Client BLI method.
Automatic Cumulative incremental backup	In a cumulative incremental backup, NetBackup for SAP backs up all data blocks that changed since the most recent full backup. This type of backup can be used only with the Snapshot Client BLI method or for RMAN stream-based incremental. Cumulative incremental backups reduce the work that is needed for a restore. You only need one cumulative incremental backup from any particular level at restore time. However, cumulative backups typically require more space and time than differential incremental backups. Cumulative backups duplicate the data that was captured in previous backups at same level.

Prerequisites for instant recovery backups (UNIX or Linux)

The following procedure explains the prerequisites to meet before you configure an instant recovery policy.

To perform prerequisites for an instant recovery backup (UNIX or Linux)

1. Shut down the database.
2. Copy the data files in the `/EPP/sapdata*` directory to a temporary location on another server:
3. Create multiple disk volumes and mount the disk volumes in their respective directories under `/EPP`.

   Create as many volumes as there are data file directories. For example, if you had data file directories named `sapdata1`, `sapdata2`, and `sapdata3`, create and mount the following new data file directories:

   - `/EPP/sapdata1`
   - `/EPP/sapdata2`
   - `/EPP/sapdata3`
4. Associate a snapshot mirror with all the directories you create.
5. Copy the data files from the temporary location to their respective directories.

   That is, copy the data files you moved in step 2 to the directories you created in step 3.
6. Start the database.

Configuring a snapshot policy with Instant Recovery for NetBackup for SAP

The following procedure describes how to configure a snapshot policy with Instant Recovery and optional snapshot retention and off-host backup. For information on standard snapshot policies, see the following topic.

See "Configuring a snapshot policy for NetBackup for SAP" on page 169.

To configure a snapshot policy with instant recovery for NetBackup for SAP

1. (UNIX or Linux) Perform the prerequisite configuration steps.

   See "Prerequisites for instant recovery backups (UNIX or Linux)" on page 172.
2. Open the NetBackup web UI.
3. Open the policy you want to configure.

4 Select the **Attributes** tab.

5 Select the SAP policy type.

6 Select a policy storage unit from the **Policy storage** list.

Select a policy storage unit in this step even if you plan to select **Snapshots only** later in this procedure. NetBackup for SAP uses this storage unit to back up the profile files and configuration files that are backed up in `brbackup` phase 2.

7 Select **Perform snapshot backups**.

8 (Optional) Select **Snapshot options** to choose a snapshot method.

By default, NetBackup chooses a snapshot method for you. To choose a snapshot method, click **auto** (the default) or click one of the methods that are presented in the list.

The snapshot method that you can use depends on your hardware environment and software environment. Only certain snapshot methods are supported in certain environments. See the NetBackup NAS Administrator's Guide and the NetBackup Snapshot Manager for Data Center Administrator's Guide. Also see the Hardware and Cloud Storage Compatibility List (HCL).

You can configure only one snapshot method per policy. For example, assume that you want one snapshot method for clients a, b, and c, and a different method for clients d, e, and f. Then you need to create two policies for each group of clients and select one method for each policy.

9 (Optional) Select **Perform off-host backup**.

By default, the client that hosts the database performs the backup. If you want to reduce the I/O processing load on the client that hosts the database, specify an alternate client to perform the backup.

10 (Conditional) Select an off-host backup method.

The following off-host backup methods are available:

Alternate client (UNIX and Windows clients)	If you select **Alternate client**, also specify the name of the client to perform the backup. This option may require additional configuration. The alternate client must be a client that shares the disk array.
Data mover (UNIX clients only)	See "Configuring a NAS snapshot policy for NetBackup for SAP on UNIX/Linux" on page 175.

11 Select the **Schedules** tab.

12 Select **Add**.

13 Configure an Automatic schedule for the database files.

These files are backed up in `brbackup` phase 1.

14 Go to **Instant Recovery** and select **Snapshots only**.

This setting suppresses NetBackup's default behavior, which is to copy the snapshot to a storage unit. When you select **Snapshots only**, NetBackup creates the on-disk snapshot copy of the database, but it does not copy the snapshot to a storage unit. The on-disk snapshot becomes the only backup copy. Note that the on-disk snapshot is not considered to be a replacement for a traditional backup.

15 Configure an Application backup schedule that backs up profile and configuration files, which are backed up in `brbackup` phase 2.

NetBackup for SAP copies the profile files or configuration files to the storage unit you select.

16 On the **Clients** tab, specify the clients to be included in this policy.

17 On the **Backup selections** tab, specify a backup script.

See "About NetBackup for SAP backup types " on page 171.

18 Configure other attributes and add any additional schedules and backup selections.

Additional configuration required for NetBackup for SAP with Network Attached Storage (NAS) on UNIX or Linux

Additional configuration is required if you want to use Network Attached Storage (NAS).

Table 12-3 Additional configuration for NAS

Backup type	Configuration
RMAN proxy backups with NAS	Keep all database data files on the NAS. Other database files should be located on the local file system. (These files include the online redo logs, control files, archive redo logs, Oracle executables, and the configuration files.)

Table 12-3 Additional configuration for NAS *(continued)*

Backup type	Configuration
BACKINT backups with NAS	Keep all database data files, online redo logs and all copies of control files on the NAS. Other database files including Oracle executables, configuration files (e.g. `SAPBACKUP` directory, `initCER.utl`) and the archive redo logs should be under the local file system.
Online backups	Use the `BR_CNTRL_DIR` environment variable to define the control file location for `brbackup`. The control file should be located on the NAS. (By default `brbackup` makes a copy of the control file in the `SAPBACKUP` directory, which is on the local file system. `brbackup` then provides that copy for backup).
Offline backups	No special configuration is required.

The following is an example of how to use the `BR_CNTRL_DIR` variable.

```
bash# export  BR_CNTRL_DIR=/oracle/CER/cntrl_loc
```

Where `/oracle/CER/cntrl_loc` should be on the NAS. It should not be located where the actual control file resides.

Configuring a NAS snapshot policy for NetBackup for SAP on UNIX/Linux

This topic describes how to configure a NetBackup for SAP NAS snapshot policy.

To configure a NetBackup for SAP NAS snapshot policy

1. Open the NetBackup web UI.
2. Open the policy you want to configure.
3. Select the **Attributes** tab.
4. Select the SAP policy type.
5. Select a storage unit.
   - NetBackup creates the snapshot on disk regardless of which storage unit you select.

- NetBackup for SAP uses the policy storage unit for backups of the profile and the configuration files. (The files that are backed up as part of `brbackup` phase 2). You can override this storage unit in the Application Backup schedule.

6 Select **Perform snapshot backups** and **Retain snapshots for instant recovery or SLP management**.

7 Select **Perform off-host backup**.

8 From the **Use** list, select **Data mover**.

9 From the **Machine** list, select **Network Attached Storage**.

See the policy configuration topic of the NetBackup NAS Administrator's Guide and the NetBackup Snapshot Manager for Data Center Administrator's Guide.

When the policy runs, NetBackup automatically selects the NAS_Snapshot method for creating the snapshot.

As an alternative, you can manually select the NAS_Snapshot method in the **Snapshot options** on the **Attributes** tab. Information about the **Maximum Snapshots (Instant Recovery only)** parameter is available.

10 Select the **Schedules** tab.

11 Select **Add**.

Configure both an Automatic backup schedule and an Application backup schedule, as follows:

- The Automatic backup schedule is for the database files.
  Locate **Instant recovery** and select **Snapshots only**. This option suppresses NetBackup's default behavior, which is to copy the snapshot to a storage unit.
- The Application backup schedule is for the archived redo logs and the control files.

12 Select the **Clients** tab.

Specify clients to back up with this policy.

13 Select the **Backup selections** tab.

14 Specify a backup script.

Review the information for how to use scripts for a NetBackup for SAP policy with Snapshot Client.

See "About NetBackup for SAP backup types " on page 171.

15 Configure other attributes and add any additional schedules and backup selections.

About configuring NetBackup for SAP block-level incremental backups on UNIX

If only a small portion of a database changes on a daily basis, full database backups are costly in terms of time and media. The Block-Level Incremental (BLI) Backup interface extends the capabilities of NetBackup to back up only the file system blocks that contain changed data blocks.

A database BLI backup is done at the file system block level, which means only changed file blocks are backed up. Unchanged blocks within the files are not backed up. The Arctera InfoScale Storage Checkpoint facility tracks changed blocks in real time. Accordingly, a BLI backup does not need to search the entire volume for the modified blocks at backup time. BLI backup saves time, decreases the amount of backup media that is required, and significantly reduces CPU and network overhead during backups. In addition, BLI backup allows more frequent backups, so backup images are more up to date.

BLI backup is particularly useful for any large databases that are sized in terms of hundreds of gigabytes or terabytes. Most traditional methods for database backup require that any change in the database—no matter how small—requires that the entire database is backed up. With BLI backup, only modified blocks (or file) need to be backed up.

See “Configuring policies for BLI backups with NetBackup for SAP” on page 180.

How BLI works with NetBackup for SAP (UNIX)

BLI backup supports two types of incremental backups: differential and cumulative. Full, differential incremental, and cumulative incremental backups are specified as part of the policy schedule configuration. When a restore is performed, NetBackup restores an appropriate full backup. Then it applies the changed blocks from the incremental backups.

Restoring any of the incremental backup images requires NetBackup to restore the last full backup image and all the subsequent incremental backups. The restore process continues until the specified incremental backup image is restored. NetBackup performs this restore process automatically, and it is completely transparent. The media that stored the last full backup and the subsequent incremental backups must be available, or the restore cannot proceed.

Note that restoring a file rewrites all blocks in that file. The first subsequent differential incremental backup and or all subsequent cumulative incremental backups back up all the blocks in the restored file. After an entire database is restored, the first subsequent backup results in a full backup.

The restore destination can be a Arctera InfoScale, UFS (Solaris), JFS (AIX), or HFS (HP-UX) file system. The destination VxFS file system does not need to support the Storage Checkpoint feature to restore files. However, a Arctera InfoScale file system with the Storage Checkpoint feature is needed to perform BLI backups of the restored data.

This topic uses the following terms to describe BLI backups:

- Full Backup.
  A backup in which NetBackup backs up each database file completely, not just data blocks that have changed since the last full or incremental backup.
- Cumulative BLI Backup.
  This type of backup is a backup of all the changed blocks in the database files since the last full backup. A cumulative BLI backup image contains only the data blocks of database files that changed since the last full backup. A cumulative BLI backup can reduce the number of incremental backup images that must be applied during a restore operation. This speeds up the restore process.
- Differential BLI backup.
  A backup in which NetBackup performs a backup of only those data blocks (within the database files) that changed since the last backup. The previous backup can be of type full, cumulative incremental, or differential incremental.

When NetBackup initiates BLI backups, it creates, manages, and uses the appropriate Storage Checkpoints of the filesystem(s) hosting the Oracle data file systems. These Storage Checkpoints identify and maintain a list of modified blocks.

About Nodata Storage Checkpoint and NetBackup for SAP

The Nodata Storage Checkpoint sets a bit to indicate if a file block changed. When you use Nodata Storage Checkpoints, the data files are left in backup mode for the duration of the backup. The amount of redo logs generated depends on the number of changes that were made during the backup.

To support BLI backup, the Arctera InfoScale file systems need extra disk space to keep track of the block change information. The space that is required depends on the database workload while the backup is running. For Nodata Storage Checkpoints, the additional space requirement by each file system is about 1% of the file system size.

Note: The default option that NetBackup uses for backups is Fulldata Storage Checkpoint. With this option, the NetBackup for SAP keeps the Oracle data files in backup mode only for the time that is needed to create a Storage Checkpoint.

About Fulldata Storage Checkpoint and NetBackup for SAP

The Fulldata Storage Checkpoint, once started makes a copy of the original file block at the moment before the block is about to be changed. When you use Fulldata Storage Checkpoints, tablespaces or data files are in backup mode for only a few seconds while the Storage Checkpoint is created. The database is put in this mode so the extra archive log space that is used is smaller.

However, space in the file system is needed to keep a copy of the original block of data that changed. If the workload is light during the backup an additional 10% of the file system size is usually sufficient. Also, if the backup window is relatively short (such as for incremental backups), an additional 10% of the file system size is usually sufficient. If the database has a heavy workload while a full backup is running, the file systems may require more space. The space requirement depends on the rate of change.

While archive log mode is required when the database is online, this mode provides the best recoverability for taking offline Storage Checkpoints, too.

Storage Checkpoint configuration on the NetBackup for SAP client

By default, the NetBackup for SAP with Snapshot Client for BLI backups uses the Fulldata Storage Checkpoint. When Fulldata Storage Checkpoint is in effect, the NetBackup for SAP agent keeps the Oracle data files in backup mode. The Oracle data files are kept in backup mode only for the time that is needed to create a Storage Checkpoint.

To change the default option to use Nodata Storage Checkpoint, a user must create the following file, which can remain empty:

```
/usr/openv/netbackup/ext/db_ext/NODATA_CKPT_PROXY
```

If the agent finds this file during run time, it uses Nodata Storage Checkpoint, and it keeps the data files in backup mode. The data files are kept in backup mode for the duration of the backup.

See “About Nodata Storage Checkpoint and NetBackup for SAP” on page 178.

See “About Fulldata Storage Checkpoint and NetBackup for SAP” on page 179.

Configuration requirements for BLI backups with NetBackup for SAP

Before you configure BLI backups, make sure that your configuration meets the following requirements:

- NetBackup for SAP is installed, licensed, and configured.

- NetBackup Snapshot Client is installed and configured, and the primary server must have a valid license for this option.
- Arctera InfoScale for Oracle must be installed and configured.
- Arctera InfoScale must have Storage Checkpoint licensed.

Configuring policies for BLI backups with NetBackup for SAP

This topic explains how to configure BLI backups for SAP policies. BLI backups do not back up all database objects. Include the policies to perform snapshot and stream-based backups.

Your backup configuration must ensure that the entire database can be successfully restored.

To configure a policy for BLI backups, configure the following:

- The BLI backup method on the policy attributes dialog box.
- An **Automatic backup** schedule to perform full and incremental snapshot backups of the data files.
- An **Application backup** schedule to back up profile and configuration files. These files are backed up during `brbackup` phase 2.

To configure a policy for BLI backups

1. On the left, select **Protection > Policies**.
2. Select **Add**.
3. Select the **Attributes** tab.
4. From the **Policy type** list, select **SAP**.
5. Select the **Policy storage**.
6. Select **Perform block-level incremental backups**.
7. To configure schedules, select the **Schedules** tab.

   SAP does not support snapshot backups of archive logs or the database profile and configuration files in the `brbackup` phase 2.

   To perform a whole database backup, configure the following:

   - One or more Automatic backup schedules to perform BLI backups of the data files.

     This backup automatically includes the control file.
   - An Application backup schedule type to back up the control files and archive logs and phase 2 profile and configuration files.

8 On the **Clients** tab, specify the clients to back up with this policy.

9 On the **Backup selections** tab, specify the script.

About the types of NetBackup for SAP BLI backups

NetBackup performs BLI backups with Automatic Full Backup, Automatic Differential Incremental Backup, and Automatic Cumulative Incremental Backup schedules.

NetBackup for SAP supports BLI backups only through `backint`. If you attempt to perform a BLI backup through RMAN proxy, the agent issues the following message:

```
BLIB with RMAN proxy is not a valid use case. Perform BLIB
through backint.
```

NetBackup for SAP checks that a full backup was performed before it proceeds with an incremental backup. If the NetBackup scheduler or user initiates an incremental backup, and NetBackup for SAP finds no record of a full backup using the same policy, it performs a full backup.

To ensure that it has a proper set of images to restore, NetBackup performs a full backup when it encounters the following situations:

- If NetBackup does not have a valid full backup image for the same policy in its database. For example, this situation can occur if images were expired.
- If a new file was added to or deleted from the list of files for an incremental backup.
- If the number of backup streams that is specified has changed from the previous backup.

NetBackup for SAP always initiates a full backup under these conditions, even if you want to perform an incremental backup.

About restoring individual files from a NetBackup for SAP snapshot backup

Data that is backed up with Snapshot Client methods is restored in the same way as data that is backed up without Snapshot Client methods.

See "Performing SAP restores using the SAPDBA utility (SAP on Oracle databases only)" on page 109.

Use this procedure for the files that were backed up with, or without, instant recovery enabled. In all cases, SAP determines the files that were backed up, and it initiates a corresponding restore request to the database agent.

If instant recovery is enabled,NetBackup attempts to restore the file by using the unique restore methods available with the instant recovery feature. The type of restore method thatNetBackup uses depends on your environment and the type of backup performed. IfNetBackup is unable to use any of the instant recovery methods, it restores the file in the typical manner. Data is copied from the snapshot to the primary file system. Information on the instant recovery methods that NetBackup uses is available. See the NetBackup Snapshot Manager for Data Center Administrator's Guide.

About NetBackup for SAP restores of volumes and file systems using snapshot rollback

You can request that an entire volume or an entire file system be restored from an instant recovery Snapshot backup. This type of a restore is called a point in time rollback. All the data in the snapshot is restored; single file restore is not available in a rollback.

You can perform a snapshot rollback from an instant recovery backup that was made with the following methods:

- UNIX: NAS_Snapshot
- UNIX: VxFS_Checkpoint snapshot
- vxvm snapshot
- FlashSnap snapshots

See the NetBackup Snapshot Manager for Data Center Administrator's Guide.

The following considerations are relevant for NetBackup for SAP restores:

- Snapshot rollback overwrites the entire volume.
- With NetBackup for SAP, snapshot rollback always performs file verification. The agent checks for the following:
  - The requested files (number and names) are identical to those in the snapshot
  - The primary volume does not contain any files that were created after the snapshot was made

  If verification fails, the rollback aborts with status 249.
- Use snapshot rollback with database files only. Control files and archive redo logs should exist on different file systems or volumes.

Performing a NetBackup for SAP snapshot rollback

You can use the `brrestore` command to perform a snapshot rollback restore. The additional configuration that is required depends on whether you used `backint` or RMAN for the snapshot backup, as follows:

- For the backups that are performed with `backint`, set the following environment variable in the script or on the command line:

  `SAP_RESTORE=rollback`

- For the backups that are performed with RMAN, add the following entry in the `init`*`SID`*`.sap` file:

  `rman_send = "'NB_ORA_PC_RESTORE=rollback'"`

Notes on NetBackup for SAP instant recovery restores

An SAP database instance generally contains three instances of control files. These reside in the following data and archive log directories:

UNIX or Linux:

`…/saparch/cntrl/cntrl`*`SID`*`.dbf`
`…/sapdata1/cntrl/cntrl`*`SID`*`.dbf`
`…/sapdata2/cntrl/cntrl`*`SID`*`.dbf`

Windows:

`…\saparch\cntrl\cntrl`*`SID`*`.dbf`
`…\sapdata1\cntrl\cntrl`*`SID`*`.dbf`
`…\sapdata2\cntrl\cntrl`*`SID`*`.dbf`

If the data files in `sapdata1` and `sapdata2` directories are on one volume, a snapshot of this volume contains these control files. However, at the time of the restore, `brrestore` provides only data files for restore and not control files. Therefore, the instant recovery with rollback method fails. For the rollback method, the files in the snapshot and the files in the restore file list should match. In this case, the files do not match, so instant recovery rollback restores method fail.

To avoid this problem, move the control files under locations `sapdata1` and `sapdata2` to some other location after you install your SAP database instance. First stop the database, then move the control files to different locations, adapt the profile file or `spfile` accordingly, and restart the database.

Performing SAP restores of volumes and file systems using block-level restore (Unix/Linux) or Fast File Resync (Windows)

This method requires the original primary data file to be present at the time of restore. The `brrestore` command deletes database files before restores to avoid permission problems. You can suppress this action by setting an environment variable. For example, use: `BR_NFD = 1` or `brrestore` command parameter `-NFD`.

Troubleshooting NetBackup for SAP rollback restores

If the rollback restore fails, it may be because the database still has a file open. Shut down and restart the database to try to correct this problem.

About NetBackup for SAP sample backup scripts (UNIX or Linux)

NetBackup for SAP installs sample scripts in the following location:

`/usr/openv/netbackup/ext/db_ext/sap/scripts/sap_oracle`

You can use some of the sample scripts as they are, or with a little modification. You can use them to perform snapshot backups.

The scripts are as follows:

- `sap_online_backup`
  This script sets environment variables and calls `brbackup` with the appropriate options to perform online backups of the Oracle database through `backint`.
- `sap_rman_backup`
  By default, this script performs offline backups of the Oracle database through RMAN.
  You can use this script to perform online snapshot backups of Oracle through RMAN proxy. Change the `brbackup` commands in this script to read as follows:

```
if [ $SAP_FULL -eq 1 ]; then
        CMD_LINE="$SAP_ENV brbackup -d rman_util -t online -m full -c"
elif [ $SAP_CINC -eq 1 ]; then
        CMD_LINE="$SAP_ENV brbackup -d rman_util -t online -m incr -c"
fi
```

If you want to perform backups by using the `brbackup` command through RMAN proxy, add following lines in the `init`*`SID`*`.sap` file:

```
rman_proxy = yes
rman_send = "'NB_ORA_PC_SCHED=auto_sched_name'"
```

About NetBackup for SAP sample backup scripts (Windows)

NetBackup for SAP installs sample scripts in the following location:

install_path\NetBackup\DbExt\SAP\samples

You can use some of the sample scripts as they are, or with a little modification. You can use them to perform snapshot backups.

The scripts are as follows:

- sap_online_backup.cmd
  This script sets environment variables and calls brbackup with the appropriate options to perform online backups of the Oracle database through backint.
- sap_rman_backup.cmd
  By default, this script performs offline backups of the Oracle database through RMAN.
  You can use this script to perform online snapshot backups of Oracle through RMAN proxy. Change the brbackup commands in this script to read as follows:

```
@if "%SAP_FULL%" EQU "1" @set CMD_LINE=%BRBACKUP% -u internal/ -c -d rman_util -t
online -m full
@if "%SAP_CINC%" EQU "1" @set CMD_LINE=%BRBACKUP% -u internal/ -c -d rman_util -t
online -m incr
```

If you want to perform backups by using the brbackup command through RMAN proxy, add following lines in the init*SID*.sap file:

```
rman_proxy = yes
rman_send = "'NB_ORA_PC_SCHED=auto_sched_name'"
```

Mixing RMAN stream and RMAN proxy NetBackup for SAP backups

You can use the split-mirror full backups that use RMAN proxy with the standard incremental backups that use RMAN stream methods. For example, use RMAN proxy file-based backups for weekly full backups using snapshot methods. You can

then use RMAN stream-based incremental backups for daily backups. To mix backup methods, perform the following additional configuration procedure.

To mix RMAN stream and RMAN proxy NetBackup for SAP backups

1. Copy `initSID.sap` to the `initFullSID.sap` and `initIIncrSID.sap` files.
2. Edit the `init_fullSID.sap` and `initI_incrSID.sap`.

   Edit the files as follows:

   - Specify the `rman_parms` parameter.
   - Set the `NB_ORA_SAP` environment variable to the `initSID.utl` path.
3. Edit the `init_fullSID.sap` file.

   Do the following:

   - Enable the `rman_proxy` flag.
   - Set the following, where “auto_sched” is the name of the automatic schedule:

     ```
     rman_send = "'NB_ORA_PC_SCHED=auto_sched'"
     ```
4. Copy following sample script to a different location on your client:

   on UNIX or Linux:

```
/usr/openv/netbackup/ext/db_ext/sap/scripts/sap_oracle/sap_rman_backup
```

   on Windows:

```
install_path\NetBackup\DbExt\SAP\samples\sap_rman_backup.cmd
```

5. On UNIX or Linux, modify the `if..elif..fi` condition as follows:

```
if [ $SAP_FULL -eq 1 ]
then CMD_LINE="$SAP_ENV brbackup -d rman_util -t offline -p initFullSID.sap -m full -c"
elif [ $SAP_CINC -eq 1 ]
then CMD_LINE="$SAP_ENV brbackup -d rman_util -t offline -p initIncrSID.sap -m incr -c"
```

6. On Windows, modify the condition as follows:

```
@if "%SAP_FULL%" EQU "1" @set CMD_LINE=%BRBACKUP% -u internal/ -c -d rman_util -t
offline -m full -p initFullSID.sap
@if "%SAP_CINC%" EQU "1" @set CMD_LINE=%BRBACKUP% -u internal/ -c -d rman_util -t
offline -m incr -p initIncrSID.sap
```

7. Specify the script from step 4 in the backup selection of your NetBackup for SAP policy

Performing user-directed snapshot backups with NetBackup for SAP

For scheduled backups, the NetBackup scheduler automatically exports or sets all the environment variables and system settings necessary to run the backups. However, if you want to initiate backups, you must set certain environment variables.

You can initiate a backup in one of the following ways:

- By running the `brbackup` command from the command line.
- By running a script that contains the `brbackup` command.

If you want to enable user-initiated backups, set the following in your environment:

- Set the following environment variable to the name of the automatic schedule to use for the phase 1 backup:

  SAP_SNC_SCHED=schedule_name_of_backup_type_to_perform

- Set the schedule parameter in the `init`*SID*`.utl` file to Application Backup schedule. NetBackup for SAP uses this schedule for backing up control files, configuration files, and archive redo log files.
- Additionally, if you plan to run the `sap_rman_backup` script from the command line, export the following before you run the script:
  `SAP_FULL = 1` (for performing full backups)
  `SAP_CINC = 1` (for performing cumulative incremental backups)

See "Mixing RMAN stream and RMAN proxy NetBackup for SAP backups" on page 185.

See "About NetBackup for SAP sample backup scripts (UNIX or Linux)" on page 184.

See "About NetBackup for SAP backup types " on page 171.

Chapter 13

NetBackup for SAP on MaxDB databases

This chapter includes the following topics:

About NetBackup for SAP on MaxDB databases

NetBackup for SAP requires different configuration for a MaxDB database than for an Oracle database.

For more information about MaxDB databases, see your MaxDB documentation.

Note: NetBackup documentation uses the terms “media” or “medium” to refer to removable media tape in a storage unit. NetBackup writes a backup file to a storage medium. MaxDB database documentation, however, uses the term “backup medium” to include the named pipes and other entities that are needed for performing backups. This manual uses the term backup medium to refer to backing up MaxDB database files.

About configuring NetBackup for SAP on MaxDB databases

Configure NetBackup for SAP for the MaxDB database environment. Follow only the instructions in that section that pertain to MaxDB databases. MaxDB databases also require you to create a backup medium.

Creating a backup medium for NetBackup for SAP

As part of the configuration process, also create at least one backup medium. Create this medium before you perform any backups. databases require at least one backup medium before a backup or restore can be performed.

To create a backup medium

- Use the DBM CLI interface and enter the following commands:

```
OS_prompt% dbmcli -d database_name -u usr,passwd
dbmcli> medium_put medium_name medium_path PIPE backup_type
```

where

database_name	Name of the SAP database
usr	user name.
passwd	password.
medium_name	Name of the backup medium. The first four characters must be `BACK`. The *medium_name* must be in format `BACK`*name*. For example, `BACKData`.

medium_path	Full path to the named pipe. For example:
	`/export/home/medium1`
	or
	`\\.\PIPE\medium1`

backup_type	Specify one of the following: ■ `DATA` for a full data backup ■ `PAGES` for an incremental data backup ■ `LOG` for a log backup

Using NetBackup for SAP on a MaxDB database to perform backups and restores

SAP on a MaxDB database does not include the BR*Tools. Instead SAP provides completely different methods of initiating a backup or restore.

You can initiate backup and restore operations through the MaxDB administrative interface when you invoke one of the following interfaces:

- DBM GUI
- DBM CLI
- Web DBM

For example, when you use the MaxDB administrative interface through the DBM CLI, the following occurs:

- The DBM CLI communicates with the MaxDB server, which creates one or more named pipes.
- The MaxDB server streams the data from the MaxDB instance to the named pipes.
- The MaxDB server calls NetBackup for SAP with a list of named pipes.
- NetBackup for SAP reads named pipes and sends data to NetBackup.

In case of a restore, the process is same except that data goes from NetBackup for SAP to the MaxDB server by the named pipes.

Performing a MaxDB database backup

The following procedure shows you how to perform a MaxDB database backup using NetBackup for SAP.

To perform a backup

1. Make sure that you have modified the proper configuration files and created a backup medium.
2. Type the following command to start a utility session:

```
dbmcli> util_connect
```

3. Type the following command to initiate the backup:

```
dbmcli> backup_start medium_name backup_type
```

where:

medium_name	Name of the backup medium.
backup_type	Specify either `DATA`, `PAGES`, or `LOG`.

Performing a MaxDB database query

Prior to performing a restore, use the following procedure to query backup information and to obtain a listing of backup information. This information is needed for a restore. The following procedure shows you how to perform a MaxDB database query using NetBackup for SAP.

To perform a query

- Type the following commands to obtain the BID:

```
dbmcli> backup_ext_ids_get medium_name database_name server
dbmcli> backup_ext_ids_list
```

where:

medium_name	Name of the backup medium.
database_name	Name of the SAP database.
server	Name of the server that hosts the MaxDB database.

Performing a MaxDB database restore

The following procedure shows you how to perform a MaxDB database restore using NetBackup for SAP.

To perform a restore

1. To obtain the external backup identification information, perform a query.

   See "Performing a MaxDB database query" on page 191.

2. Type the following command to initiate the restore:

   ```
   dbmcli> recover_start medium_name backup_type ExternalBackupId
   "database_name BID medium_path"
   ```

   where:

medium_name	Name of the backup medium.
backup_type	Specify either `DATA`, `PAGES`, or `LOG`.
BID	The external backup identifier. This information is returned when you perform a backup.
database_name	Name of the SAP database.
medium_path	Specify full path to the named pipe. For example, /export/home/medium1 or \\.\PIPE\medium1.

Performing SAP backups and restores using parallel medium groups

Parallel backups and restores are used to improve backup and restore performance.

Note: Parallel backups can only be used with backup_type data, not PAGES, or log.

- You can combine individual backup mediums to form parallel mediums. When parallel mediums are used for backup, the MaxDB database passes two or more named pipes to NetBackup for SAP. NetBackup for SAP starts as many backup streams as there are named pipes. It processes each named pipe independently from the rest of the named pipes in the backup job. If the policy and the storage

unit are configured for multiplexing, the streams can be interleaved on the same tape.

- Even if a backup is done in parallel, you can perform a restore of that data in either the parallel or sequential mode. In other words, parallel backups do not require parallel restores.
  When the MaxDB database requests a parallel restore by listing multiple named pipes in an input file, NetBackup for SAP processes each pipe independently from the rest of the named pipes. Note that the MaxDB database requires that the number of media in a group of parallel media equal the number of media used during backup.

Chapter 14

Troubleshooting NetBackup for SAP and SAP HANA

This chapter includes the following topics:

- Parent and grandparent jobs of SAP HANA backups fail

About troubleshooting NetBackup for SAP

NetBackup for SAP contains several processes and resources that can help you to troubleshoot database backups. These resources include the logs and reports that NetBackup, NetBackup for SAP, and the SAP tools provide. These reports are useful for finding the errors that are associated with those applications.

NetBackup debug logs and reports

Debug logs

The NetBackup server and client software let you enable detailed debugging logs. The information in these log files can help you troubleshoot the problems that occur outside of either the database agent or SAP commands.

Note the following for these logs:

- These logs do not reveal the errors that occur when SAP commands is running unless those errors also affect NetBackup. SAP may (or may not) write errors in the application to the NetBackup logs. Your best sources for SAP error information are the logs provided by SAP.
- Generally, each debug log corresponds to a NetBackup process and executable.

Information about the debugging log files is available in the NetBackup Troubleshooting Guide.

Also refer to the following file:

```
UNIX: /usr/openv/netbackup/logs/README.debug file
```

Reports

NetBackup provides other reports that are useful in isolating problems. One such report is **All logs entries** on the server. For more information see the NetBackup Web UI Administrator's Guide.

Enable the debug logs automatically (SAP client) (Windows)

You can enable debug logging by running a batch file that creates each log directory. To create all log file directories automatically, run the following:

```
install_path\NetBackup\logs\mklogdir.bat
```

Or, you can manually create the directories for the log files you want created.

Enable the debug logs manually (NetBackup for SAP) (Windows)

To enable the NetBackup for SAP for Windows database agent logs manually

1. Create the following directories on the client:

   install_path\NetBackup\logs\bpbackup

   install_path\NetBackup\logs\bpbkar32

   install_path\NetBackup\logs\bphdb

   install_path\NetBackup\logs\bprestore

   install_path\NetBackup\logs\tar32

   install_path\NetBackup\logs\backint

   For example:

   ```
   cd install_path\NetBackup\logs
   mkdir bphdb
   ```

2. (Conditional) If you are running NetBackup for SAP on Oracle with RMAN, create the following additional directory:

   install_path\NetBackup\logs\dbclient

3. Verify that the `user_ops` directory, each of the log directories, and any subdirectories exist and are accessible for all applications to operate correctly.
4. Enable logging for the `nbpem`, `nbjm`, and `nbrb` scheduling processes, which use unified logging.

   NetBackup writes unified logs to *install_path*\NetBackup\logs.

   You do not need to create log directories for processes that use unified logging. For information on how to use logs and reports, see the NetBackup Troubleshooting Guide.

Enable the debug logs manually (SAP) (UNIX/Linux)

This topic describes how to manually create the directories that are used for debug logging. More information is available on how on how to use logs and reports.

See the NetBackup Troubleshooting Guide.

To enable the debug logs manually

1. Create the following directories on the client:

   `/usr/openv/netbackup/logs/bpbkar`

   `/usr/openv/netbackup/logs/bphdb`

   `/usr/openv/netbackup/logs/tar`
   `/usr/openv/netbackup/logs/backint`

   For example:

   ```
   cd /usr/openv/netbackup/logs
   mkdir bphdb
   ```

2. (Conditional) If you run NetBackup for SAP on Oracle with RMAN, create the following additional directory:

   `/usr/openv/netbackup/logs/dbclient`

3. Verify that the `user_ops` directory, each of the log directories, and any subdirectories exist and are accessible for all applications to operate correctly.

   See "About permissions for NetBackup for SAP log files (UNIX)" on page 72.

4. Enable logging for the `nbpem`, `nbjm`, and `nbrb` scheduling processes that use unified logging.

   NetBackup writes unified logs to `/usr/openv/logs`.

   You do not need to create log directories for the processes that use unified logging.

About the bphdb directory on the Windows database client

The *install_path*`\NetBackup\logs\bphdb` directory contains log files.

The following types of logs exist:

- `sap_stdout.`*mmddyy*`.`*hhmmss*`.txt`

  Unless it is redirected elsewhere, NetBackup writes SAP script output to this file.

- `sap_stderr.`*log*`.`*mmddyy*`.`*hhmmss*`.txt`

  Unless it is redirected elsewhere, NetBackup writes SAP script errors to this file.

- *mmddyy*`.log`

This log contains debugging information for the `bphdb` process. `bphdb` is the NetBackup database backup binary. It is invoked when an automatic backup schedule is run. NetBackup for SAP uses this client process for SAP script execution.

About the bphdb directory on the UNIX database client

The `/usr/openv/netbackup/logs/bphdb` directory contains logs.

The following types of logs exist:

- `sap_stdout.`*`mmddyy`*

  Unless it is redirected elsewhere, NetBackup writes SAP script output to this file.

- `sap_stderr.`*`mmddyy`*

  Unless it is redirected elsewhere, NetBackup writes SAP script errors to this file.

- *`mmddyy`*

  This log contains debugging information for the `bphdb` process. `bphdb` is the NetBackup database backup binary. It is invoked when an automatic backup schedule is run. NetBackup for SAP uses this client process for SAP script execution.

About the backint directory on the UNIX database client

The `/usr/openv/netbackup/logs/backint` directory contains run logs.

The following run log exists:

- *`log.mmddyy`*

See "Enable the debug logs manually (SAP) (UNIX/Linux) " on page 196.

About the backint directory on the Windows database client

The *`install_path`*`\NetBackup\logs\backint` directory contains execution logs, as follows:

This log contains debugging information and execution status for the SAP NetBackup client processes program provided with NetBackup for SAP.

Set the debug level on a SAP client (Windows)

To control the amount of information that is written to the debug logs, change the Database debug level. Typically, the default value of 0 is sufficient. However, technical support may ask you to set the value higher to analyze a problem.

To set the debug level

1. Open the **Backup, Archive, and Restore** interface.
2. Select **File > NetBackup Client Properties**.
3. Click the **Troubleshooting** tab.
4. Set the **General** debug level.
5. Set the **Verbose** debug level.
6. Set the **Database** debug level.
7. Click **OK** to save your changes.

Set the debug level on a SAP client (UNIX)

To control the amount of information that is written to the debug logs, change the "Database" debug level. Typically, the default value of 0 is sufficient. However, Technical Support may ask you to set the value higher to analyze a problem.

The debug logs are located in `/usr/openv/netbackup/logs`.

To set the debug level

- Enter the following line in the `bp.conf` file.

  `VERBOSE = X`

  Where *X* is the debug level you want.

sapdba logs and messages (Oracle-based SAP environments only)

The SAP tools log provides information on SAP operations. You can check the log files to determine the ultimate success or failure of database backups and restores.

You can view the backup and restore logs in the following ways:

- Through the following `sapdba` menu options: `Show>Cleanup` and `Show log files>profiles`.
- In the directories for `brbackup` and `brrestore` log information and `brarchive` log information.

NetBackup for SAP backup and restore log files

The following directories contain log files for different types of backups and restores:

Windows:

```
%SAPDATA_HOME%\sapbackup
```

UNIX or Linux:

```
$SAPDATA_HOME/sapbackup
```

The files in this directory are named according to the following pattern:

- The summary log file is named `back`*`SID`*`.log`, where *SID* is the unique name for the Oracle database instance.
- The detail log files are named *`encoded_timestamp`*`.`*`xyz`*, where:

encoded_timestamp	A timestamp that is used in each detail log name that guarantees a unique file name.
xyz for backup logs:	
x	`a` represents all, `i` represents incremental, `p` represents partial.
y	`n` represents online. `f` represents offline.
z	`f` represents `util_file` backup, `r` represents `rman_util` backup
xyz for restore logs:	
xyz	`rsb` represents restore backup files.
xyz	`rsa` represents restore archive files.
xyz	`rsf` represents restore individual files.

NetBackup for SAP archive log files

The following directories contain log files for different types of archive activities:

Windows:

```
%SAPDATA_HOME%\saparch
```

UNIX or Linux:

```
$SAPDATA_HOME/saparch
```

The files in this directory are named according to a pattern. The *encoded_timestamp* is a timestamp used in each detail log name that guarantees a unique file name.

Table 14-1 lists the suffixes that differentiate the different types of archive logs.

Table 14-1 File suffixes

File name	Meaning
encoded_timestamp.sve	Original saved.
encoded_timestamp.svd	Original saved and deleted.
encoded_timestamp.cpy	Original copied or saved a second time.
encoded_timestamp.cpd	Original copied or saved a second time and deleted.
encoded_timestamp.dcp	Deleted that were saved twice.
encoded_timestamp.dsv	Deleted that were saved.

NetBackup SAP Oracle backup jobs fail with error 90

In case of `sort_backup_type custom` the backup job fails with status code 90:

- Ensure that the custom sort file stream is less than the number of SAP oracle DB table spaces.
- If the backup is running on a single volume (SAP device) with custom sort option, ensure that the custom sort file must have less than or equal to seven streams.

Minimize time-out failures on large database restores

Large database restores sometimes fail when multiple restore sessions compete for resources. In this situation, a restore session can be delayed while NetBackup waits for media or device access. If the delay is too long, the restore session times out. Use the following procedure to minimize session time-outs and to allow the restore jobs to complete successfully.

To minimize time out failures on large database restores

1. Open the NetBackup web UI.
2. On the left, select **Hosts > Host properties**.
3. Select the check box for the client.
4. If necessary, select **Connect**.
5. Select **Edit client**.
6. Select **Timeouts**.
7. Set the **Client read timeout** property to a large value.

   The default for the **Client read timeout** setting is 300 seconds (5 minutes). For database agent clients, increase the value significantly from the recommended value.

   See the NetBackup Web UI Administrator's Guide.

   For example, change this setting to 30-60 minutes to minimize time-out errors.
8. Select **Save**.

Note: This change may delay detecting problems during subsequent backups. Consider putting the original value back in place once any restore that requires a change is complete.

About troubleshooting NetBackup for SAP HANA

The following sections contain several processes and resources that can help you to troubleshoot the issues that are related to SAP HANA.

NetBackup SAP HANA backup jobs fail with error 41 and 25

Tape backups on a multi-node and single-node SAP HANA appliance fail. The solution is to change the **Client connect timeout** and **Client read timeout** to 3000 seconds.

The default **Client read timeout** is 300 seconds.

To change the timeout period

1. Open the NetBackup web UI.
2. On the left, click **Hosts > Host properties**.
3. Select the primary server. If necessary, click **Connect**. Then click **Edit primary server**.

4. Click **Timeouts**. Set the **Client connect timeout** and the **Client read timeout** property to a value of **3000** seconds.

   See the NetBackup Web UI Administrator's Guide for more information on this setting.

5. Click **Save**.

Recovery of SAP HANA database from tape storage fails

Recovery of SAP HANA database from tape storage fails to complete successfully as the restore jobs hang indefinitely. This failure is specifically seen during the logs restore phase.

By default, SAP HANA does not read all the pipes that it has opened during recovery. As NetBackup processes the pipes randomly, it is possible that the pipe where NetBackup tries to write is not open. This situation creates a deadlock situation and hence the jobs are stalled in the Activity Monitor.

To streamline the read and write operations on the pipe, set the following parameter in the `global.ini` file:

In HANA Studio, double-click on instance tab.

Instance-> **Configuration**-> **global.ini**->**backup**-> enter **1** as the max_recovery_backint_channels value (by default, this value is 64).

This setting ensures that only a single request is sent to NetBackup. Hence only a single pipe is open for reading and NetBackup writes on that pipe only.

Note: This issue is not seen in case of disk recovery. Hence no change is required for max_recovery_backint_channels in case of disk recovery.

Log backup fails with status 50 for SAP HANA

Log backup fails with status 50 for SAP HANA

A log backup fails if a recovery job is also initiated at the same time. A recovery job shuts down the database and log backups fail with error 50.

When a recovery job is initiated, the SAP HANA Studio forces a shutdown of the system and the backup jobs fail. To avoid the shutdown and subsequent backup failure, stop the system. You can then start the recovery job.

Disaster recovery of an SAP HANA database

Use the SAP HANA studio for disaster recovery of SAP HANA SPS 09 and later. For disaster recovery, follow the same steps as used for the redirected restore of an SAP HANA instance using the SAP HANA studio.

See "Using SAP HANA studio for SAP HANA instance redirected restore" on page 144.

Cannot browse backup catalog during SAP HANA restore operations via HDBStudio

The following error is seen:

Node name list file not found.

Description:

The node_names.txt file located on the SAP Hana client /usr/openv/netbackup/ext/db_ext/sap/ path is required for SAP Hana database recovery operations.

It should contain the hostname (shortname and FQDN) of the SAP Hana client if on single or multi hosts and node host names of the cluster.

The following error messages are displayed in HDBStudio during a restore attempt:

```
line 1: syntax error: unexpected end of file /usr/bin/sh: error
importing function definition for `which' Error: Node name list file
not found. Exiting. , [110088] Error reading backup from 'BACKINT'
'/usr/sap/RH1/SYS/global/hdb/backint/DB_TENANTDB/log_backup_0_0_0_0',
[110202] Not all data could be read from
/usr/sap/RH1/SYS/global/hdb/backint/DB_TENANTDB/log_backup_0_0_0_0.
Expected to read 4096, but received 0 bytes
```

Solution or workaround:

Ensure that the `/usr/openv/netbackup/ext/db_ext/sap/node_names.txt` file exists with the appropriate host names on the SAP HANA clients.

SAP HANA databases that are no longer active are not immediately removed from NetBackup asset

Description:

NetBackup assets are not immediately updated when databases are dropped or become inactive. NetBackup cleanup any asset not discovered in the last 30 days that has not been added to a policy.

Solution or workaround:

Remove the inactive database using the delete action.

Credential validations fail during instance registration even though credentials appeared to be correct

Description:

SAP HANA instance credential validation requires access or connection to the NetBackup clients or SAP HANA database servers and login to the SAP HANA SYSTEMDB database.

When manually creating the SAP instance in NetBackup, the host name, NetBackup client name and NetBackup host name may not match.

Solution or workaround:

Ensure that the host names specified in the SAP HANA instance matches with the NetBackup host name or client name and there are no network host name resolution issues.

Make sure the primary server can connect to the target SAP HANA clients and the credentials being used are correct.

Use NetBackup tool to test connectivity within NetBackup hosts:

```
bptestbpcd -host <sap host name> -bpclntcmd_test -M <Primary server name> - verbose
```

Make sure the SAP HANA database server instance is running and test login using the same credentials like the following command:

```
hdbsql -n <host name> -i <instance#> -u SYSTEM -p <pw> -d SYSTEMDB
```

SAP HANA databases are not automatically discovered after instance registration

Description:

SAP HANA databases are not automatically discovered after a successful instance registration. It requires a user action to explicitly discover the databases before it can be added in to the policy for backups to run successfully.

Solution or workaround:

Manually discover the databases after a successful instance registration. This can be done via API or web UI:

Go to **Workloads > SAP Databases** tab, click **Discover** databases.

Backup failures occur when SAP HANA databases are included in multiple intelligent policies

Description:

Including SAP HANA databases in multiple intelligent policies results in backup failures and will temporarily leave backint UTL files inactive on the client.

NetBackup does not restrict databases from being included in multiple SAP intelligent policies, resulting in backup failures due to overlapping backup windows that trigger parallel executions.

Solution or workaround:

Configure one SAP intelligent policy per instance or database and ensure backup window schedules do not overlap to avoid parallel execution.

Parent and grandparent jobs of SAP HANA backups fail

Description:

Parent and Grandparent jobs of SAP HANA backups fail with the following error:

```
The backup failed to back up the requested files (6)
```

Snippet of job details:

```
Info bphdb (pid=40259) INF - SAP HANA backup statement: SQL execution
failed with status: [SAP AG][LIBODBCHDB SO][HDBODBC] General error;447
backup could not be completed: [110090] Error while backing up backup
catalog, [110507] Backint exited with exit code 1 instead of 0.
console output: /bin/sh: which: line 1: syntax error: unexpected end
o.
```

Solution or workaround:

The SAP intelligent policy sets the configuration changes to `global.ini` for SYSTEM and sets the dynamic path for `utl` files with the `policyname prefix`.

The file will be persisted in the following pat:h `/usr/openv/tmp/`

If DATABASE has the path of `initSAP.utl` file, it takes the precedence over SYSTEM and backup fails with error 6.

The solution to resolve the issue is to clear the path for the `backup_parameter_path` configuration option of DATABASE.

Appendix A

backint command line interface

This appendix includes the following topics:

About backint command line interface for SAP

The NetBackup for SAP `backint` interface communicates instructions from the SAP tools to NetBackup. The `backint` interface implements the SAP system's BC-BRI BACKINT Interface specification.

See "About the BC-BRI BACKINT Interface" on page 211.

Note: Cohesity recommends that you use the `backint` command only with guidance from a technical support staff member.

The `backint` interface performs the following functions:

- The backup function. The backup function of the `backint` interface supports and defines the SAP `brbackup` and `brarchive` tools to NetBackup. `brbackup` and `brarchive` communicate with the `backint` interface through an *in_file* and an *out_file* parameter. The *in_file* parameter includes a list of files to be backed up or archived. The *out_file* parameter reports the status for each file and assigns a backup ID (BID) to each file. In the event of a partial backup, this function can identify successfully backed up files to the user.
- The restore function. The restore function of the `backint` interface supports and defines the `brrestore` tool to NetBackup. It communicates with the `backint`

interface through the *in_file* parameter and *out_file* parameter. The *in_file* parameter includes a list of files to be restored through NetBackup. It also includes the BID assigned during the backup function. The *out_file* parameter contains the status of the restore for each file. When the NetBackup restore operation is complete, the restore function lists successfully restored files. It also lists BIDs used during the operation.

During the backup function, NetBackup assigns the BID. It can identify one or more backup runs, a single file backup, or a group of files. During a backup function, the BID is submitted to the *out_file* parameter. During the restore and the inquiry functions, the BID can only be set in the *in_file* parameter.

If the BID is not set, the restore function uses the BID of the last backup. As an option, this function can also include a list of directories into which files are restored.

- The inquiry function. The inquiry function supports and defines the `sapdba` tool to NetBackup. `sapdba` uses the *in_file* parameter and the *out_file* parameter to collect backup information. The *in_file* parameter contains optional BIDs and file names.

  If only a `#NULL` is received on the *in_file* parameter, a list of BIDs is generated to the *out_file* parameter. If a BID is received, a list of files belonging to the BID is generated. If a file name is entered along with the `#NULL`, a list of BIDs containing that file is listed.

The `backint` command line uses the following syntax:

```
backint -u user_id -f function [-t type] -p par_file [-i in_file] [-o out_file]
```

Table A-1 shows the `backint` command options.

Table A-1 `backint` command options

Option	Arguments and purpose
`-u user_id`	Required. UID for backup utility user. No default.
`-f function`	Required. Defines the function that SAP tools requests of NetBackup for SAP: ■ `backup` - If backup is specified, NetBackup backs up the files in the list that SAP tools provide. ■ `restore` - If restore is specified, NetBackup restores the files in the list that SAP tools provide. ■ `inquiry` - If inquiry is specified, NetBackup returns the saved, or the not saved status of the files in the list that SAP tools provide.

Table A-1 `backint` command options *(continued)*

Option	Arguments and purpose
`-t type`	Optional. It defines the type of backup that NetBackup for SAP should perform. If no *type* is specified, the default value of `file` is used. Specify one of the following arguments: ■ `file` – All data files are either offline or in backup mode. NetBackup for SAP can back them all up without coordination with SAP tools. ■ `file_online` – Allows NetBackup for SAP to request that SAP tools set each tablespace into `#BEGIN/#END` backup mode when a related file backup takes place. It is used for online backups only. The architecture is based on the switch files that are defined in the `-p par_file` parameter.
`-p par_file`	Required. A text file that contains parameters (required and optional), and parameter values. These parameters determine the backup and restore procedure between the SAP tools and NetBackup for SAP and between NetBackup for SAP and NetBackup. The following list explains the components of *par_file*: ■ A pound character (#) in column one denotes a comment. Any other character in column one is considered valid. ■ All required parameters must be specified with a valid value before the NetBackup for SAP `backint` interface can run correctly. ■ You can comment out an optional parameter by placing a # in column one. If optional parameters are specified, they must have valid values for the NetBackup for SAP `backint` interface to run correctly. If an invalid parameter name is found, the NetBackup for SAP `backint` interface reports a warning message and continues. The SAP tools parameter file specifies the location of this *par_file*. For an example file, see the following: on UNIX or Linux: `/usr/openv/netbackup/ext/db_ext/sap/scripts/sap_oracle/initSAP.utl` on Windows: *install_path*`\NetBackup\dbext\sap\samples\initSAP.utl` See “About parameters used in initSID.utl” on page 227.

Table A-1 `backint` command options *(continued)*

Option	Arguments and purpose
`-i in_file`	Optional. Specifies a text file, the contents of which list the files to backup, inquire, or restore. The format can vary. If this option is not specified, the content of this file is data from standard input. See "About backint -i in_file contents" on page 216.
`-o out_file`	Optional. Specifies a text file that contains each file that is processed and associated status. May also contain other process messages. If this option is not specified, the output is written to standard output. See "About backint -o out_file contents" on page 218.

About the BC-BRI BACKINT Interface

One of the key outputs and then input parameters of the BC-BRI BACKINT Interface specification is an external backup ID (BID). The BID is a unique identifier that a successful backup NetBackup for SAP creates. After a backup, it is returned to SAP tools and stored for use for future inquire or restore functions.

The parameter has two formats:

- VXF<ctime> - NetBackup for SAP Oracle
- VXP<ctime> - NetBackup for SAP MaxDB

It is similar to the backup ID created by NetBackup (*client_name_ctime*) for all backup images but is not stored within NetBackup. The *ctime* in the BID matches the *ctime* in the NetBackup backup ID. Therefore allowing NetBackup for SAP to make efficient image queries to the image database on the primary server.

About backint command line interface for SAP HANA

To run backup and recovery you need to provide the `backint` command with additional command line options.

Table A-2 lists the `backint` command options

Table A-2 `backint` command options

Parameter	Description	Values
`-u <user_id>`	The user ID defines a common name space for file names in a multi-host SAP HANA database. The user ID is bound to a database instance and used by all the servers in this instance.	
`-f <function>`	The requested function	backup, restore, inquire, delete
`-p <par_file>`	Name of the vendor supplied parameter file	
`-i <in_file>`	The name of the input file containing the parameters for the corresponding function (-f). If this option is not set, the input is read from `stdin`.	
`-o <out_file>`	The `backint` uses this name of the output file to write return values and messages. If this option is not set, the output is written to `stdout`.	
`-s <database_backup_id>`	The database backup ID which can be used to determine which calls to `backint` belong to a specific multi-file backup. This option is only usable if the function is backup.	
`-c <number_of_objects>`	The number of objects that belong to the database backup ID specified using the `-s` option. This option is only usable if the function is backup and in combination with option `-s`.	complete log
`-l <backup_level>`	The backup level indicates the type of backup. This option is only usable if the function is backup.	

Table A-2 `backint` command options *(continued)*

Parameter	Description	Values
`-v2`	This option requests the version to be printed to `stdout`. One line describing `backint` API version and `backint` implementation version. It should be the same string as #SOFTWAREID in the output file. For example: `backint 1.04 "ACME Corp backint 0815.2012.06`.	
`-V2`	Detailed version information, same as text as for option '`-v`' plus additional lines with detailed version and or licensing information.	

Note: The SAP HANA database does not rely on any action that is related to the parameters that are passed with options `-s`, `-c`, or `-l`. These are for informational purposes only.

The SAP HANA database never calls the options `-v` and `-V`. They are for diagnosis only. The implementation of `-V` is optional.

Appendix B

Input and output files for SAP HANA

This appendix includes the following topics:

About input and output files for SAP HANA

The input file contains information necessary to perform the requested function (-f). It usually contains a list of objects.

The output file contains result information for the database.

Each line must start with a keyword defined in the following sections. The lines that do not start with a keyword are considered as comments. These lines may contain any plain text. For example, to indicate progress or to support error analysis.

The character encoding of both files is UTF-8.

The formatting requirements for the parameters following keywords are listed:

- Parameters may be quoted with double quotation marks ".
- Parameters containing a space must be quoted with double quotation marks.
- If a parameter contains a double quote, the double quote must be escaped with a backslash.
- A backslash escapes double quotes only. A backslash in a parameter value must not be escaped.

The following are examples of valid input or output lines:

```
#PIPE /var/tmp/backup_Monday_cw47_2012

#PIPE "/var/tmp/backup Monday cw47 2012"
```

```
#EBID "Hb12\"NAB" "/var/tmp/myTool.#YSOwa"
```

The input files with multiple lines handling objects by pipe like `#PIPE` in function `BACKUP` or `#EBID` and `#NULL` in function `RESTORE` must be processed in sequential order starting from top to bottom. One `backint` may handle multiple pipes in parallel. However, it should not skip entries for later processing or pause processing of entries depending on events in subsequent entries.

The input and the output files are provided to `backint` by the SAP HANA database.

Appendix C

backint -i in_file contents

This appendix includes the following topics:

About backint -i in_file contents

This appendix explains how to create a `backint` input file. You specify the name of this input file as the argument to the `backint` command's `-i` *`in_file`* option. The input file consists of one or more directive lines that specify how to perform a backup, restore, or inquiry.

The input file format differs depending on the function you want `backint` to initiate and on the underlying database.

Table C-1 lists the directives that you can include in `in_file` and indicates the databases that use them.

Table C-1 Directives in the *in_file*

Directive	The underlying database
Directives for an input file that requests a backup:	
file	Oracle
`pipe #PIPE`	MaxDB
special_file size	Oracle
Directives for an input file that requests a restore:	
`BID file [dest_dir]`	Oracle

Table C-1 Directives in the *in_file (continued)*

Directive	The underlying database
`#NULL file [dest_dir]`	Oracle
`BID pipe [dest_name]`	MaxDB
`#NULL pipe [dest_name]`	MaxDB
Directives for an input file that performs an inquiry:	
`#NULL`	Oracle and MaxDB
BID	Oracle and MaxDB
`#NULL file`	Oracle
`#NULL pipe`	MaxDB
BID file	Oracle
BID pipe	MaxDB

Table C-2 lists the variables that you need to specify in the `in_file`:

Table C-2 Variables in the *in_file*

Variable	Specification
file	The full path name of the file that you want to back up, restore, or inquire about.
pipe	A named pipe. Specify the keyword `#PIPE` after the pipe name.
special_file	The name of a special file, such as a raw device file.
size	The file size as specified to the Oracle database.
BID	The external backup identifier, as generated by NetBackup for SAP.
dest_dir	The name of a directory or folder to which the restore is written.
dest_name	The name of a pipe to which the restore is written.

If you use the `#NULL` keyword, `backint` uses the most recent version of the file or pipe.

Appendix D

backint -o out_file contents

This appendix includes the following topics:

About backint -o out_file contents

When the `backint` command finishes, it writes an output file to the file that the `-o` *`out_file`* option specifies on its command line. The contents of the output text file change depending on the function `backint` performed.

The following table shows the `backint` output that is returned from a backup, restore, or inquiry request.

Table D-1 `backint` -o out_file contents

Output	Description
Information returned for a backup	If a backup completes successfully, the output file includes the following: ■ The external backup ID (BID) assigned to the backup by NetBackup for SAP. ■ The files, pipes, directories, or folders that are backed up and their sizes. If a backup function fails, the output file lists the files or pipes that it could not back up.

Table D-1 `backint` -o out_file contents *(continued)*

Output	Description
Information returned for a restore	If a restore completes successfully, the output file entry includes the following: ■ The BID. ■ The files or pipes restored. If a restore function fails, the output file lists the files and pipes not found. It also lists the files and pipes that were not successfully restored.
Information returned for an inquiry	If an inquiry completes successfully, the output file entry includes the following: ■ The BID. ■ The files or pipes that are backed up.

Appendix E

NetBackup for SAP environment variables for backint

This appendix includes the following topics:

- NetBackup for SAP backint interface environment variables

NetBackup for SAP backint interface environment variables

Table E-1 describes the environment variables that the NetBackup for SAP `backint` interface recognizes.

Note: These environment variables apply only to the `backint` interface and are not applicable to the phase 1 backup that SAP performs with RMAN. The `initSID.sap` settings are applicable to the phase 2 backup.

Table E-1 NetBackup for SAP `backint` interface environment variables

UNIX/Linux	Windows	Purpose
`$SAP_CLIENT`	`%SAP_CLIENT%`	Sets the name of the NetBackup client. It can be used to override the current client and perform an alternate restore of a different client. This variable is same as the `client` parameter in the `init`*`SID`*`.utl` file. This variable also overrides the `CLIENT_NAME` option in the registry or `bp.conf` file.

Table E-1 NetBackup for SAP `backint` interface environment variables *(continued)*

UNIX/Linux	Windows	Purpose
`$SAP_DRIVES`	`%SAP_DRIVES%`	Sets the number of simultaneous backups or restore operations the NetBackup for SAP `backint` interface can run for a single operation. This environment variable overrides the `drives` parameter in the `init`*`SID`*`.utl` file.
`$SAP_POLICY`	`%SAP_POLICY%`	Sets the name of the NetBackup policy. Used to define which policy to use for different types of database backup or inquiry functions. You can use one policy to perform offline database backups and another policy to perform archive log backups. This environment variable overrides as the `policy` parameter in the `init`*`SID`*`.utl` file. This variable also overrides the `BPBACKUP_POLICY` option in the registry or `bp.conf` file.
`$SAP_RECOVERY`	`%SAP_RECOVERY`	Set this environment variable to **disaster** when the disaster recovery is used with`BRRECOVER`. If the `init`*`SID`*`.sap` file is not present, `BRRECOVER` calls `BACKINT` without the "-p `init`*`SID`*`.utl`" parameter. If NetBackup for SAP is called without the `util` file parameter, it expects that the SAP_RECOVERY environment variable is set to **disaster**.
`$SAP_RESTORE`	`%SAP_RESTORE`	Set this environment variable to **rollback** to specify a snapshot rollback restore from BACKINT-based backups.
`$SAP_SCHED`	`%SAP_SCHED%`	Sets the name of the SAP backup policy schedule. This schedule provides an easy way to switch to a different schedule for each SAP database backup. This environment variable overrides the `schedule` parameter in the `init`*`SID`*`.utl` file. This variable also overrides the `BPBACKUP_SCHED` option in the registry or `bp.conf` file.
`$SAP_SNC_SCHED`	`%SAP_SNC_SCHED%`	Sets the name of the SAP backup policy schedule. This schedule is used to perform file-based snapshot and off-host backups of the data file in `brbackup` phase 1. The following types of schedules can be used for this type of backup: Automatic Full, Automatic Differential Incremental, or Automatic Cumulative Incremental.

Table E-1 NetBackup for SAP `backint` interface environment variables *(continued)*

UNIX/Linux	Windows	Purpose
`$SAP_SERVER`	`%SAP_SERVER%`	Sets the name of the NetBackup primary server. It can be used to override the current server and perform a backup to an alternative server. This environment variable overrides the `server` parameter in the `init`*`SID`*`.utl` file. This variable also overrides the `SERVER` option in the registry or `bp.conf` file.
`$SAPSWITCH`	`%SAPSWITCH%`	Set this environment variable to specify the location of the `.switch` files (`.switch.lis`, `.switch.sem`, `.switch.log`). This variable gets higher precedence than the switch parameters (`switch_list`, `switch_sem`, `switch_log`) mentioned in `int`*`SID`*`.utl` file.

Appendix F

NetBackup for SAP configuration or bp.conf file settings

This appendix includes the following topics:

NetBackup for SAP configuration or bp.conf file settings

This topic describes the NetBackup for SAP configuration parameters. You can specify parameters in several places.

Table F-1 describes the order of precedence, in the case of conflicts.

Table F-1 Order of precedence for NetBackup for SAP parameters

Setting	Precedence
Environment variables	Highest precedence. Overrides the values in the `init`*`SID`*`.utl` and in the `bp.conf` file or registry.
`init`*`SID`*`.utl` file	Medium precedence. Values in environment variables override the values that are specified in this file, but values in this file override those in the `bp.conf` file or registry.

Table F-1 Order of precedence for NetBackup for SAP parameters *(continued)*

Setting	Precedence
(UNIX or Linux) `/usr/openv/netbackup/bp.conf` file	Lowest precedence. Values in environment variables and in `init`*`SID`*`.utl` override values that are specified in the `bp.conf` file.
(Windows) NetBackup registry	Lowest precedence. Values in environment variables and in `init`*`SID`*`.utl` override values that are specified in the registry.

Table F-2 lists the variable names and definitions you can use in the `bp.conf` file (UNIX or Linux) or the registry (Windows).

Table F-2 NetBackup for SAP variable names and definitions

UNIX or Linux	Windows	Meaning
`$SERVER`	`%SERVER%`	Specifies the NetBackup primary server. This option is the same as the `$SAP_SERVER` (`%SAP_SERVER%`) environment variable and the `server` parameter in the `init`*`SID`*`.utl` file.
`$CLIENT_NAME`	`%CLIENT_NAME%`	Specifies the NetBackup client. This option is the same as the `$SAP_CLIENT` (`%SAP_CLIENT%`) environment variable and the `client` parameter in the `init`*`SID`*`.utl` file.
`$BPBACKUP_POLICY`	`%BPBACKUP_POLICY%`	Specifies the NetBackup policy. This option is the same as the `$SAP_POLICY` (`%SAP_POLICY%`) environment variable and the `policy` parameter in the `init`*`SID`*`.utl` file.

Table F-2 NetBackup for SAP variable names and definitions *(continued)*

UNIX or Linux	Windows	Meaning
`$BPBACKUP_SCHED`	`%BPBACKUP_SCHED%`	Specifies the NetBackup schedule. This option is the same as the `$SAP_SCHED` (`%SAP_SCHED%`) environment variable and the `schedule` parameter in the `init`*`SID`*`.utl` file.

For more information, see the NetBackup Administrator's Guide, Volume II.

Appendix G

Parameters used in initSID.utl

This appendix includes the following topics:

- sort_backup_type <value>
- sort_restore_type <value>
- switch_list <control_file_path>
- switch_log <control_file_path>
- switch_sem <control_file_path>

About parameters used in initSID.utl

This appendix shows the parameters you can specify in the `init`*`SID`*`.utl` configuration file. The SAP tools pass these parameters to the `backint` interface through its `-p` *`parfile`* option.

Unless otherwise noted, the parameters in this appendix section apply to the following environments:

- NetBackup for SAP on Oracle databases without RMAN
- NetBackup for SAP on Oracle databases with RMAN
- NetBackup for SAP on MaxDB databases

Where necessary, the parameter descriptions indicate whether they apply to only one or two of these environments.

initSID.utl parameter summary

Table G-1 summarizes the parameters and the database environments that support them.

Table G-1 NetBackup for SAP parameters and the database environments

Parameter name	Underlying database and usage notes
`backup_stream_buffersize`	MaxDB See "backup_stream_buffersize <size>" on page 229.
`client`	See "client <client_name>" on page 229.
`custom_sort_file`	Oracle without RMAN See "custom_sort_file <file_path>" on page 229.

Table G-1 NetBackup for SAP parameters and the database environments *(continued)*

Parameter name	Underlying database and usage notes
`drives`	Oracle without RMAN See "drives <number_of_drives>" on page 233.
`inquire_query_period`	See "inquiry_query_period <months>" on page 233.
`master_time_offset` (UNIX or Linux)	See "master_time_offset <minutes> (UNIX or Linux)" on page 234.
`multistream_restore`	Oracle without RMAN See "multistream_restore" on page 234.
`policy`	See "policy <policy_name>" on page 234.
`policy2`	See "policy2 <policy_name>" on page 235.
`restore_stream_buffersize`	MaxDB See "restore_stream_buffersize <size>" on page 235.
`retry_backup`	See "retry_backup <number_of_retries>" on page 235.
`schedule`	See "schedule <schedule_name>" on page 235.
`schedule2`	See "schedule2 <schedule_name>" on page 236.
`server`	See "server <server_name>" on page 236.
`sort_backup_type`	Oracle without RMAN See "sort_backup_type <value>" on page 236.
`sort_restore_type`	Oracle without RMAN See "sort_restore_type <value>" on page 239.
`switch_list`	Oracle without RMAN See "switch_list <control_file_path>" on page 241.

Table G-1 NetBackup for SAP parameters and the database environments *(continued)*

Parameter name	Underlying database and usage notes
switch_log	Oracle without RMAN See "switch_log <control_file_path>" on page 241.
switch_sem	Oracle without RMAN See "switch_sem <control_file_path>" on page 241.

backup_stream_buffersize <size>

This parameter specifies the buffer size, in bytes, for stream-based backups. NetBackup receives data from MaxDB through the stream (pipe). When MaxDB passes the data to NetBackup, it uses this buffer size. Also see the `restore_stream_buffersize` parameter.

client <client_name>

This parameter specifies the host name by which the primary server knows this SAP client, as entered into the policy. In some cases the server and the client host are same hosts. The following is an example:

```
client saturn
```

If the NetBackup for SAP `backint` interface finds the `$SAP_CLIENT` (`%SAP_CLIENT%`) environment variable, the `$SAP_CLIENT` environment variable value overrides the `client` parameter value.

If the `client` parameter is not specified, and there is no environment variable, the `client` parameter defaults to the value that is specified for the `CLIENT_NAME` option in the `bp.conf` file or the registry. If the value is not specified there, the NetBackup for SAP `backint` interface uses the value the `gethostname()` library function returns.

custom_sort_file <file_path>

Specify this parameter only when the `sort_backup_type` or the `sort_restore_type` parameters are set to `custom`.

When `custom` is specified, set the `custom_sort_file` parameter to a valid file. The *file_path* value must be a full path name to a custom sort file, and it must have public permissions. This file can be created by end user manually or can be created automatically by using the `nbsapcustomsort` utility.

Linux:

```
/usr/openv/netbackup/bin/nbsapcustomsort
```

Windows:

```
install_path\NetBackup\bin\nbsapcustomsort.exe
```

Note: For `nbsapcustomsort the brspace` utility need to be available in BRTOOLS in the environment.

The following is an example argument to the `custom_sort_file` parameter:

UNIX or Linux:

```
<custom_sort_file_path>/sap_custom_sort_file
```

Windows:

```
<custom_sort_file_path>/sap_custom_sort_file
```

Note: Not applicable in MaxDB database environments or in any Oracle database environments that use RMAN.

The custom sort file must include two fields. The first field groups a set of files into a particular job. The second field is a file path name that maps the SAP backup file list to a group ID.

The following is an example of a custom sort file for Windows:

```
1 c:\oracle\sap\sapdata1\btabd_1\btabd.data1
1 c:\oracle\sap\sapdata2\btabi_1\btabi.data1
1 c:\oracle\sap\sapdata2\clud_1\clud.data1
1 c:\oracle\sap\sapdata1\ddicd_1\ddicd.data1
1 c:\oracle\sap\sapdata5\ddici_1\ddici.data1
1 c:\oracle\sap\sapdata4\el30cd_1\EL30cd.data1
1 c:\oracle\sap\sapdata1\el30ci_1\el30ci.data1
1 c:\oracle\sap\sapdata6\es30cd_1\es30cd.data1
1 c:\oracle\sap\sapdata2\poold_1\poold.data1
1 c:\oracle\sap\sapdata1\pooli_1\pooli.data1
1 c:\oracle\sap\sapdata4\protd_1\protd.data1
```

```
2 c:\oracle\sap\sapdata1\roll_1\roll.data1
2 c:\oracle\sap\sapdata2\sourced_1\sourced.data1
2 c:\oracle\sap\sapdata3\stabd_1\stabd.data1
2 c:\oracle\sap\sapdata2\stabi_2\stabi.data2
2 c:\oracle\sap\sapdata1\temp_1\temp.data1
2 c:\oracle\sap\sapdata4\user1d_1\user1d.data1
2 c:\oracle\sap\sapdata2\user1i_1\user1i.data1
2 c:\oracle\sap\sapdata1\system_1\system.data1
2 c:\oracle\sap\saplog1\log_g1_m1\log1_m1.dbf
2 c:\oracle\sap\saplog1\log_g2_m1\log2_m1.dbf
2 c:\oracle\sap\saplog1\log_g3_m1\log3_m1.dbf
2 c:\oracle\sap\saplog1\log_g4_m1\log4_m1.dbf
2 c:\oracle\sap\dbs\cntrlSAP.dbf
```

The following is an example of a custom sort file for UNIX or Linux:

```
1 /oracle/sap/sapdata1/btabd_1/btabd.data1
1 /oracle/sap/sapdata2/btabi_1/btabi.data1
1 /oracle/sap/sapdata2/clud_1/clud.data1
1 /oracle/sap/sapdata1/ddicd_1/ddicd.data1
1 /oracle/sap/sapdata5/ddici_1/ddici.data1
1 /oracle/sap/sapdata4/el30cd_1/EL30cd.data1
1 /oracle/sap/sapdata1/el30ci_1/el30ci.data1
1 /oracle/sap/sapdata6/es30cd_1/es30cd.data1
1 /oracle/sap/sapdata2/poold_1/poold.data1
1 /oracle/sap/sapdata1/pooli_1/pooli.data1
1 /oracle/sap/sapdata4/protd_1/protd.data1

1 /dev/rdsk/c0t4d0s6

2 /oracle/sap/sapdata1/roll_1/roll.data1
2 /oracle/sap/sapdata2/sourced_1/sourced.data1
2 /oracle/sap/sapdata3/stabd_1/stabd.data1
2 /oracle/sap/sapdata2/stabi_2/stabi.data2
2 /oracle/sap/sapdata1/temp_1/temp.data1
2 /oracle/sap/sapdata4/user1d_1/user1d.data1
2 /oracle/sap/sapdata2/user1i_1/user1i.data1
2 /oracle/sap/sapdata1/system_1/system.data1
2 /oracle/sap/saplog1/log_g1_m1/log1_m1.dbf
2 /oracle/sap/saplog1/log_g2_m1/log2_m1.dbf
2 /oracle/sap/saplog1/log_g3_m1/log3_m1.dbf
2 /oracle/sap/saplog1/log_g4_m1/log4_m1.dbf
2 /oracle/sap/dbs/cntrlSAP.dbf
```

Based on the custom sort file, if SAP tools submit the entire file list to `backint` for backup, two jobs are created. The first job includes all the files that have a “1” in the first field. The second job includes all the files that have a “2” in the first field.

The following is a list of jobs and associated files:

- Backup job and restore job 1
  Windows:

```
c:\oracle\sap\sapdata1\btabd_1\btabd.data1
c:\oracle\sap\sapdata2\btabi_1\btabi.data1
c:\oracle\sap\sapdata2\clud_1\clud.data1
c:\oracle\sap\sapdata1\ddicd_1\ddicd.data1
c:\oracle\sap\sapdata5\ddici_1\ddici.data1
c:\oracle\sap\sapdata4\el30cd_1\EL30cd.data1
c:\oracle\sap\sapdata1\el30ci_1\el30ci.data1
c:\oracle\sap\sapdata6\es30cd_1\es30cd.data1
c:\oracle\sap\sapdata2\poold_1\poold.data1
c:\oracle\sap\sapdata1\pooli_1\pooli.data1
c:\oracle\sap\sapdata4\protd_1\protd.data1
c:\dev\rdsk\c0t4d0s6
```

  UNIX or Linux:

```
/oracle/sap/sapdata1/btabd_1/btabd.data1
/oracle/sap/sapdata2/btabi_1/btabi.data1
/oracle/sap/sapdata2/clud_1/clud.data1
/oracle/sap/sapdata1/ddicd_1/ddicd.data1
/oracle/sap/sapdata5/ddici_1/ddici.data1
/oracle/sap/sapdata4/el30cd_1/EL30cd.data1
/oracle/sap/sapdata1/el30ci_1/el30ci.data1
/oracle/sap/sapdata6/es30cd_1/es30cd.data1
/oracle/sap/sapdata2/poold_1/poold.data1
/oracle/sap/sapdata1/pooli_1/pooli.data1
/oracle/sap/sapdata4/protd_1/protd.data1/dev/rdsk/c0t4d0s6
```

- Backup job and restore job 2:
  Windows:

```
c:\oracle\sap\sapdata1\roll_1\roll.data1
c:\oracle\sap\sapdata2\sourced_1\sourced.data1
c:\oracle\sap\sapdata3\stabd_1\stabd.data1
c:\oracle\sap\sapdata2\stabi_2\stabi.data2
c:\oracle\sap\sapdata1\temp_1\temp.data1
c:\oracle\sap\sapdata4\user1d_1\user1d.data1
```

```
c:\oracle\sap\sapdata2\user1i_1\user1i.data1
c:\oracle\sap\sapdata1\system_1\system.data1
c:\oracle\sap\saplog1\log_g1_m1\log1_m1.dbf
c:\oracle\sap\saplog1\log_g2_m1\log2_m1.dbf
c:\oracle\sap\saplog1\log_g3_m1\log3_m1.dbf
c:\oracle\sap\saplog1\log_g4_m1\log4_m1.dbf
c:\oracle\sap\dbs\cntrlSAP.dbf
```

UNIX or Linux:

```
/oracle/sap/sapdata1/roll_1/roll.data1
/oracle/sap/sapdata2/sourced_1/sourced.data1
/oracle/sap/sapdata3/stabd_1/stabd.data1
/oracle/sap/sapdata2/stabi_2/stabi.data2
/oracle/sap/sapdata1/temp_1/temp.data1
/oracle/sap/sapdata4/user1d_1/user1d.data1
/oracle/sap/sapdata2/user1i_1/user1i.data1
/oracle/sap/sapdata1/system_1/system.data1
/oracle/sap/saplog1/log_g1_m1/log1_m1.dbf
/oracle/sap/saplog1/log_g2_m1/log2_m1.dbf
/oracle/sap/saplog1/log_g3_m1/log3_m1.dbf
/oracle/sap/saplog1/log_g4_m1/log4_m1.dbf
/oracle/sap/dbs/cntrlSAP.dbf
```

drives <number_of_drives>

This parameter specifies the number of jobs into which `backint` should divide the file list when `backup_sort_type` is drive or size.

Note: Not applicable in MaxDB database environments or in any Oracle database environments that use RMAN.

The following is an example entry:

```
drives 5
```

inquiry_query_period <months>

In the case of `#NULL` based inquiry queries, NetBackup for SAP restricts the catalog search to the last 6 months by default. You can use this parameter to specify a

different query period. Specify an integer value for *months*. For example, if you want to set the query to the last 10 months, set this parameter as follows:

```
inquiry_query_period 10
```

master_time_offset <minutes> (UNIX or Linux)

This parameter restores old backups if there was a time difference between the primary server and the client hosts.

Use this option only when the date and times are out of sync between the server and the client hosts.

The parameter value that is specified in minutes is subtracted from the start time and added to the end time for a restore or inquire. The following is an example of an entry:

```
master_time_offset 3
```

multistream_restore

This parameter is optional. The parameter can be set to 0 or 1. Specify 0 so that NetBackup primary server determines the most efficient way to create restore jobs from the selected backup images. All these run as one job. Specify 1 to create restore streams based on the values of `drives` and `sort_restore_type`. Each stream then becomes a separate job.

Note: Some sites find it easier to monitor and restart failed streams for large restores if they are separate jobs.

A value of 0 is the default setting.

```
multistream_restore 0
```

Note: Not applicable in MaxDB database environments or in any Oracle database environments that use RMAN.

policy <policy_name>

This parameter specifies the name of an SAP policy that is defined in NetBackup. The SAP policy must have an Application Backup schedule defined in order for the NetBackup for SAP `backint` interface to work. The following is an example entry:

```
policy sap_nb
```

policy2 <policy_name>

This parameter specifies the name of a policy to be used for the secondary SAP (phase 2) backup. A secondary backup is performed for each SAP database backup on any files that are needed to track SAP backup information. You can use this option to save the backup information on a different media. If `policy2` is not specified and `schedule2` is specified, NetBackup uses the value that is specified for the `policy` parameter.

restore_stream_buffersize <size>

This parameter specifies the buffer size, in bytes, for stream-based restores. MaxDB receives data from NetBackup through the stream (pipe). When MaxDB receives the data from NetBackup and copies it to the pipe, it uses this buffer size. For more information, see the `backup_stream_buffersize` parameter.

retry_backup <number_of_retries>

This parameter specifies the number of retries for a failed backup. `backint` retries a failed job the specified number of times. For example:

```
retry_backup 2
```

schedule <schedule_name>

This parameter specifies the name of an Application Backup schedule that is associated with an SAP policy type. The schedule can define aspects of the backup such as how long NetBackup retains images, maximum multiplexing per drive, storage unit, and volume pool. The following is an example entry:

```
schedule sap_full_backup
```

If the NetBackup for SAP `backint` interface finds the `$SAP_SCHED` (`%SAP_SCHED%`)environment variable, the `$SAP_SCHED` (`%SAP_SCHED%`) environment variable value overrides the `schedule` parameter value. If the `schedule` parameter is not specified and there is no environment variable, then the `schedule` parameter value defaults to the `BPBACKUP_SCHED` option in the NetBackup `bp.conf` file or the NetBackup registry. If `BPBACKUP_SCHED` is not specified there, NetBackup uses the first Application Backup schedule it finds in the SAP policy.

schedule2 <schedule_name>

This parameter specifies name of an Application Backup schedule to be used for the secondary SAP (phase 2) backup. If it is not specified, the `schedule` parameter value is used.

For each SAP backup, NetBackup performs two individual backups. The first backup backs up database data. The second backup backs up the log files that are needed to track SAP backup information. You can use this parameter to save SAP log files to a different media. This can make it easier to restore a database. You can also use this option to save the backup information to a different volume pool. The following is an example entry:

```
schedule2  sap_backup_information
```

See "About parameters used in initSID.utl" on page 227.

See "initSID.utl parameter summary" on page 227.

server <server_name>

This parameter specifies the network routable hostname of the NetBackup primary server. The server provides most of the administration and control for NetBackup operations. It contains the NetBackup database. The following is an example entry:

```
server jupiter
```

sort_backup_type <value>

This parameter specifies one of four different backup sort parameter values: `custom`, `device`, `drive`, or `size` (default).

Note: Not applicable in MaxDB database environments or in any Oracle database environments that use RMAN.

See "sort_backup_type custom" on page 237.

See "sort_backup_type device" on page 237.

See "sort_backup_type drive" on page 237.

See "sort_backup_type size (default)" on page 238.

sort_backup_type custom

Specifies that a customized sort file be used.

If `sort_backup_type custom` is specified, also specify the `custom_sort_file` parameter with a file path.

See "custom_sort_file <file_path>" on page 229.

sort_backup_type device

Specifies that NetBackup create jobs based on a file's device ID. This causes the `drives` parameter in the `int`*`SID`*`.utl` file to be ignored.

For example, if 12 files are requested for backup, and they reside on two different devices (`X` and `Y`), then two jobs are created. The first job contains all the files that are associated with device `X`, and the second job contains all the files on device `Y`.

Note: The implementation is based on the *st_dev* value from the `stat()` function. This value identifies a file partition.

The following UNIX or Linux example shows how jobs are created when device is specified:

- Input file list from SAP (`brbackup`, `sapdba`):
- Backup job 1 (all have the device ID x):
- Backup job 2 (all have the same device ID y):

The following Windows example shows how jobs are created when device is specified:

- Input file list from SAP (`brbackup`, `sapdba`):
- Backup job 1 (all have the device ID x):
- Backup job 2 (all have the same device ID y):

sort_backup_type drive

Specifies that backint should distribute the input file list in a cyclic fashion. The distribution will be amongst a number of jobs equal to the drives parameter specified in the `initSID.utl` file

For example, on UNIX or Linux, if there are three tape drives and 10 SAP files, the following distribution occurs:

- Input file list from SAP:

- Backup job and restore job 1:
- Backup job and restore job 2:
- Backup job and restore job 3:

For example, on Windows, if there are three tape drives and 10 SAP files, the following distribution occurs:

- Input file list from SAP:
- Backup job and restore job 1:
- Backup job and restore job 2:
- Backup job and restore job 3:

sort_backup_type size (default)

This value is the default for the `sort_backup_type` parameter.

If `sort_backup_type size` is specified, backint sorts the files by size before distributing them amongst the jobs based upon the drives that are specified in the `init`*`SID`*`.utl` file.

The number of drives that are specified does not have to equal the number of physical storage units. The number of drives correlates to the number of simultaneous jobs created by the NetBackup for SAP `backint` interface. For example, if you have 10 SAP files and three tape drives, you can specify 10 drives in the `init`*`SID`*`.utl` file. This value creates 10 jobs with one file for each `bpbackup` job. NetBackup handles all of the job scheduling. Initially, three jobs are active and the other seven jobs are queued.

Note: Not applicable in MaxDB database environments or in any Oracle database environments that use RMAN.

For example, if three tape drives are specified, the files are divided evenly into three `bpbackup` jobs based on size. So, if there are 25 input files from SAP and three tape drives, then three `bpbackup` jobs run at the same time.

The following files are in each job, on UNIX or Linux:

- Input file list from SAP (`brbackup`, `sapdba`):
- Backup job 1:
- Backup job 2:
- Backup job 3:

The following files are in each job, on Windows:

- Input file list from SAP (`brbackup`, `sapdba`):
- Backup job 1:
- Backup job 2:
- Backup job 3:

sort_restore_type <value>

To use this parameter, the `multistream_restore` parameter must be set to `1`. NetBackup can use the `sort_restore_type` parameter to restore the backup images.

`sort_restore_type` specifies one of three different restore sort options: `custom`, `drive`, or `image` (default).

The following sections describe the possible parameter values.

Note: Not applicable in MaxDB database environments or in any Oracle database environments that use RMAN.

---

sort_restore_type custom

Specifies that a customized sort file be used.

If `sort_restore_type custom` is specified, also specify the `custom_sort_file` parameter with a file path.

See “custom_sort_file <file_path>” on page 229.

sort_restore_type drive

Specifies that NetBackup create jobs based on the `drives` parameter in the `init`*`SID`*`.utl` file. The following examples show the distribution if there are 3 tape drives and 10 SAP files.

UNIX or Linux:

- Input file list from SAP:
- Backup job and restore job 1:
- Backup job and restore job 2:
- Backup job and restore job 3:

Windows:

- Input file list from SAP:
- Backup job and restore job 1:
- Backup job and restore job 2:
- Backup job and restore job 3:

sort_restore_type image (default)

Specifies that `backint` create groups of files based on their backup image IDs and create a job for each group. This behavior is the default if the `sort_restore_type` parameter is not set in the `int`*`SID`*`.utl` file.

For example, if nine files were backed up by two `bpbackup` jobs, each file is associated with one of two backup image IDs. If all nine files are restored, the NetBackup for SAP `backint` interface creates the two jobs, one job for each image. The files are grouped the way they were backed up. The following are examples of a restore.

Note: Restore creates a separate job for raw partition files. The partition files cannot be grouped with regular files.

UNIX or Linux:

Input file list from SAP (`brrestore`, `sapdba`):

- Image 1:
- Image 2:
- Restore job 1:
- Restore job 2:

Windows:

Input file list from SAP (`brrestore`, `sapdba`):

- Image 1:
- Image 2:
- Restore job 1:
- Restore job 2:

switch_list <control_file_path>

This parameter specifies a control file that communicates from the NetBackup for SAP `backint` interface to `brbackup` for online backups. A switch list file is created each time `backint` wants to take a snapshot, or back up a file. The switch list file is also created when it wants to indicate that a snapshot or backup is finished.

Note: Not applicable in any Oracle database environments that use RMAN.

The `switch_list` parameter must specify a file path in one of the following locations:

The following are examples of valid entries:

switch_log <control_file_path>

This parameter specifies a control file that communicates to the NetBackup for SAP `backint` interface from `brbackup`. After the switch semaphore file has been deleted, the NetBackup for SAP `backint` interface opens and reads the switch log file. to determine if `brbackup` was successful at making the file ready for snapshot or backup.

Note: Not applicable in any Oracle database environments that use RMAN.

The `switch_log` parameter must specify a file path in one of the following locations:

The following are examples of valid entries:

switch_sem <control_file_path>

This parameter specifies a control file that communicates between the NetBackup for SAP `backint` interface and `brbackup`. After the switch list file is created and closed, the NetBackup for SAP `backint` interface creates the switch semaphore file and waits until `brbackup` deletes it.

Note: Not applicable in any Oracle database environments that use RMAN.

The `switch_sem` parameter must specify a file path in the following location:

The following is an example of a valid entry:

Appendix H

Configuring split mirror backups

This appendix includes the following topics:

About configuring split mirror backups

The following snapshot methods are available for configuring split mirror backups for NetBackup for SAP.

- Local-host
  See "Local-host snapshot method: nbu_snap (Solaris SPARC platform only)" on page 243.
  See "Local-host snapshot method: VxVM (Solaris SPARC, HP-UX, Windows)" on page 244.
  See "Local-host snapshot method: VxFS_Checkpoint (UNIX or Linux)" on page 245.
- Off-host

See "Off-host alternate client, FlashSnap method" on page 246.
See "Off-host alternate client, VVR method (Solaris SPARC, HP-UX)" on page 248.

- Hardware array-based (UNIX or Linux)
  See "Hardware array-based snapshot method: Hitachi Shadow Image, EMC TimeFinder, HP Business Copy (UNIX or Linux)" on page 251.

Some of the configuration methods are platform-specific.

Local-host snapshot method: nbu_snap (Solaris SPARC platform only)

The `nbu_snap` snapshot method is for Solaris (SPARC) clients only. It makes copy-on-write snapshots for any SAP Oracle databases that reside in UFS or Veritas File Systems (VxFS).

For the `nbu_snap` methods, you must identify a cache device for the copy-on-write process. The cache device is a raw disk partition, either a logical volume or physical disk. It stores the portions of the client's data that the incoming write requests change while the copy-on-write is in progress.

For the cache device, do not select an active partition that contains valuable data. Any data in that partition is lost when the snapshot is complete. Specify the raw partition as the full path name of either the character special device file or the block device file.

For example, specify the following:

- Solaris (SPARC) raw partition: `/dev/rdsk/c2t0d3s3` or `/dev/dsk/c2t0d3s3`
- VxVM volume: `/dev/vx/rdsk/diskgroup_1/volume_3` or `/dev/vx/dsk/diskgroup_1/volume_3`
  Do not specify wildcards (such as `/dev/rdsk/c2*`) as paths.

The cache partition must be unmounted and must reside on the same host as the snapshot source (the client's data to back up). The partition must have enough space to hold all the writes to the partition that can occur during the backup. Note that backups during off-peak periods normally require a smaller cache than a backup during peak activity.

For more information, see the NetBackup Snapshot Client Administrator's Guide.

Configuring NetBackup for SAP with the local-host snapshot method (nbu_snap)

1. On the production database host, install NetBackup Client, which includes the NetBackup Snapshot Client, and NetBackup for SAP.
2. Identify a cache device for the copy-on-write process. Manually select the snapshot method on the **Advanced Snapshot Options** dialog box.
3. Specify the raw cache partition using one of the following methods:
   - Specify the raw partition in the default cache device path for snapshots field. (Open **NetBackup Management > Host Properties > Clients**. Then in the **Client Properties** dialog box, expand **UNIX Client > Client Settings**). This setting applies to the client in all policies.
   - In the **Advanced Snapshot Options** dialog box, specify the cache device path value field. This cache setting applies to all clients in the current policy and overrides the cache setting in the client settings dialog box.

Local-host snapshot method: VxVM (Solaris SPARC, HP-UX, Windows)

Use the following procedure to configure.

To configure the local host snapshot method

1. Install and configure the following on the production database host:
   - Install NetBackup Client, NetBackup Snapshot Client, and NetBackup for SAP Agent.
   - Configure the production database (Oracle).
   - You can install NetBackup primary server on a separate host.
2. Attach the external disk D1 to the primary client (production database host) and create a VM disk group using the following command:

```
root@primary# vxdg init diskgroup [cds=on|off] diskname=devicename
```

3. Create a volume in primary disk D using the following command:

```
root@primary# vxassist -g diskgroup make volume size
```

4. Add a DCO log to the volume using the following command:

```
root@primary# vxassist -g diskgroup addlog volume logtype=dco
```

5 Enable FastResync on the volume using the following command:

```
root@primary# vxvol -g diskgroup set fastresync=on volume
```

6 Create a VxFS file system on this volume and mount it using the following commands:

```
root@primary# mkfs -F vxfs /dev/vx/rdsk/diskgroup/volume
root@primary# mount -F vxfs /dev/vx/dsk/diskgroup/volume mnt
```

7 Create the SAP production database on this volume.

8 Start the snapshot for the data volume in the primary client using the following command:

```
root@primary# vxassist -g diskgroup -b snapstart volume
```

Local-host snapshot method: VxFS_Checkpoint (UNIX or Linux)

Configure the Oracle database in the Veritas File System (VxFS) in the production host. The same requirements apply to the VxFS_Checkpoint method as for BLI backups.

See “About configuring NetBackup for SAP block-level incremental backups on UNIX” on page 177.

Off-host alternate client, FlashSnap method

Figure H-1 Typical SAP split mirror off-host backup with FlashSnap

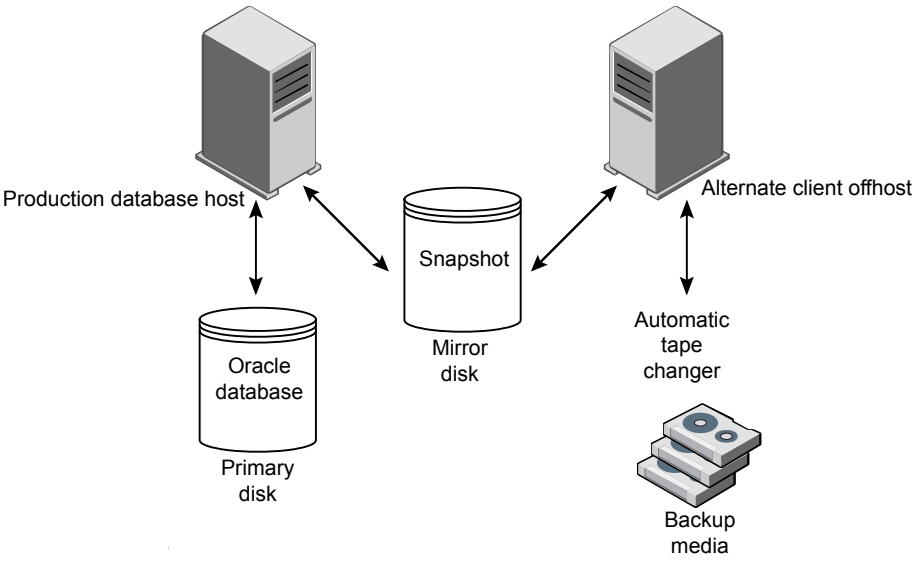

Based on this FlashSnap configuration, the typical flow during the Snapshot Client backup for SAP is as follows:

- The database shuts down or the tablespaces are put in backup mode in the production host.
- The alternate client off-host takes the snapshot of the production database from primary disk D1 to mirror disk D2.
- The production database in primary disk D1 is created online.
- The mirror disk D2 is split to perform backup from the snapshot to the backup media.
- The snapshot from the mirror disk D2 is then joined back to the volume in primary disk D1.

The user identification and group identification numbers (UIDs and GIDs) associated with the files to be backed up must be available to the primary client and the alternate backup client. The UID on the primary client and the alternate backup client must be the same. Similarly, the GID on the primary client and the alternate backup client must be the same.

To configure the local host snapshot method

1 Install and configure the following on the production database host:

- Install NetBackup Client, NetBackup Snapshot Client, and the NetBackup for SAP Agent.
- Configure the production database (Oracle).
- You can install the NetBackup primary server on a separate host.

2 Install the following on the alternate client (off-host):

- Install NetBackup Client and NetBackup Snapshot Client. The NetBackup for SAP Agent is not required.
- The NetBackup primary server. You can install the primary server on the same host or on a separate host where you want to attach the backup media.
- The UID number can be different than the GID number.

3 Attach the external disk D1 to the primary client (production database host) and create a VM disk group using the following command:

```
root@primary# vxdg init diskgroup [cds=on|off] diskname=devicename
```

4 Create a volume in primary disk D using the following command:

```
root@primary# vxassist -g diskgroup make volume size
```

5 Add a DCO log to the volume using the following command:

```
root@primary# vxassist -g diskgroup addlog volume logtype=dco
```

6 Enable FastResync on the volume using the following command:

```
root@primary# vxvol -g diskgroup set fastresync=on volume
```

7 Create a VxFS file system on this volume and mount it using the following commands:

```
root@primary# mkfs -F vxfs /dev/vx/rdsk/diskgroup/volume
root@primary# mount -F vxfs /dev/vx/dsk/diskgroup/volume mnt
```

8 Create the SAP production database on this volume.

9 Connect the external mirror disk D2 to the primary client and the alternate client, so the disk is accessible to both hosts.

10 Start the snapshot for the data volume in the primary client using the following command:

```
root@primary# vxassist -g diskgroup -b snapstart volume
```

Off-host alternate client, VVR method (Solaris SPARC, HP-UX)

Figure H-2 Typical SAP split mirror off-host backup with VVR

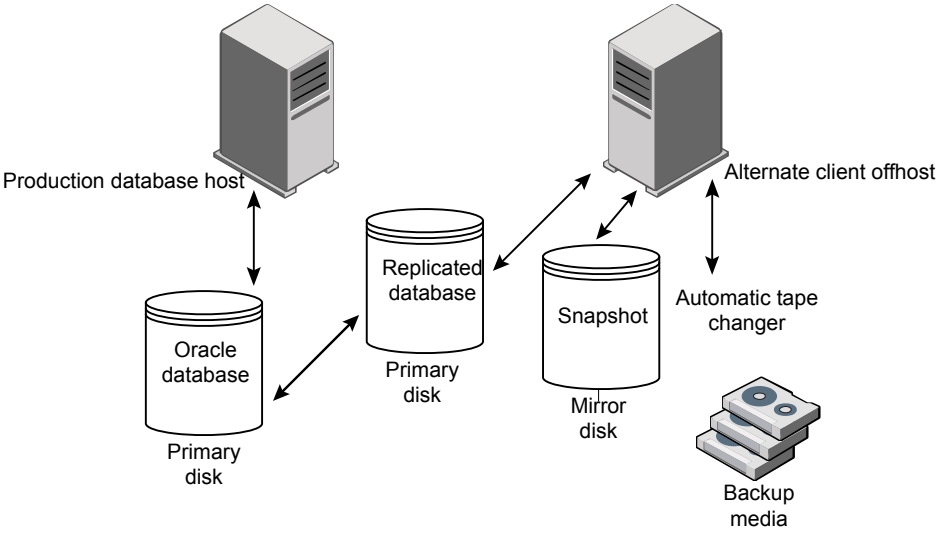

Based on this VVR configuration, the typical flow during the Snapshot Client backup for SAP:

- The database shuts down or tablespaces are put in backup mode in the production host.
- Replication is stopped between primary disk D1 and replication disk D2.
- The alternate client off-host takes the snapshot of the replication database from disk D2 to mirror disk D3.
- The production database in primary disk D1 is created online.
- The mirror disk D3 is split to perform backup from the snapshot to the backup media.
- The snapshot from the mirror disk D3 is then joined back to the volume in replication disk D2.
- Then the storage replication log (SRL) resync happens in primary disk D1 and secondary disk D2.
- Finally, replication resumes.

Before configuring this method, install and configure the following:

- Install and configure the following on the production database host:
  - Install NetBackup Client, NetBackup Snapshot Client, and the NetBackup for SAP Agent
  - Configure the production database (Oracle).
  - You can install the NetBackup primary server on a separate host
- Install the following on the alternate client (off-host):
  - Install NetBackup Client and NetBackup Snapshot Client. The NetBackup for SAP Agent is not required.
  - The NetBackup media server. You can install the media server on the same host or on separate host where you want to attach the backup media.

The following procedures describe how to configure both the volume replicator primary and the volume replicator secondary.

To configure the volume replicator primary (production database host)

1. Attach the external primary disk D1 to the volume replicator primary, and create a VM disk group using the following command:

```
root@primary# vxdg init diskgroup [cds=on|off] diskname=devicename
```

2. Create a data volume in primary disk D1 using the following command:

```
root@primary# vxassist -g diskgroup make volume size
```

By default, it creates two copies that require two disks in one group. To override the default, enter the following command:

```
root@primary # vxassist -g diskgroup addlog volume logtype=dcm nlog=1
```

3. Create the VxFS file system on this volume and mount it using the following commands:

```
root@primary# mkfs -F vxfs /dev/vx/rdsk/diskgroup/volume
root@primary# mount -F vxfs /dev/vx/dsk/diskgroup/volume mnt
```

4. Create the SAP production database on this volume.
5. Create a Storage Replicator Log (SRL) volume in primary disk D1 using the following command:

```
root@primary# vxassist -g diskgroup make volume_srl size
```

To configure the volume replicator secondary (alternate client off-host)

1. Attach the external replication disk D2 to the volume replicator secondary and create the VM disk group using the following command:

```
root@secondary# vxdg init diskgroup [cds=on|off] diskname=devicename
```

2. Create a data volume in replication disk D2 using the following command:

```
root@secondary # vxassist -g diskgroup make volume size
```

   By default, it creates two copies that require two disks in one group. To override the default, enter the following command:

```
root@secondary# vxassist -g diskgroup addlog volume logtype=dcm nlog=1
```

3. Create a Storage Replicator Log (SRL) volume in the replication disk D2 using the following command:

```
root@secondary# vxassist -g diskgroup make volume_srl size
```

4. Create the primary replicated volume group (RVG) specifying the data volume(s) and the SRL volume. On the volume replicator primary (production database host), run the following command:

   The names of the disk group and volumes must be same as on the primary host.

```
root@primary# vradmin -g diskgroup createpri rvg_name data_volumes(s) srl_volume
```

5. Ensure that authentication can complete.

   Before you add a new secondary host (by using the `addsec`) or overwrite data on a remote host (by using the `syncvol`), `vradmin` performs some authentication. This process checks that the `/etc/vx/vras/.rdg` file on the remote host contains an entry for the primary disk group ID. The `vradmin addsec` or `syncvol` command fails if the `/etc/vx/vras/.rdg` file on the remote host does not have such an entry.

   If authentication cannot complete, do the following:

   - Add the primary disk group ID to the `/etc/vx/vras/.rdg` file on the remote host.
   - To find disk group ID, run the `vxprint -l diskgroup_name` command on the primary.

- Ensure that the `/etc/hosts` file has entries of primary and secondary host names in both hosts.

6. Add the secondary RVG to form the replicated data set (RDS) using the following command:

```
root@primary# vradmin -g diskgroup addsec rvg_name primary_hostname secondary_hostname
```

7. Start replication between the volume replicator primary and secondary using the following command:

```
root@primary# vradmin -g group_1 -a startrep v1_rvg
```

8. Attach the external mirror disk D3 to the alternate client for taking a snapshot mirror.
9. On the volume replicator secondary (alternate client off-host), start the snap for the data volume using the following command:

```
root@secondary# vxassist -g diskgroup -b snapstart volume
```

Hardware array-based snapshot method: Hitachi Shadow Image, EMC TimeFinder, HP Business Copy (UNIX or Linux)

Ensure that the vendor-related library files are installed in the following directory: `/usr/lib`.

Appendix I

Register authorized locations

This appendix includes the following topics:

Registering authorized locations used by a NetBackup database script-based policy

During a backup, NetBackup checks for scripts in the default script location and any authorized locations. The default, authorized script location for UNIX is `usr/openv/netbackup/ext/db_ext` and for Windows is *`install_path`*`\netbackup\dbext`. If the script is not in the default script location or an authorized location, the policy job fails. You can move any script into the default script location or any additional authorized location and NetBackup recognizes the scripts. You need to update the policy with the script location if it has changed. An authorized location can be a directory and NetBackup recognizes any script within that directory. An authorized location can also be a full path to a script if an entire directory does need to be authorized.

If the default script location does not work for your environment, use the following procedure to enter one or more authorized locations for your scripts. Use `nbsetconfig` to enter an authorized location where the scripts reside. You can also use `bpsetconfig`, however this command is only available on the primary or the media server.

Note: One recommendation is that scripts should not be world-writable. NetBackup does not allow scripts to run from network or remote locations. All scripts must be stored and run locally. Any script that is created and saved in the NetBackup `db_ext` (UNIX) or `dbext` (Windows) location needs to be protected during a NetBackup uninstall.

For more information about registering authorized locations and scripts, review the knowledge base article:

https://support.cohesity.com/s/article/article-100039639

To add an authorized location

1. Open a command prompt on the client.
2. Use `nbsetconfig` to enter values for an authorized location. The client privileged user must run these commands.

   The following examples are for paths you may configure for the Oracle agent. Use the path that is appropriate for your agent.

   - On UNIX:

     ```
     [root@client26 bin]# ./nbsetconfig
     nbsetconfig>DB_SCRIPT_PATH = /Oracle/scripts
     nbsetconfig>DB_SCRIPT_PATH = /db/Oracle/scripts/full_backup.sh
     nbsetconfig>
     <ctrl-D>
     ```

   - On Windows:

     ```
     C:\Program Files\Cohesity NetBackup\NetBackup\bin>nbsetconfig
     nbsetconfig> DB_SCRIPT_PATH=c:\db_scripts
     nbsetconfig> DB_SCRIPT_PATH=e:\oracle\fullbackup\full_rman.sh
     nbsetconfig>
     <ctrl-Z>
     ```

Note: Review the NetBackup Command Reference Guide for options, such as reading from a text file and remotely setting clients from a NetBackup server using `bpsetconfig`. If you have a text file with the script location or authorized locations listed, `nbsetconfig` or `bpsetconfig` can read from that text file. An entry of `DB_SCRIPT_PATH=none` does not allow any script to run on a client. The `none` entry is useful if an administrator wants to completely lock down a server from running scripts.

3 (Conditional) Perform these steps on any clustered database or agent node that can perform the backup.

4 (Conditional) Update any policy if the script location was changed to the default or authorized location.