

NetBackup™ Web UI Microsoft SQL Server Administrator's Guide

Release 10.1.1



Last updated: 2023-02-14

Legal Notice

Copyright © 2023 Veritas Technologies LLC. All rights reserved.

Veritas, the Veritas Logo, and NetBackup are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas is required to provide attribution to the third party ("Third-party Programs"). Some of the Third-party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the Third-party Legal Notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Veritas Technologies LLC and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. Veritas Technologies LLC SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

<https://sort.veritas.com/documents>

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

NB.docs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1	About NetBackup for SQL Server	7
	Overview of NetBackup for SQL Server	7
Chapter 2	Installation and host configuration	10
	Planning the installation of NetBackup for SQL Server	10
	NetBackup server and client requirements	11
	Requirements for using NetBackup for SQL Server in a NetBackup cluster	12
	License for NetBackup for SQL Server	12
	Configuring SQL Server hosts and user permissions	13
	Installing the Veritas VSS provider for vSphere	14
	Configuring the NetBackup services for SQL Server backups and restores	14
	Configure local security privileges for SQL Server	16
	Assign an RBAC role to the SQL Server administrator	17
	Configuring the primary server host name for the SQL Server agent	17
	Configure the number of jobs allowed for backup operations	18
	Configure the Maximum jobs per client setting	19
Chapter 3	Managing SQL Server discovery and credentials	21
	About discovery of SQL Server objects	21
	Discover advanced or basic availability groups on demand	22
	Discover databases on demand	22
	Discover read-scale availability groups	23
	About SQL Server credentials	23
	Add a credential for SQL Server	25
	Select a credential for a SQL Server instance or replica	26
	View the credential name that is applied to an asset	27
	Edit or delete a named credential	27
	Remove SQL Server instances	28
	Manually add a SQL Server instance	28

Chapter 4	Managing protection plans for SQL Server	30
	About protecting SQL Server availability groups	30
	Create a protection plan to protect SQL Server assets	31
	Schedules	34
	Performance tuning and configuration options	34
	Using copy-only snapshot backups to affect how differentials are based	40
	Snapshot methods	40
	Protect a SQL Server availability group that crosses NetBackup domains	41
Chapter 5	Viewing SQL Server asset details	44
	Browse SQL Server assets	44
	View the protection status of databases, instances, or availability groups	46
Chapter 6	Protecting SQL Server	48
	Add SQL Server assets to a protection plan	48
	Customize protection settings for a Microsoft SQL Server asset	50
	Remove protection from SQL Server assets	51
Chapter 7	Restoring SQL Server	53
	Requirements for restores of SQL Server	53
	Perform a complete database recovery	54
	Recover a single recovery point	57
	Options for SQL Server restores	60
	Restore a database (non-administrator users)	61
	Select a different backup copy for recovery	62
	Restore a SQL Server availability database to a secondary replica	65
	Restore a SQL Server availability database to the primary and the secondary replicas	67
Chapter 8	Using instant access with SQL Server	70
	Prerequisites when you configure an instant access SQL Server database	70
	Hardware and configuration requirements of instant access	71
	Things to consider before you configure an instant access database	72
	Configure Samba users for SQL Server instant access	73

	Configure an instant access database	75
	View the livemount details of an instant access database	77
	Delete an instant access database	78
	Options for NetBackup for SQL Server instant access	78
	NetBackup for SQL Server terms	80
	Frequently asked questions	80
Chapter 9	Protecting SQL Server with VMware backups	85
	About protecting an application database with VMware backups	85
	Limitations of VMware application backups	86
	Create a protection plan to protect SQL Server data with a VMware backup	87
	Backup options and Advanced options	88
	Exclude disks from backups	89
	Snapshot retry options	90
	Protect SQL Server data with a VMware backup	91
	Restore SQL Server databases from a VMware backup	92
Chapter 10	Performance and troubleshooting	93
	NetBackup for SQL Server performance factors	93
	Troubleshooting credential validation	96
	Troubleshooting VMware backups	97
	SQL Server log truncation failure during VMware backups of SQL Server	99

About NetBackup for SQL Server

This chapter includes the following topics:

- [Overview of NetBackup for SQL Server](#)

Overview of NetBackup for SQL Server

The NetBackup web UI provides the capability for backups and restores of SQL Server databases. You can perform the following operations:

- View discovered instances, databases, or availability groups.
- Back up SQL Server instances, databases, and availability groups.

NetBackup offers the following types of SQL Server backup methods:

- **Protection plans**
Instances are automatically discovered in the NetBackup environment. The SQL Server administrators can select one or more protection plans that contain the wanted storage, backup, and tuning settings.
- **SQL Server Intelligent Policies**
A single policy protects multiple SQL Server instances that are spread over multiple clients. You select instances for a policy from a list of instances that are automatically discovered in the NetBackup environment.
- **Policies, using clients and batch files**
These policies include a list of SQL database clients and a batch file. The batch file contains SQL Server backup commands to run when the backup is scheduled.
- **Restore the databases that are protected with protection plans.**

- Monitor backup and restore operations.

In this guide, Microsoft SQL Server is referred to as SQL Server. NetBackup for Microsoft SQL Server is referred to as NetBackup for SQL Server.

Table 1-1 NetBackup for SQL Server features

Feature	Description
Integration with NetBackup role-based access control (RBAC)	The NetBackup web UI provides RBAC roles to control which NetBackup users can manage SQL Server operations in NetBackup. The user does not need to be a NetBackup administrator to manage SQL Server operations. However, the user still must be a member of the Windows administrator group and have the SQL Server “sysadmin” role.
Protection plans and SQL Server Intelligent Policy	The following benefits are included: ■ Use a single protection plan or policy to protect the following: Multiple SQL Server instances, instance databases, availability groups, or availability databases. Instances can be spread over multiple clients. ■ Include a full, differential, and transaction log backup in the same policy. ■ Schedule frequent backups of transaction logs. ■ You are not required to know SQL Server commands or to write and use batch files. Instead, this feature automatically generates the batch files at run-time.
Management of SQL Server assets	NetBackup automatically discovers SQL Server instances and availability groups in the environment. You can also perform manual discovery. After instances or replicas are registered, the SQL Server workload administrator can select one or more protection plans to protect the SQL Server assets.
Authentication and credentials	SQL Server protection plans and SQL Server Intelligent Policy support the following: ■ Windows authentication ■ Windows Active Directory authentication ■ With the proper configuration, you do not have to run the NetBackup service account as a privileged SQL Server user on the client.

Table 1-1 NetBackup for SQL Server features (*continued*)

Feature	Description
Backup and restore features.	The following features are available for backups and restores: ■ Backups and 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. ■ The NetBackup web UI supports the backup and restore of databases and transaction logs from one interface. Note: SQL Server recovery with the web UI requires that the SQL Server client is at version 8.3 or later. ■ Backup schedules for full, differential, or transaction log backups. ■ Manual backups and copy-only backups. ■ Support for high availability (HA) environments, including SQL Server clusters and availability groups. ■ Restore SQL Server objects to different locations (redirected restores). ■ Ability to use multiple stripes during a backup. ■ Tuning options that can improve the performance of backups.
Stream-based backups and restores	Stream-based backup and restore of SQL Server objects with SQL Server's high-speed virtual device interface.
Snapshot backups and instant access databases	NetBackup can perform backups of SQL Server with snapshot methodology. Backups with policies also offer off-host backups, Instant Recovery, and backups with a hardware provider. You can also create an instant access database from a NetBackup backup image. The database is available almost instantaneously, achieving a near-zero recovery time objective. NetBackup mounts the database's snapshot directly on the backup storage device and treats the snapshot as a normal database.
Support for VMware backups that protect SQL Server	VMware protection plans and VMware intelligent policies offer support for application-consistent, full backups of VMware computers using snapshots. The VMware intelligent policy also supports Replication Director (RD) snapshots. Use of NetBackup Accelerator can increase the speed of backups. See the following for more information. Use Accelerator to back up virtual machines Application protection options (VMware) NetBackup Replication Director Solutions Guide
Multistreaming	(Backups with policies) Ability to use multiple stripes during a backup.
Legacy SQL Server policies	Support for the legacy backup policies that use batch files and a list of clients.

Installation and host configuration

This chapter includes the following topics:

- [Planning the installation of NetBackup for SQL Server](#)
- [Configuring SQL Server hosts and user permissions](#)
- [Configuring the NetBackup services for SQL Server backups and restores](#)
- [Configure local security privileges for SQL Server](#)
- [Assign an RBAC role to the SQL Server administrator](#)
- [Configuring the primary server host name for the SQL Server agent](#)
- [Configure the number of jobs allowed for backup operations](#)
- [Configure the Maximum jobs per client setting](#)

Planning the installation of NetBackup for SQL Server

To use the new features that are included in NetBackup for SQL Server in NetBackup 10.1.1, upgrade your NetBackup for SQL Server clients to NetBackup 10.1.1. The NetBackup media server must use the same version as or a higher version than the NetBackup for SQL Server client.

[Table 2-1](#) shows the installation steps that are required to run NetBackup for SQL Server.

Table 2-1 Installation steps for NetBackup for SQL Server

Step	Action	Description
Step 1	Verify the operating system and platform compatibility.	See the NetBackup Compatibility Lists .
Step 2	Verify that primary server has a valid license for NetBackup for SQL Server and any NetBackup options or add-ons that you want to use.	See “ License for NetBackup for SQL Server ” on page 12.
Step 3	Install the NetBackup client software on the computers that have the databases that you want to back up.	See “ NetBackup server and client requirements ” on page 11.
Step 4	To protect a read-scale availability group, you must have the SQL Server Native Client version 11.0.7462 ODBC driver or later installed on the availability group replicas.	This version of the driver lets you discover and browse databases on a read-scale availability group.
Step 5	To use NetBackup for SQL Server in a NetBackup cluster, verify that your cluster environment is supported and that the NetBackup cluster is configured correctly.	See “ Requirements for using NetBackup for SQL Server in a NetBackup cluster ” on page 12.

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 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:

- To use the new features that are included in NetBackup for SQL Server in NetBackup 10.1.1, you must upgrade your NetBackup for SQL Server clients to NetBackup 10.1.1. The NetBackup media server must use the same version as the NetBackup for SQL Server client or a higher version than the client.

Requirements for using NetBackup for SQL Server in a NetBackup cluster

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

- NetBackup supports your cluster environment.
 See 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 SQL Server must exist on each node where NetBackup server resides.

License for NetBackup for SQL Server

The NetBackup for SQL Server 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 Administrator's Guide, Volume I](#).

Configuring SQL Server hosts and user permissions

The following table contains the prerequisites for users to run SQL Server backups and restores.

Table 2-2 Prerequisites for NetBackup hosts and user permissions

Step	Action	Description
Step 1	If you plan to perform VMware backups to protect SQL Server, install the Veritas VSS provider.	See “Installing the Veritas VSS provider for vSphere” on page 14.
Step 2	Assign users to the necessary RBAC roles.	See “Assign an RBAC role to the SQL Server administrator” on page 17.
Step 3	Add the necessary SQL Server credentials.	Add the SQL Server credentials that you need for database discovery and the credentials to perform recovery. See “About SQL Server credentials” on page 23. Note: To use gMSA credentials, you must use the credential option Use credentials that are defined locally on the client.
Step 4	Configure the NetBackup Client Service and the NetBackup Legacy Network Service.	This configuration allows access to the SQL Server when NetBackup performs backups and restores. See “Configuring the NetBackup services for SQL Server backups and restores” on page 14.
Step 5	Configure any necessary local security privileges.	For SQL Server credentials that use the option Use these specific credentials, an account other than Local System requires additional local security privileges. These privileges are necessary since the NetBackup for SQL Server agent logs on as the SQL Server user when it accesses data. See “Configure local security privileges for SQL Server” on page 16.

Table 2-2 Prerequisites for NetBackup hosts and user permissions
(continued)

Step	Action	Description
Step 6	Approve each valid host mapping that NetBackup discovers.	NetBackup automatically discovers many shared names and cluster names that are associated with the NetBackup hosts in your environment. Perform this configuration in Security > Host mappings on the primary server. Refer to the information on configuring host mappings in the NetBackup Web UI Administrator's Guide. Or, contact your NetBackup administrator for assistance.

Installing the Veritas VSS provider for vSphere

To use the Veritas VSS provider you must install it manually following installation of the NetBackup for Windows client. If the VMware VSS provider is installed, the installation program removes it and may require a restart of the computer.

To install the Veritas VSS provider

- 1 Browse to the following location:

```
install_path\Veritas\NetBackup\bin\goodies\
```

- 2 Double-click on the **Veritas VSS provider for vSphere** shortcut.
- 3 Follow the prompts.
- 4 When the utility has completed, restart the computer if prompted.
- 5 Following the restart, the utility resumes. Follow the prompts to complete the installation.

To uninstall the Veritas VSS provider

- 1 In the Control Panel, open **Add or Remove Programs** or **Programs and Features**.
- 2 Double-click on **Veritas VSS provider**.

The uninstall program does not automatically reinstall the VMware VSS provider.

Configuring the NetBackup services for SQL Server backups and restores

For policies and protection plans with the NetBackup web UI, NetBackup uses the NetBackup Client Service and the NetBackup Legacy Network Service to access the SQL Server when it performs backups and restores.

Note the following requirements for the NetBackup services logon account:

- The account has the fixed server role “sysadmin”. You can use a domain account, a member of BUILTIN\Administrators, or another account that has this role.
- (non-VMware backups) If you want to use Local System for the logon account, apply the SQL Server sysadmin role manually to the NT AUTHORITY\SYSTEM or the BUILTIN\Administrators group.
- (VMware backups) You must use an account other than the Local System account as the logon account. Both services must use the same logon account.
- (VMware backups) If you choose to truncate logs, ensure that the account that runs the Microsoft SQL Server Service has full permissions for the NetBackup Legacy Network Service `temp` directory.
This directory is `C:\Users\user\AppData\Local\Temp`. *User* is the account that runs the NetBackup Legacy Network Service.
- To use a gMSA account for backups and restores, you must create a credential with the option **Use credentials that are defined locally on the client**.

To configure the NetBackup services for SQL Server backups and restores

- 1 Log on to the Windows host with the account that has the SQL Server sysadmin role and any necessary local security privileges.
- 2 In the Windows Services application, open the **NetBackup Client Service**.
- 3 Configure the account as follows:
 - (non-VMware backups) Confirm that **Local System account** or a SQL Server administrator account is configured.
If you use the setting **Use credentials that are defined locally on the client** for instance credentials, both services must use the same logon account. If you use the setting **Use these specific credentials** for instance credentials, the services can use the same logon or separate logon accounts.
 - (VMware backups) Provide the name of the logon account and click **OK**.
The account must include the domain name, followed by the user account, **domain_name\account**. For example, **recovery\netbackup**.
- 4 Open the **NetBackup Legacy Network Service**.
- 5 Configure the account as follows:
 - (non-VMware backups) Confirm that **Local System account** or a SQL Server administrator account is configured.
If you use the setting **Use credentials that are defined locally on the client** for instance credentials, both services must use the same logon account. If you use the setting **Use these specific credentials** for instance

credentials, the services can use the same logon or separate logon accounts.

- (VMware backups) Provide the name of the logon account and click **OK**. Configure the same logon account for this service as you did for the NetBackup Client Service.
- 6 If you selected a different logon account, restart the services.
 - 7 If you selected the option **Use these specific credentials** for instance or for replica credentials, an account other than Local System requires certain local security privileges.

See [“Configure local security privileges for SQL Server”](#) on page 16.

- 8 For virtual environments, configure the services on the necessary services.
 - For VMware backups, configure the services for each host that you use to browse for backups and perform restores.
 - For a SQL Server cluster, configure the services on each node in the cluster.
 - For availability groups, configure the services on all replicas in the availability group where you want to run backups.

Configure local security privileges for SQL Server

If you use the option **Use these specific credentials** for instance or for replica credentials, an account other than Local System requires certain local security privileges. These privileges are necessary since the NetBackup for SQL Server agent logs on as the SQL Server user when it accesses data.

Note: This configuration applies to local security privileges only. For domain-level privileges, contact your domain administrator.

To configure the local security privileges

- 1 Open the **Local Security Policy**.
- 2 Click **Local Policies**.
- 3 In the **User Rights Assignment**, add the account to the following policies:
 - **Impersonate a client after authentication**
 - **Replace a process level token**
- 4 Restart the SQL Server.

- 5 If the NetBackup Client Service and the NetBackup Legacy Network Service use this account to log on, restart these services.
- 6 (non-VMware backups) For a SQL Server cluster, configure the local security privileges on each node in the cluster. For SQL Server availability groups, configure the services on all replicas where you want to run backups.

Assign an RBAC role to the SQL Server administrator

To protect and restore SQL Server assets, workload administrators must have permissions to access those assets and to assign credentials to those assets. The RBAC role named Default Microsoft SQL Server Administrator provides these permissions for a SQL Server administrator. Or, that user must have the equivalent permissions in a custom RBAC role.

For information on the RBAC permissions and default roles, see the NetBackup API documentation at <http://sort.veritas.com/>.

Configuring the primary server host name for the SQL Server agent

In some environments, you may need to override the host name that NetBackup for SQL Server uses for server-directed backup and restores. Specifically, when the primary server knows itself by one host name and the client must connect to a different host name to reach the primary server. For example, when the primary server has more than one IP address or associated host name. In this case some client hosts may not resolve and network route to the host name by which the primary server knows itself.

The SQL Server agent obtains the host name for the primary server from several sources, in the following order:

- **NBSERVER value.**
 For protection plans, this name is the host name by which the primary server identifies itself. Or the host name in the operation that the SQL Server backup administrator configured.
- **SQL Server agent registry setting.**
 The primary server name (**Current NetBackup Server**) in the NetBackup client properties of the NetBackup MS SQL Client interface. This setting corresponds to the following registry entry:

HKEY_CURRENT_USER\Software\Veritas\NetBackup\NetBackup for
 Microsoft SQL Server\DEFAULT_SQL_NB_MASTER_SERVER

- The first `SERVER` entry in the NetBackup registry on the client host.
 This setting is located in the following registry entry:
 HKEY_LOCAL_MACHINE\Software\VERITAS\NetBackup\CurrentVersion\Config\Server
- Domain server value.
 The host name of the primary server from which the client last requested a host ID certificate. This value is the "serverName" for the primary server in the `certmapinfo.json` file.

Alternatively, you can set `USE_REQUESTED_MASTER = FALSE` on the client to give the `NBSERVER` value lower precedence:

- SQL Server agent registry value
- Primary server value
- `NBSERVER` value
- Domain server value

To change the `USE_REQUESTED_MASTER` setting to `FALSE`

- 1 Add the following statement to a text file (for example, `new_config.txt`).

```
USE_REQUESTED_MASTER = FALSE
```

- 2 On the primary or the media server, enter the following command:

```
# bpsetconfig -h ClientA new_config.txt
```

NetBackup sets the configuration change on client host `ClientA`.

Configure the number of jobs allowed for backup operations

Last updated: 2023-02-14 09:52:02

When NetBackup starts a backup of SQL Server, a number of jobs are created. Depending on the policy configuration, additional jobs are created if you configure settings such as **Number of backup stripes** 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. It is configured in the primary server host properties in the Global Attributes . See “ Configure the Maximum jobs per client setting ” on page 19.
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 . See the NetBackup Administrator's Guide, Volume I

Configure the Maximum jobs per client setting

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 host and click **Connect**.
- 3 Click **Edit primary server**.
- 4 Click **Global attributes**.
- 5 Change the **Maximum jobs per client** value to the wanted value.
The default is 1.

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*

Refer to the following definitions:

number of database_objects 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. If striping is not used, each separate stream starts a new backup job on the client. If striping is used, each new job uses one stream per stripe.
- 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.

Managing SQL Server discovery and credentials

This chapter includes the following topics:

- [About discovery of SQL Server objects](#)
- [About SQL Server credentials](#)
- [Remove SQL Server instances](#)
- [Manually add a SQL Server instance](#)

About discovery of SQL Server objects

NetBackup discovery runs regularly and gathers information for instances and for advanced and basic availability groups in your environment. (Read-scale availability groups must be discovered manually.) The data expires after one hour. The NetBackup Discovery Service (`nbdisco`) runs “shallow” discovery every 8 hours for instances and availability groups 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.

Deep discovery includes discovery of databases and is performed in the following circumstances:

- After a full backup, an incremental backup, or a restore occurs
The client sends details when database data is changed and not more than every 15 minutes.
- When you run a manual discovery of databases or availability groups
- After you add credentials for the instances or replicas

By default, this service reports to the primary server when it finds SQL Server 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](#).

The client maintains a cache file `NB_instancename_cache_v1.0.dat` in the `NetBackup\dbext\mssql` directory for each instance. The file can be deleted and NetBackup recreates it after the next full backup when deep discovery data is sent again.

Confirmation messages in the web UI

A message `Starting the discovery of databases...` displays after you click **Discover databases** or **Discover availability groups**. 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 advanced or basic availability groups on demand

You can manually start the NetBackup discovery process if you want to immediately discover advanced or basic availability groups or replicas or discover databases in your environment. The instances or replicas must have credentials before you can perform on-demand discovery.

To discover advanced or basic availability groups

- 1 On the left, click **Workloads > Microsoft SQL Server**.
- 2 Click on the **Availability groups** tab.
- 3 Click **Discover availability groups**.
- 4 Select the host and the instance that is associated with a replica in the availability group.

Note that only registered replicas are shown in this list.

- 5 Click **Discover**.

Discover databases on demand

You can manually start the NetBackup discovery process if you want to immediately discover instance databases or availability databases in your environment.

To discover databases

- 1 On the left, click **Workloads > Microsoft SQL Server**.
- 2 Click on the **Databases** tab.
- 3 Click **Discover databases**.
- 4 Select the host and the instance that is associated with the databases.
Note that only registered instances are shown in this list.
- 5 Click **Discover**.

Discover read-scale availability groups

Read-scale availability groups are not discovered automatically. You must specify one of the replicas in the availability group and manually start discovery.

To discover read-scale availability groups

- 1 On the left, click **Workloads > Microsoft SQL Server**.
- 2 Click on the **Instances** tab.
- 3 Select one of the replicas that is part of the availability group and click **Manage credentials**.
- 4 Follow the prompts to provide the credentials for the replica.
- 5 Click on the **Availability groups** tab.
- 6 Click **Discover availability groups**.
- 7 Select the host and the instance that is associated with a replica in the availability group.
Note that only registered replicas are shown in this list.
- 8 Click **Discover**.

About SQL Server credentials

To protect SQL Server, you must add (or register) credentials to the SQL Server instances or availability replicas. The NetBackup web UI supports Windows authentication and Windows Active Directory authentication. It does not support Mixed Mode or SQL Server authentication. Credentials are not supported at the database or the availability group level.

Table 3-1 Options to register credentials

Option to register credentials	Environment and configuration
Use these specific credentials (recommended)	■ The SQL Server DBA provides the NetBackup administrator with the SQL Server user credentials. ■ The SQL Server DBA does not want the NetBackup services running as a privileged SQL Server user on the client. Configuration requirements The user account that is used to register credentials must have the SQL Server “sysadmin” role and be a member of the Windows Administrators group. The NetBackup services can use the Local System logon account. If you want to use a different logon account, that account must also have certain local security privileges.
Use credentials that are defined locally on the client	■ The NetBackup services run as a privileged SQL Server user on the client. ■ The SQL Server DBA does not want to provide credentials to register instances or replicas. ■ The NetBackup administrator does not have access to the SQL Server credentials. ■ You want to use gMSA credentials. Configuration requirements The user account that is used to register credentials must have the SQL Server “sysadmin” role and be a member of the Windows Administrators group. You must also configure the logon account for the NetBackup Client Service and the NetBackup Legacy Network service.

Registering instances when SQL Server hosts are clustered or use multiple NICs

When NetBackup discovers a SQL Server 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 SQL Server 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 SQL Server 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 SQL Server cluster that uses multiple NICs, add credentials to the instance with the private virtual name of the SQL Server cluster.

Registering Microsoft SQL Server failover cluster instances (FCIs)

NetBackup discovers and displays failover cluster instances (FCIs) under the cluster name and the physical node names. For example, instance `FCI` is enumerated with both its physical nodes `hostvm10` and `hostvm11` and with its cluster name `sql-fci`. Databases that exist for FCIs are also enumerated with the node names and the cluster name. Depending on how you want to protect a database, add credentials to either the cluster name (that are valid for all nodes) or to a physical node name.

Validation of credentials

After you add credentials, NetBackup validates the credentials and starts the database and availability group discovery. When discovery completes, the results are displayed on the **Databases** or the **Availability group** tab.

For a SQL Server cluster or if an availability group instance is part of SQL Server cluster, NetBackup validates the credentials on the active node. The credentials must be valid for all nodes in the cluster. For a SQL Server availability group, replicas are registered and validated individually. 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.

Add a credential for SQL Server

This type of credential allows NetBackup to access a SQL Server.

To add a credential for SQL Server

- 1 On the left, click **Credential management**.
- 2 On the **Named credentials** tab, click **Add** and 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*)
- 3 Click **Next**.
- 4 Select **Microsoft SQL Server**.
- 5 Provide the authentication details that are needed to connect to the SQL Server. See [“About SQL Server credentials”](#) on page 23.
- 6 Click **Next**.
- 7 Add one or more RBAC roles that you want to have access to the credential.
 - Click **Add**.

- Select the role name.
 - Select the credential permissions that you want the role to have.
- 8 Click **Next** and follow the prompts to complete the wizard.

Select a credential for a SQL Server instance or replica

To allow for full discovery of SQL Server assets, you must select the credential that you want to use for the instances or replicas. You can select from the existing server credentials or create a new credential.

Note: The database and the availability group discovery begin after the credentials are validated. However, these assets may not appear in the web UI immediately. They appear after the discovery process completes. The date reflects when the credential was added or updated; it does not indicate if the credential is valid.

Review the recommendations and requirements for the type of credentials that you want to use for authentication.

See [“About SQL Server credentials”](#) on page 23.

See [the section called “Select an existing credential for a SQL Server instance or replica”](#) on page 26.

See [the section called “Add a new credential for a SQL Server instance or replica”](#) on page 27.

Select an existing credential for a SQL Server instance or replica

For availability groups, each replica in must be registered with credentials.

To select an existing credential for a SQL Server instance or replica

- 1 On the left, click **Workloads > Microsoft SQL Server**.
- 2 Click on the **Instances** tab.
- 3 Select the check box for the instances or replicas and click **Manage credentials**.
- 4 Click **Select from existing credentials**.
- 5 Click **Next**.
- 6 Select the credential that you want to use for the selected assets and click **Next**.

Add a new credential for a SQL Server instance or replica

For availability groups, each replica in must be registered with credentials.

To add a new credential for a SQL Server instance or replica

- 1 On the left, click **Workloads > Microsoft SQL Server**.
- 2 Click on the **Instances** tab.
- 3 Select the check box for the instances or replicas and click **Manage credentials**.
- 4 Select **Add credentials** and click **Next**.
- 5 Select the authentication option.

See [“About SQL Server credentials”](#) on page 23.

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.

To view credentials for SQL Server instances

- 1 On the left, select **Workloads > Microsoft SQL Server**.
- 2 On the **Instances** tab, locate the **Credential name** column.

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.

Note: Ensure that the credential category used for **AHV cluster** is *AHV* and for **Nutanix Prism Central** is *Prism Central*.

To edit a named credential

- 1 On the left, click **Credential management**.
- 2 On the **Named credentials** tab, locate and click on the credential that you want to edit.

- 3 Click **Edit** and update the credential as needed.
- 4 Review the changes and click **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, click **Credential management**.
- 2 On the **Named credentials** tab, locate and click on the credential that you want to delete.
- 3 Click **Delete**.

Remove SQL Server instances

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

To remove a SQL Server instance

- 1 On the left, click **Workloads > Microsoft SQL Server**, then click the **Instances** tab.
- 2 Locate and select the checkbox 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.

Manually add a SQL Server instance

Newly discovered SQL Server 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 SQL Server instance

- 1 On the left, click **Workloads > Microsoft SQL Server**, then click the **Instances** tab.
- 2 Click **Add**.

- 3** Provide the **Host** name where the instance resides and the **Instance name**.
 - For a SQL Server cluster, the host name is the virtual name of the SQL Server 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 SQL Server host or of the virtual SQL Server.
 - For a failover cluster instance, enter the virtual name of the SQL Server cluster.
NetBackup enumerates the FCI under the physical node names and the cluster name.
- 4** Click **Next**.
- 5** Review the roles that have access to the instance. Click **Add** to give additional roles access to the instance.
- 6** Click **Manage credentials** to add the credentials for this instance.
See [“Select a credential for a SQL Server instance or replica”](#) on page 26.
You may omit credentials at this time. The instance is marked as unregistered and the **Registered** column is empty.
- 7** Click **Finish**.

Managing protection plans for SQL Server

This chapter includes the following topics:

- [About protecting SQL Server availability groups](#)
- [Create a protection plan to protect SQL Server assets](#)
- [Protect a SQL Server availability group that crosses NetBackup domains](#)

About protecting SQL Server availability groups

NetBackup for SQL Server supports backups and restores of SQL Server Always On and read-scale availability groups. For information on supported versions and environments, see the [Application/Database Agent Compatibility List](#).

You can protect an availability group environment in the following ways:

- With a protection plan that protects the preferred or the primary replica.
- If an availability group crosses multiple NetBackup domains, you can use Auto Image Replication (A.I.R.) to replicate the backup to the other NetBackup domains.

See [“Protect a SQL Server availability group that crosses NetBackup domains”](#) on page 41.

Note the following before you configure the protection plan:

- NetBackup can only fully protect the availability group environment if each replica on which backups occur is registered with credentials.
- NetBackup runs a backup job on each replica in the availability group. On the replicas which are not the backup source, the job skips the backup.

Limitations

Note the following limitations for backups of availability groups:

- NetBackup does not support the following types of backups for availability databases:
 - Backups of filegroups or files
 - VMware backups
 - A grouped snapshot backup
 - Backups of non-readable secondary replicas
 NetBackup can only back up databases in a replica when you allow user connections for the replica.
 If a secondary replica is the preferred replica and it is non-readable, the backup fails. If a secondary replica is not the preferred replica, NetBackup skips the backup of that replica.

SQL Server does not support the following types of backups on a secondary replica:

- Full backups
 If a full backup takes place on a secondary replica, NetBackup converts the full backup to a copy-only backup.
- Differential backups
 Backups of this type result in a failed backup.
- Copy-only transaction log backups
 Backups of this type result in a failed backup.

Create a protection plan to protect SQL Server assets

You can create a protection plan to perform scheduled backups of SQL Server assets.

To create a protection plan to protect SQL Server assets

- 1 On the left, click **Protection > Protection plans** and then click **Add**.
- 2 In **Basic properties**, enter a **Name**, **Description**, and select **Microsoft SQL Server** from the **Workload** list.

(Optional) Add a policy name prefix. NetBackup automatically creates a policy when a user subscribes assets to a protection plan; NetBackup appends this prefix to that policy name.

3 In Schedules and retention, click Add.

You can select the frequency and the retention backup. You can set up the following backup schedules: **Full**, **Differential incremental**, or **Transaction log**.

In the **Attributes** tab:

- Select the **Backup type**, how often it runs, and how long to keep the backup for this schedule.

In the **Start window** tab:

- Define a **Start day**, **Start time**, **End day**, and **End time** for this schedule using the options available on the screen. Or you can drag your cursor over the time boxes to create the schedule.
- Use the options on the right to duplicate, remove, or undo changes to a schedule.

Click **Save** after all options are selected in the **Attributes** and the **Start window** tabs.

Review the **Backup schedule preview** window and verify that all schedules are set correctly.

- In **Storage options**, configure the storage type per schedule you configured in step 3.

The options vary depending on storage options currently setup to work with NetBackup.


A protection plan can only use the storage that a NetBackup 8.1.2 or newer media server can access.

Storage option	Requirements	Description
Perform snapshot backups		Performs a point-in-time, read-only, disk-based copy of a client volume. NetBackup backs up data from the snapshot, not directly from the client's primary or original volume. Snapshots cannot be used to perform differential backups or transaction log backups. In those cases NetBackup performs a stream-based backup. You can select from the following methods: Automatic, VxVM, or VSS. See "Snapshot methods" on page 40. Note that SQL Server dynamic file allocation can reduce the likelihood that any of the component files contain large areas of empty space.
Backup storage	OpenStorage is required for this option. Tape, storage unit groups, and Replication Director not supported.	Click Edit to select the storage target. Click Use selected storage after selecting the storage target.
Transaction log options		When you configure a transaction log schedule, you can choose to use the same storage that is used for database backups. Or, you can choose unique storage for transaction logs.

- In **Backup options**, configure the options that you want.

See ["Performance tuning and configuration options"](#) on page 34.

Click the **Database options** tab to select options for databases. Click the **Transaction log options** tab to select options for transaction logs.



Note: For availability groups, ensure that you select a **Availability database backup preference** setting for databases and for transaction logs.

- 6 In **Permissions**, review the roles that have access to the protection plan.
To give another role access to this protection plan, click **Add**. Select the **Role** in the table and customize the role by adding or removing permissions in the **Select permissions** section.
- 7 In **Review**, verify that the protection plan details are correct and click **Save**.

Schedules

The following schedule settings are included in a protection plan.

Note that when you customize a protection plan for an asset, you can only edit the following schedule settings:

- Start window
- (SQL Server transaction logs) Recurrence
- (SQL Server transaction logs) Keep for

Table 4-1 Schedule options for protection plans

Option	Description
Backup type	The type of backup that the schedule controls.
Recurrence (frequency)	How frequently or when to run the backup.
Keep for (retention)	How long to keep the files that were backed up by the schedule.
Replicate this backup	Replicates the snapshot to another volume.
Duplicate a copy immediately to long-term retention	Immediately after the schedule is created, a copy is duplicated to the media that is selected for long-term storage.
Start window	On this tab, set the window during which a backup can start.

Performance tuning and configuration options

[Table 4-2](#) describes the tuning and configuration options that are available that are available in protection plans. Some options cannot be changed when you edit an existing plan or when you subscribe an asset to the plan.

Note: Filegroup backups are only available for policies and not for protection plans.

Table 4-2 Performance tuning and configuration options

Field	Description
Backup stripes	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 SQL Server dumps data faster than your tape drive is capable of writing. The default value for this option is 1. Range is 1–32.
Client buffers per stripe	(Stream-based backups only) This option affects buffer space availability. NetBackup uses this parameter to decide how many buffers to allocate for reading or writing each data stream during a backup operation. By allocating a greater number of buffers, you can affect how quickly NetBackup can send data to the NetBackup primary server. The default value for this option is 2, which allows double buffering. You may get slightly better performance by increasing this value to a higher value. Range is 1–32.
Maximum transfer size	(Stream-based backups only) This option is the buffer size used by SQL Server for reading and writing backup images. Generally, you can get better SQL Server performance by using a larger value. This option can be set for each backup operation. Calculated as $64 \text{ KB} * 2^{\text{MAX_TRANSFER_SIZE}}$. It ranges in size from 64 KB to 4 MB. The default is 4 MB.
Backup block size	This option applies to stream-based backups only. Sets the incremental size that SQL Server uses for reading and writing backup images and can be set for each backup operation. Calculated as $512 \text{ bytes} * 2^{\text{BLOCK_SIZE}}$. The value for this option ranges from 0.5 KB to 64 KB. The default is 64 KB.
Parallel backup operations	This option is the number of backup operations to start simultaneously, per database instance. Range is 1–32. The default is 1.
VDI timeout (seconds)	(Databases only) Determines the time-out interval for SQL Server Virtual Device Interface. The selected interval is applied to backups and restores of databases and of transaction logs. The default value for backups is 300. The default value for restore jobs is 600. Range is 300–2147483647.
Use Microsoft SQL Server compression	Enable this option to use SQL Server to compress the backup image. If you enable SQL Server compression, do not enable NetBackup compression. SQL Server compression is not supported for snapshot backups.

Table 4-2 Performance tuning and configuration options (*continued*)

Field	Description
Skip unavailable (offline, restoring, etc.) databases	NetBackup skips any database with a status that prevents NetBackup from successfully backing up the database. These statuses include offline, restoring, recovering, and emergency mode, etc. NetBackup skips the backup of the unavailable database, but continues with the backup of the other databases that are subscribed to the protection plan. The backup completes with a status 0 and the job details indicate that the database was skipped.
Create copy-only backup	(Databases only) This option allows SQL Server to create an out-of-band backup so that it does not interfere with the normal backup sequence. See “Using copy-only snapshot backups to affect how differentials are based” on page 40.
Perform Microsoft SQL Server checksum	Choose one of the following options for SQL Server backup checksums: ■ None. Disables the backup checksums. ■ To verify the checksums before the backup, choose one of the following options. Note that these options impose a performance penalty on a backup or restore operation. ■ Continue on error. If the backup encounters a verification error, the backup continues. ■ Fail on error. If the backup encounters a verification error, the backup stops.

Table 4-2 Performance tuning and configuration options (*continued*)

Field	Description
Convert incremental backup to full backup	(Databases only) If no previous full backup exists for the database, then NetBackup converts a differential backup to a full backup. The agent determines if a full backup exists for each database. If no previous full backup exists, a differential backup is converted to a full as follows: ■ If you select a database for a differential backup, the backup is converted to a full database backup. (For policies) If the Skip read-only file groups option is selected the backup is converted to a full read/write filegroup backup. ■ (For policies) If you select a filegroup for a differential backup, NetBackup does the following: ■ If the filegroup is the default database filegroup, NetBackup converts the backup to a full filegroup backup. ■ If the filegroup is a secondary filegroup and a backup of the primary filegroup does not exist, NetBackup converts the backup to a partial full database backup. This backup contains the selected filegroup and default filegroup. ■ If the filegroup is a secondary filegroup and a backup of the primary filegroup does exist, NetBackup converts the backup to a full filegroup backup of the selected filegroup. ■ For snapshot backup policies, a Full schedule must exist to successfully convert differential backups to full backups. Note: NetBackup only converts a differential backup if a full backup was never performed on the database. If a full backup does not exist in the NetBackup catalog but SQL Server detects an existing full LSN, NetBackup performs a differential backup and not a full. In this situation, you can restore the full backup with native tools and any differentials with the NetBackup MS SQL Client. Or, if you expired the backup in NetBackup, you can import the full backups into the NetBackup catalog. Then you can restore both the full and the differential backups with the NetBackup MS SQL Client.

Table 4-2 Performance tuning and configuration options (*continued*)

Field	Description
Convert transaction log backup to full backup	(Transaction logs only) If no previous full backup exists for the database, then NetBackup converts a transaction backup to a full backup. This option also detects if a full recovery database was switched to the simple recovery model and back to the full recovery model. In this scenario, the log chain is broken and SQL Server requires a differential backup before a subsequent log backup can be created. If NetBackup detects this situation, the backup is converted to a differential database backup. Note: NetBackup only converts a transaction log backup if a full backup was never performed on the database. If a full backup does not exist in the NetBackup catalog but SQL Server detects an existing full LSN, NetBackup performs a transaction log backup and not a full. In this situation, you can restore the full backup with native tools and any differentials and log backups with the NetBackup MS SQL Client. Or, if the backup is expired, you can import the full backups into the NetBackup catalog. Then you can restore the full, differential, and log backups with the NetBackup MS SQL Client.

Table 4-2 Performance tuning and configuration options (*continued*)


Field	Description
Availability database backup preference	This option determines where backups of availability groups occur. Select a setting for databases in the Database options tab. Select a setting for transaction logs on the Transaction logs tab. ■ Disabled (for policies, the equivalent setting is None) Performs the backup on the specified instance. Use this option when you intend to protect individual availability databases. ■ Protect primary replica Backups always occur on the primary replica. This option applies to availability replicas and to instances that have both standard databases and availability databases. ■ Protect preferred replica Honors your SQL Server backup preferences. These preferences include the preferred replica, backup priority, and excluded replicas. Note that NetBackup initiates a backup job on each replica. The backup is skipped on any replica that isn't the intended backup source. This option applies to availability replicas and to instances that have both standard databases and availability databases. ■ Skip availability databases Skips any availability databases on the instance. Use this option when you intend to protect any instances that contain both standalone databases and availability databases and only want to protect the standalone databases. Note: Do not select this option if you intend to protect availability groups. Backup preference for individual availability databases Note the following behavior when you select a protection plan to protect individual availability databases. ■ If the preference for Databases is set to Skip availability databases, scheduled backups cannot succeed. Databases must have the setting None, Protect preferred replica, or Protect primary replica. ■ When a user selects Backup now to back up an availability database, the backup is performed on the selected node. The image is cataloged under the cluster name.
Truncate logs after backup	(Transaction logs only) This option backs up the active part of the transaction log and then marks it inactive or empty. This option is enabled by default.

Using copy-only snapshot backups to affect how differentials are based

When you use both full backups and snapshot backups to protect SQL Server, the previous snapshot backup expires after the next snapshot backup is created. If you require a point in time restore before the latest backup, the differentials are based on a snapshot backup that no longer exists. Alternatively, NetBackup lets you create copy-only backups that are out-of-band so the backup does not reset the differential baseline. Differential backups are then based on the last full backup.

If a failure occurs and is detected immediately, you can restore the last full backup. Then you can replay the necessary transaction logs to achieve recovery. However, if a failure is not detected until after the next full backup, then there are no snapshot backups available to restore. When you use copy-only backups, each differential is instead based on the last full backup that is not copy-only. You can restore the last full backup, restore the latest differential backup, then restore the necessary transaction log backups before the error occurred.

Figure 4-1 Recovering after an error when using full and copy-only backups



Snapshot methods

The following snapshot methods and options are available for snapshot backups. For more details see the [NetBackup Snapshot Client Administrator's Guide](#).

Table 4-3

Method	Description
Automatic	NetBackup selects a snapshot method when the backup starts. If necessary, NetBackup selects a different method for assets in the protection plan.

Table 4-3 (continued)

Method	Description
VSS	VSS uses the Volume Shadow Copy Service of Windows. VSS is for local backup and it selects the actual snapshot method depending on which snapshot provider is configured on the client. Provider type: ■ Automatic. NetBackup selects the available provider in this order: Hardware, Software, System. ■ System. Use the Microsoft system provider for a block-level copy on write snapshot. ■ Use the software provider to intercept the I/O requests at the software level between the file system and the volume manager. ■ Use the hardware provider for your disk array. Snapshot attribute: ■ Automatic. NetBackup selects the attribute. ■ Differential. Use a copy-on-write type of snapshot. ■ Plex. Use a clone or a mirror type of snapshot.
VxVM	For any snapshots with any data that is configured over Volume Manager volumes. ■ Resynchronize mirror in background. Select this option to allow more efficient use of backup resources. If two backups need the same tape drive, the second can start even though the resynchronize operation for the first job has not completed. ■ Wait for mirror sync completion. Select this option if full-sized instant snapshots are not available for backup until the mirror synchronization is complete. If the backup starts before the snapshot disks are fully synchronized with the source and the server does not have access to the source disks, then the backup fails. ■ Maximum number of volumes to resynchronize. The number of volume pairs that are resynchronized simultaneously. Accept the default if the I/O bandwidth in your clients and disk storage cannot support simultaneous synchronization of volumes. For the configurations that have sufficient I/O bandwidth, multiple volumes can be resynchronized simultaneously, to complete resynchronization sooner. A major factor in I/O bandwidth is the number and speed of HBAs on each client.

Protect a SQL Server availability group that crosses NetBackup domains

When you have an availability group that crosses NetBackup domains, you can use Auto Image Replication (A.I.R.) to replicate backup images to another NetBackup domain. The following configuration requirements exist:

- Configure the storage in the NetBackup source and target domains:
 - For OpenStorage, a disk appliance of the same type in each domain. The disk appliance type must support NetBackup Auto Image Replication (A.I.R.).

- For NetBackup deduplication, the storage that NetBackup can use for a Media Server Deduplication Pool in each domain.
- Configure the domain where the backups occur as the source domain. Then configure the domain where you want to restore the backups as the target domain.

To create a protection plan to protect a SQL Server availability group that crosses domains

- 1** On the left, click **Protection > Protection plans** and then click **Add**.
- 2** In **Basic properties**, enter a **Name** and **Description**.
- 3** From the **Workload** list, select **Microsoft SQL Server**.
- 4** In **Schedules and retention**, click **Add**.

You can set up a full, differential, or transaction log backup.

In the **Attributes** tab:

- Select the **Backup type**, how often it runs, and how long to keep the backup for this schedule.
- Select **Replicate this backup**.
 - The backup storage must be a source in a targeted A.I.R. environment. The **Replication target** is configured in step [5](#).
 - For more information about replication, review *About NetBackup Auto Image Replication* in the [NetBackup Administrator's Guide, Volume I](#).

In the **Start window** tab:

- Define a start window for this schedule using the options available on the screen. You can add multiple schedule windows for this schedule if needed.

Review the **Backup schedule preview** and verify that all schedules are set correctly.

- 5 In **Storage options**, configure the storage type per schedule you configured in step 5.

A protection plan can only use the storage that a NetBackup 8.1.2 or newer media server can access.

Storage option	Requirements	Description
Backup storage	OpenStorage is required for this option. Tape, storage unit groups, and Replication Director not supported.	Click Edit to select the storage target. Click Use selected storage after selecting the storage target.
Replication target	The backup storage must be a source in a targeted A.I.R. environment.	Click Edit to select the replication target primary server. Select a primary server and then select a storage lifecycle policy. Click Use selected replication target to return to the storage options screen.

- 6 In **Backup options**, select the options that you want.

From the **Availability database backup preference** list, choose one of the following:

- **Protect primary replica**
- **Protect preferred replica**

See “[Performance tuning and configuration options](#)” on page 34.

(Optional) Make any other changes to the tuning parameters.

- 7 In **Permissions**, review the roles that have access to this protection plan.
- 8 In **Review**, verify that the protection plan details are correct and click **Finish**.

Additional resources

[NetBackup Administrator's Guide, Volume I](#)

[NetBackup Deduplication Guide](#)

[NetBackup OpenStorage Solutions Guide](#)

[NetBackup Administrator's Guide, Volume I](#)

[NetBackup Deduplication Guide](#)

[NetBackup OpenStorage Solutions Guide](#)

<http://www.netbackup.com/compatibility>

Viewing SQL Server asset details

This chapter includes the following topics:

- [Browse SQL Server assets](#)
- [View the protection status of databases, instances, or availability groups](#)

Browse SQL Server assets

You can browse instances, databases, and availability groups to view their details such as how they are protected and recovery points that are available.

Note: Classic policy information is displayed for databases but not for instances or availability groups. The web UI indicates if a protection plan protects the instance or replica, but not if a classic policy does. However, when a backup using a classic policy is performed on an individual database, the **Protected by** column displays the classic policy name.

Browse SQL Server instances

On the **Instances** tab you can view and manage instances, including how they are protected and the instance credentials.

To browse SQL Server instances

- 1 On the left, click **Workloads > Microsoft SQL Server**.
- 2 Click the **Instances** tab.

- 3 To view the available actions for one or more instances, select the checkbox for the instances. Note that **Backup now** is only available when you select one instance.
- 4 To view the details for an instance, click the instance. You can perform the following tasks.
 - Perform an immediate backup of the instance by clicking **Backup now**.
 - Click **Add protection** to add the instance to a protection plan.
 - Click **Remove protection** to remove an instance from a protection plan.
 - To see the databases that are discovered the instance and their protection information and status, click on the **Databases** tab.
 - To view the roles that have access to the instance, click the **Permissions** tab.

Browse SQL Server availability groups

On the **Instances** tab you can view and manage availability groups, including how database and replica details and how the availability group is protected.

To browse SQL Server availability groups

- 1 On the left, click **Workloads > Microsoft SQL Server**.
- 2 To view the available actions for one or more availability groups, select the check box for the availability groups. Note that **Backup now** is only available when you select one availability group.
- 3 Click on an availability group to view its details. You can perform the following tasks.
 - Click **Backup now** to perform an immediate backup of the instance.
 - Click **Add protection** to add the availability group to a protection plan.
 - Click **Remove protection** to remove an availability group from a protection plan.
 - To see the databases that are discovered for the availability group and their protection information and status, click on the **Databases** tab.
 - To see the replicas for the availability group and their protection information and status, click on the **Replicas** tab.
 - To view the roles that have access to the availability group, click the **Permissions** tab.

Browse SQL Server databases

Note: Databases only appear on the **Databases** tabs if they meet one of the following criteria: A backup exists of the database, the database instance has validated credentials, or a manual discovery of databases was performed.

To browse SQL Server databases

- 1 On the left, click **Workloads > Microsoft SQL Server**.
- 2 Click the **Databases** tab.
- 3 To view the available actions for one or more databases, select the check box for each database. Note that **Backup now** is only available when you select one database.
- 4 To view the details for a database, click the database. You can perform the following tasks.
 - Click **Backup now** to perform an immediate backup of the instance.
 - Click **Add protection** to add the database to a protection plan.
 - Click **Remove protection** to remove a database from a protection plan.
 - To see the available recovery points for the database, click **Recovery points**.
 - To view the restore jobs for the database, click **Restore activity**.
 - To view the roles that have access to the database, click the **Permissions** tab.

View the protection status of databases, instances, or availability groups

You can view the protections plans that are used to protect instances or availability groups.

To view the protection status of databases, instances, or availability groups

- 1 On the left, click **Workloads > Microsoft SQL Server**.
- 2 Click on one of the following tabs: **Databases**, **Instances**, or **Availability groups**.
- 3 The **Protected by** column indicates how the asset is protected.

Table 5-1 Protection status of SQL Server assets

Protection type or status	Protected by column	
	Database	Instance or availability group
Asset is protected by a classic policy	Classic policy	Not protected Go to Protection > Policies to see how classic policies are used to protect instances or availability groups.
Asset is protected by a protection plan	Protected	Protected
Asset is not protected by plan or a policy	Not protected	Not protected
A policy or protection plan protects the asset, but it is not backed up yet (no backup image exists).	Not protected Protected by column is blank.	Not protected

Protecting SQL Server

This chapter includes the following topics:

- [Add SQL Server assets to a protection plan](#)
- [Customize protection settings for a Microsoft SQL Server asset](#)
- [Remove protection from SQL Server assets](#)

Add SQL Server assets to a protection plan

The following procedure describes how to subscribe an SQL Server asset to a protection plan. When you subscribe an asset to a protection plan, you assign predefined backup settings to the asset.

Note the following:

- For backups to be successful, a SQL Server instance or replica must have a valid credential configured for it in **Instances** tab.
See [“Select a credential for a SQL Server instance or replica”](#) on page 26.
- Your user account is assigned to the RBAC role **Default Microsoft SQL Server** or another role with the same permissions for protection plans and for SQL Server.
See [Default RBAC roles](#) and [RBAC permissions](#) in the [NetBackup Web UI Administrator's Guide](#). Or, contact your NetBackup administrator for assistance.
- Ensure other requirements are met for the NetBackup environment and for non-administrator users.
See [“Configuring SQL Server hosts and user permissions”](#) on page 13.
- Databases only appear on the **Databases** tabs if they meet one of the following criteria: A backup exists of the database, the database instance has validated credentials, or a manual discovery of databases was performed.

To add SQL Server assets to a protection plan

1 On the left, click **Workloads > Microsoft SQL Server**.

2 Choose the asset or assets that you want to protect.

All the databases in an instance	■ On the Instances tab, select the box for the instance that you want to protect.
An individual database	■ On the Instances tab, click on the instance that contains the database you want to protect. ■ On the Databases tab, click the box for one or more databases.
An availability group	■ On the Availability groups tab click the box for the availability group name.
An individual availability database	■ On the Availability groups tab click on the availability group name that contains the database that you want to protect. ■ On the Databases tab, click the box for one or more databases.
A SQL Server cluster	■ On the Instances tab, select the box for the instance that belongs to the cluster. The Host name is the virtual name of the SQL Server cluster.
A SQL Server failover cluster instance (FCI)	On the Instances tab, select the instance name depending on if you want to protect the cluster or a node in the cluster: ■ The instance name, where the Host name is the cluster name of the FCI. The backup is attempted on the active node. Both nodes must be hosts of the same primary server and the instances must have valid credentials registered. ■ The instance name, where the Host name is one of the physical node names of the FCI. For the backup to succeed, this node must be the active node in the cluster. The backup is cataloged under the cluster name.
A SQL Server host that uses multiple NICs	On the Instances tab, select the instance: ■ The instance name, where the Host name is the private interface name of the SQL Server host. ■ The instance name for a SQL Server cluster that uses multiple NICs, where the Host name is the private interface name of the virtual SQL Server.

3 Click **Add protection**.

4 Select a protection plan and click **Next**.

- For a snapshot backup, look for a protection plan that lists **Snapshot options** and a **Snapshot method**.
See [“Snapshot methods”](#) on page 40.

- For an availability group, select a protection plan that has a configured **Availability database backup preference**, either **Protect primary replica** or **Protect preferred replica**.
Do not subscribe an availability group to a protection plan that has a setting of **None** or **Skip availability databases**.
- 5 You can adjust one or more of the following settings:
- **Schedules and retention**
Change the backup start window. For transaction log schedules, you can also edit the frequency and the retention.
See [“Schedules”](#) on page 34.
 - **Backup options and Configuration options**
Adjust the performance tuning options or change or enable any options for the protection plan.
See [“Performance tuning and configuration options”](#) on page 34.
- 6 Click **Protect**.
- The results of your choices appear under **Instances** or **Databases**.

Customize protection settings for a Microsoft SQL Server asset

You can customize certain settings for a protection plan, including the schedule backup window and other options.

- See [“Schedules”](#) on page 34.
- See [“Performance tuning and configuration options”](#) on page 34.

To customize protection settings for a Microsoft SQL Server asset

1 On the left, click **Workloads > Microsoft SQL Server**.

2 Do one of the following:

- | | |
|---|---|
| Edit the settings for an instance | ■ On the Instances tab, click on the instance that you want to edit. |
| Edit the settings for a database | ■ On the Databases tab, click on the database that you want to edit. |
| Edit the settings for an availability group | ■ On the Availability groups tab, click on the availability group that you want to edit. |

Edit the settings for an availability database

- On the **Databases** tab, click on the database that you want to edit.

- 3 Click **Customize protection > Continue**.
- 4 Adjust any of the following settings:
 - The backup start window.
See “[Schedules](#)” on page 34.
 - For transaction log schedules, you can edit the frequency and the retention.
See “[Schedules](#)” on page 34.
 - **Backup options**
Adjust the performance tuning options or change or enable any configuration options for the protection plan.
See “[Performance tuning and configuration options](#)” on page 34.
- 5 Click **Protect**.

Remove protection from SQL Server assets

You can unsubscribe databases, instances, or availability groups from a protection plan. When the asset is unsubscribed, backups are no longer performed.

Note: When you unsubscribe an asset from a protection plan, there is a possibility that the asset displays **Classic policy** in the web UI. This situation can happen when an asset is subscribed to a protection plan and a backup runs for that asset. Then the asset is unsubscribed from the protection plan while it has a valid backup image. The web UI displays **Classic policy**, but there may or may not be an active policy protecting the asset.

To remove protection from an instance

- 1 On the left, click **Workloads > Microsoft SQL Server**.
- 2 Select the asset that you want to unsubscribe.

Remove protection from an instance

- On the **Instances** tab, click on the instance that you want to edit.

Remove protection from a database

- On the **Databases** tab, click on the database that you want to edit.

Remove protection from an availability group ■ On the **Availability groups** tab, click on the availability group that you want to edit.

Remove protection from an availability database ■ On the **Databases** tab, click on the database that you want to edit.

3 Click **Remove protection > Yes**.

The asset is listed as **Not protected**.

Restoring SQL Server

This chapter includes the following topics:

- [Requirements for restores of SQL Server](#)
- [Perform a complete database recovery](#)
- [Recover a single recovery point](#)
- [Options for SQL Server restores](#)
- [Restore a database \(non-administrator users\)](#)
- [Select a different backup copy for recovery](#)
- [Restore a SQL Server availability database to a secondary replica](#)
- [Restore a SQL Server availability database to the primary and the secondary replicas](#)

Requirements for restores of SQL Server

To restore perform restores of SQL Server, the following requirements exist:

- NetBackup services are correctly configured.
See [“Configuring SQL Server hosts and user permissions”](#) on page 13.
- Both administrators or non-administrators can perform restores. However, additional configuration steps are required for non-administrators.
Administrators must provide during the restore a user account that is a member of the Windows administrator group and a member of the local SQL Server sysadmin role.
Non-administrators must follow these additional steps for successful recovery:
See [“Restore a database \(non-administrator users\)”](#) on page 61.

- The user that signs into the NetBackup web UI is assigned to the RBAC role **Default Microsoft SQL Server Administrator** or another role with the same restore permissions for SQL Server.
 See [Default RBAC roles](#) and [role permissions](#) in the [NetBackup Web UI Administrator's Guide](#). Or, contact your NetBackup administrator for assistance.
 See [Default RBAC roles](#) and [role permissions](#) in the [NetBackup Web UI Administrator's Guide](#). Or, contact your NetBackup administrator for assistance.
- The security administrator has configured the necessary mappings for the hosts in **Security > Hosts**.
 Refer to the information on [configuring host mappings](#) in the [NetBackup Web UI Administrator's Guide](#). Or, contact your NetBackup administrator for assistance.
- To restore to a different server (host), the following requirements and conditions exist:
 - NetBackup must have the ability to communicate with the destination client.
 - Non-administrator users can only perform restores from their own backups.

Perform a complete database recovery

A complete database recovery selects all the backup images that are necessary to restore the complete database. It also leaves the database in the recovered state, or ready to use.

To perform a complete database recovery

- 1 On the left, select **Workloads > Microsoft SQL Server**.
- 2 On the **Databases** tab, locate the database that you want to restore.

The **Host** name for the database differs depending on how the instance or the host is protected.

A database that is part of a SQL Server cluster

The **Host** name is the virtual name of the SQL Server cluster.

A database that is part of a SQL Server failover cluster instance (FCI)

The **Host** name is one of the following:

- The cluster name of the FCI
- The physical node names of the FCI

A SQL Server host that uses multiple NICs The **Host** name is one of the following:

- The private interface name of the SQL Server host
- The the private interface name of the virtual SQL Server

- 3 Click **Actions > Recover**.
- 4 On the **Recovery points** tab, locate the full, differential, or transaction log image that you want to restore.
- By default NetBackup uses the primary copy. To select a different copy, click **Copies**.
- See [“Select a different backup copy for recovery”](#) on page 62.
- 5 Click **Actions > Perform complete database recovery**.
- 6 (Conditional) For a transaction log, select one of the following options.

Recovery point selected	Restore the database to the time indicated.
Point in time	Select a different point in time to which you want to restore the database.
Transaction log mark	■ Choose whether to restore at or before the transaction mark. ■ Enter the name of the transaction mark. ■ To select a transaction mark that occurs after a certain date, select After specific date and time. Then specify the date and time. ■ Click Next.

- 7 Select the host, instance, and database for recovery. You have the following options.

Restore to the original host, instance, and database.	
Restore to a different instance.	Type the name in the Instance field.
Select a different host and instance,	Click Change instance .
Restore to a different database.	Type the name in the Database name field.

- 8 Select the path to which you want to restore the database files. You have the following options:

Restore everything to the original directory

Restores all the files to the original directory that was backed up.

Restore everything to a different directory

Restores all the files to the directory that you enter in the **Directory for restore** field.

Restore files to different paths

Restores the individual files to the path that you enter. Click **Edit file paths** and click on any directory path to edit the restore path for that file.

Example of a restore to different paths:

Logical file name	File type	Restore to
modeldev	DATAFILE	E:\Restore\model.mdf
modellog	LOGFILE	E:\Restore-log\modellog.ldf

- 9 Enter the credentials for the recovery target. Or, click **Select existing credentials** to select the credential you want to use.

The user account must be a member of the Windows administrator group and a member of the local SQL Server sysadmin role.

- 10 Click **Next**.

- 11 Select the recovery options.

- For the **Database recovery state after restore**, select **Recover**.
- Choose a **Consistency check** option to perform after the restore.
- Select any other recovery options.

See [“Options for SQL Server restores”](#) on page 60.

- 12 Click **Next**.

- 13 On the **Review** page, review the restore options that you selected.

- At the top, click on the link that follows **Recovery set** to view the backup images that are required for the restore.

- Click **Edit** to change the **Recovery target** settings or **Recovery options**.
- Click **Start recovery**.

Recover a single recovery point

Perform a recovery of a single recovery point when you want to restore backup images in separate restore operations.

To restore to a different server (host), the following requirements exist.

- RBAC permissions to restore to an alternate location.
Refer to the information on configuring host mappings in the [NetBackup Web UI Administrator's guide](#). Or, contact your NetBackup administrator for assistance.
- NetBackup must have the ability to communicate with the destination client.

To recover a single recovery point

- 1 On the left, select **Workloads > Microsoft SQL Server**.
- 2 On the **Databases** tab, locate the database that you want to restore.
The **Host** name for the database differs depending on how the instance or the host is protected.

A database that is part of a SQL Server cluster

The **Host** name is the virtual name of the SQL Server cluster.

A database that is part of a SQL Server failover cluster instance (FCI)

The **Host** name is one of the following:

- The cluster name of the FCI
- The physical node names of the FCI

A SQL Server host that uses multiple NICs

The **Host** name is one of the following:

- The private interface name of the SQL Server host
- The the private interface name of the virtual SQL Server

- 3 Click **Actions > Recover**.
- 4 On the **Recovery points** tab, locate the full, differential, or transaction log that you want to restore.

By default NetBackup uses the primary copy. To select a different copy, click **Copies**.

See [“Select a different backup copy for recovery”](#) on page 62.

- 5 Select **Actions > Restore single recovery point**.

- 6** (Conditional) For a transaction log image, select one of the following options and click **Next**.

Recovery point selected	Restore the database to the time indicated.
Point in time	Select a different point in time to which you want to restore the database.
Transaction log mark	■ Choose whether to restore at or before the transaction mark. ■ Enter the name of the transaction mark. ■ To select a transaction mark that occurs after a certain date, select After specific date and time. Then specify the date and time.

- 7** Select the host, instance, and database for recovery. You have the following options.

Restore to the original host, instance, and database.

Restore to a different instance. Type the name in the **Instance** field.


Select a different host and instance, Click **Change instance**.

Restore to a different database. Type the name in the **Database name** field.

8 Select the path to which you want to restore the database files. You have the following options:

Restore everything to the original directory	Restores all the files to the original directory that was backed up.
Restore everything to a different directory	Restores all the files to the directory that you enter in the Directory for restore field.
Restore files to different paths	Restores the individual files to the path that you enter. Click Edit file paths and click on any directory path to edit the restore path for that file.

Example of a restore to different paths:



9 Enter the credentials for the recovery target. Or, click **Select existing credentials** to select the credential you want to use.

The user account must be a member of the Windows administrator group and a member of the local SQL Server sysadmin role.

10 Click **Next**.

11 Select the recovery options.

- Select the recovery state from the **Database recovery state after restore** options.
- Select the other recovery options.
- If you select the **Recover** option, choose a **Consistency check** option to perform after the restore.

See [“Options for SQL Server restores”](#) on page 60.

12 Click **Next**.

13 On the **Review** page, review the restore options that you selected.

- At the top, click on the link that follows **Recovery set** to view the backup images that are required for the restore.
 - Click **Edit** to change the **Recovery target** settings or **Recovery options**.
 - Click **Start recovery**.
- 14 When the restore completes, continue with the restore of differential incremental or transaction log backups.
- For each intermediate backup, for the **Database recovery state after restore** select **Restoring**.
 - For the final backup image, select **Recovered**.

Options for SQL Server restores

The following options exist when you perform restores of SQL Server.

Table 7-1 Recovery options

Option	Description
Verify backup image but do not restore	NetBackup processes the image for errors, but does not perform a restore. This option does not apply to snapshot images.
Database recovery state after restore	Select the state for the database after the restore. ■ Recover Restore the last image in a restore sequence and make the database ready for use. ■ Restoring Restore an intermediate backup image. The database is left in a loading state so you can restore and apply additional backup images. ■ Standby Create and maintain a standby during a transaction log and database restore. This option requires a standby undo log, which by default is placed in the same directory as the primary datafile. The account that runs the SQL Server service must have full access permission to the <code>SQLStandBy</code> folder.

Table 7-1 Recovery options (*continued*)

Option	Description
Consistency check	The consistency check to perform after the restore. Output from the consistency check is written to the SQL Server client progress log. ■ Do not perform Do not perform consistency checking. ■ Full check, including indexes Include indexes in the consistency check. Any errors are logged. ■ Full check, excluding indexes Exclude indexes from the consistency check. If indexes are not checked, the consistency check runs significantly faster but is not as thorough. Only the data pages and clustered index pages for each user table are included in the consistency check. The consistency of the non-clustered index pages is not checked. ■ Check catalog Check for consistency in and between system tables in the specified database. ■ Physical check only Perform a low overhead check of the physical consistency of the SQL Server database. This option only verifies the integrity of the physical structure of the page and the record headers. It also verifies the consistency between the pages' object ID and index ID and the allocation structures.
Overwrite the existing database	Allows SQL Server to overwrite a database or any database files, if they already exist. If this operation is not available, contact your NetBackup administrator for the necessary RBAC permission.
VDI timeout	Determines the time out interval for SQL Server Virtual Device Interface. The selected interval is applied to backups and restores of databases and of transaction logs. The default value for backups is 300. The default value for restore operations is 600. Range is 300 - 2147483647.

Restore a database (non-administrator users)

Non-administrators can perform restores of SQL Server. However, additional requirements and configuration steps are required.

The following requirements exist for a non-administrator:

- Is a member of the domain users group.
- Has the sysadmin role on the local SQL Server.
- Has full access control to the following:
 - `install_path\NetBackup\dbext\mssql` folder

- HKLM\SOFTWARE\ODBC registry hive
- *install_path\NetBackup\logs\user_ops* folder

Restore a database (non-administrator users)

- 1 Before you can perform a restore, you must do the following:
 - Add the non-administrator credentials to the SQL Server instance.
See [“Select a credential for a SQL Server instance or replica”](#) on page 26.
 - Perform a new backup of the database.
Locate the database and click **Action > Backup Now**.
- 2 When you perform the restore, provide the credentials that you used to register the instance.

Select a different backup copy for recovery

You can restore from the primary backup copy or choose from other available backup copies.

To select a different backup copy for recovery

- 1 Locate the full, differential, or transaction log that you want to restore.
- 2 Click **Copies** and locate the copy that you want.

In the example below, there is an additional copy for the transaction log on **Tape**.

April 30, 2021

Backup images/Recovery points	Backup type	
12:00 PM - 02:00 PM	1 Full, 1 Incremental, 6 Transaction log	
12:11:54 PM	Full	Copies > ⋮
12:26:41 PM	Incremental	Copies v ⋮
Storage	Storage server	Storage type
storage1 (Primary copy)	storageserver1	MSDP ⋮
storage2	storageserver1	AdvancedDisk ⋮
E:\storage3	storageserver1	
/storage4	storageserver2	

Perform complete database recovery
Recover single recovery point

- 3 You can then click the **Actions** menu for that copy to select the restore that you want to perform.

In this example, for the copy on **AdvancedDisk**, you can select either **Perform complete database recovery** or **Recover single recovery point**.

Edit the storage for recovery

In the example below, the **Recovery source** page of the recovery wizard displays the selected storage for recovery. If the images that are needed for recovery are not available on that storage, NetBackup automatically selects the primary images on the appropriate storage. You can change the storage if you don't want to use the automatic selections.

In this case, you selected a transaction log copy on AdvancedDisk storage. Because the full and incremental images were not available on the same storage, NetBackup automatically picked the copies on MSDP storage. You can click **Edit** to change the selected storage for the **Full** image.

Figure 7-1 Storage selected for recovery

Storage for recovery			
Full			
Selected storage	Storage server	Storage type	
 storage1	storageserver1	MSDP	Edit
Incremental			
Selected storage	Storage server	Storage type	
storage2	storageserver2	AdvancedDisk	Edit

If you want to use only the primary copies for the recovery, click **Select only primary copies** (see [Figure 7-2](#)). Otherwise, you can click **Edit** to select the specific storage that you want to use (see [Figure 7-3](#)).

Figure 7-2 Select only primary copies of transaction logs

Storage for recovery			
Full			
Selected storage	Storage server	Storage type	
storage1	storageserver1	MSDP	Edit
Incremental			
Selected storage	Storage server	Storage type	
storage1	storageserver1	MSDP	Edit
Transaction log			
☒ Select only primary copies This option automatically selects only the primary copies for the transaction logs.			
Selected storage	Storage server	Storage type	Images
storage1	storageserver1	MSDP	12 of 24
storage2	storageserver2	AdvancedDisk	12 of 24

Figure 7-3 Edit storage for transaction logs

Storage for recovery

Full

Selected storage	Storage server	Storage type	
storage1	storageserver1	MSDP	Edit

Incremental

Selected storage	Storage server	Storage type	
storage1	storageserver1	MSDP	Edit

Transaction log

☐ Select only primary copies
This option automatically selects only the primary copies for the transaction logs.

Selected storage	Storage server	Storage type	Images	
storage1	storageserver1	MSDP	24 of 24	Edit

Restore a SQL Server availability database to a secondary replica

This procedure describes how to restore a SQL Server availability database to a secondary replica. Follow this procedure if a secondary replica is unavailable for an extended time and needs to be synchronized with the primary. Or follow these instructions after you add a new secondary replica to the availability group.

To restore a SQL Server availability database to a secondary replica

- Log on to the node that hosts the secondary replica and perform the following actions:
 - Close any connections to the database on the secondary replica.
 - Remove the secondary database from the availability group.
- On the left, select **Workloads > Microsoft SQL Server**.
- Click on the **Availability groups** tab and then click on the availability group name.
- On the **Replicas** tab, click on the instance that is hosted on the secondary replica.

- 5 On the **Databases** tab, click on the database that you want to restore.
- 6 Click the **Recovery points** tab and locate the latest transaction log backup.
By default NetBackup uses the primary copy. To select a different copy, click **Copies**.
See [“Select a different backup copy for recovery”](#) on page 62.
- 7 From the **Actions** menu select **Perform complete database recovery**.
- 8 Select one of the following options.

Recovery point selected	Restore the database to the time indicated.
Point in time	Select a different point in time to which you want to restore the database.
Transaction log mark	■ Choose whether to restore at or before the transaction mark. ■ Enter the name of the transaction mark. ■ To select a transaction mark that occurs after a certain date, select After specific date and time. Then specify the date and time.

- 9 If the replicas in the availability group use different paths for the database file, select **Restore files to different paths** and edit the file path.
- 10 Enter the credentials of the instance that you want to restore to and click **Next**.
The user account must be a member of the Windows administrator group and a member of the local SQL Server sysadmin role.
- 11 Select the following settings:
 - **Restoring**
 - **Overwrite existing database**
 See [“Options for SQL Server restores”](#) on page 60.
- 12 Click **Next**. Then click **Start recovery**.
- 13 When the restore completes, join the database to the availability group.

Restore a SQL Server availability database to the primary and the secondary replicas

In some situations you may need to restore the SQL Server availability databases to both the primary and the secondary replicas. These situations can include when you restore databases:

- Following a disaster recovery
- After logical corruption of the databases
- To a clone of an availability group or test environment
- To an earlier point in time

You may want to perform this restore for the primary database in parallel with the restores for the secondary databases.

To restore a SQL Server availability database to the primary and the secondary replicas

- 1 Log on to the host of the primary replica and perform the following actions:
 - In SQL Server Management Studio, suspend data movement on the database and remove the database from the availability group.
 - Close any connections to the database.
 - Remove the primary database from SQL Server.
- 2 In the NetBackup web UI, on the left select **Workloads > Microsoft SQL Server**.
- 3 Click on the **Availability groups** tab and then click on the availability group name.
- 4 On the **Replicas** tab, click on the instance that is hosted on the primary replica.
- 5 On the **Databases** tab, click on the database that you want to restore.
- 6 Click the **Recovery points** tab and locate the latest transaction log backup. By default NetBackup uses the primary copy. To select a different copy, click **Copies**.
See [“Select a different backup copy for recovery”](#) on page 62.
- 7 From the **Actions** menu select **Perform complete database recovery**.
- 8 Select one of the following options.

Recovery point selected	Restore the database to the time indicated.
-------------------------	---

Point in time	Select a different point in time to which you want to restore the database.
Transaction log mark	■ Choose whether to restore at or before the transaction mark. ■ Enter the name of the transaction mark. ■ To select a transaction mark that occurs after a certain date, select After specific date and time. Then specify the date and time.

- 9 Enter the credentials of the instance that you want to restore to and click **Next**.
The user account must be a member of the Windows administrator group and a member of the local SQL Server sysadmin role.
- 10 Select the following settings:
 - **Recover**
 - **Overwrite existing database**
 See [“Options for SQL Server restores”](#) on page 60.
- 11 Click **Next**. Then click **Start recovery**.
- 12 When the restore completes, add the database to the availability group using the **Skip initial data synchronization** option.
- 13 Log on to the host of the secondary replica and complete the following steps:
 - Close any connections to the database on the secondary replica.
 - Remove the secondary database from SQL Server.
- 14 In the NetBackup web UI, on the left select **Workloads > Microsoft SQL Server**.
- 15 Click on the **Availability groups** tab and then click on the availability group name.
- 16 On the **Replicas** tab, click on the instance that is hosted on the secondary replica.
- 17 On the **Databases** tab, click on the database that you want to restore.
- 18 Click the **Recovery points** tab and locate the image that you restored to the primary replica.
- 19 From the **Actions** menu select **Perform complete database recovery**.
- 20 For the transaction log, select the same point in time or log mark that you did for the primary replica.

- 21 Enter the credentials of the instance that you want to restore to and click **Next**.
The user account must be a member of the Windows administrator group and a member of the local SQL Server sysadmin role.
- 22 Select the following settings:
 - **Restoring**
 - **Overwrite existing database**See [“Options for SQL Server restores”](#) on page 60.
- 23 Click **Next**. Then click **Start recovery**.
- 24 When the restore completes, join the database to the availability group.
- 25 Repeat step 13 through step 24 for additional replicas in the availability group.

Using instant access with SQL Server

This chapter includes the following topics:

- [Prerequisites when you configure an instant access SQL Server database](#)
- [Things to consider before you configure an instant access database](#)
- [Configure Samba users for SQL Server instant access](#)
- [Configure an instant access database](#)
- [View the livemount details of an instant access database](#)
- [Delete an instant access database](#)
- [Options for NetBackup for SQL Server instant access](#)
- [NetBackup for SQL Server terms](#)
- [Frequently asked questions](#)

Prerequisites when you configure an instant access SQL Server database

The prerequisites are only applicable to Microsoft SQL Server instant access Build Your Own (BYO).

Prerequisites:

- The BYO server operating system version must be Red Hat Enterprise Linux 7.6 or later.

- Ensure that the Samba service is installed and the Samba share permission is allowed in the selinux policy using the following command


```
setsebool -P samba_export_all_rw=1
```
- The storage server with NGINX installed.
 - The NGINX version must be same as the one in the corresponding official RHEL version release. You need to install it from the corresponding RHEL yum source (epel).
 - Ensure that the `polycoreutils` and `polycoreutils-python` packages are installed from the same RHEL yum source (rhel server). Then run the following commands:
 - `semanage port -a -t http_port_t -p tcp 10087`
 - `setsebool -P httpd_can_network_connect 1`
- Ensure that the `/mnt` folder on the storage server is not mounted by any mount points directly. User mount points should be mounted to its subfolders.
- Enable the logrotate permission in selinux using the following command:


```
semanage permissive -a logrotate_t
```
- Instant access is only supported for SQL Server backup images when the following conditions are met:
 - Snapshots are enabled in the policy or the protection plan.
 - The backup is a full database backup.
 - The primary server, media server, storage server, and client must be at version 8.3 or later.

For instant access using backup copies from cloud LSU (logical storage unit), the primary server and media server must be at version 10.0.1 or later. For more information about limitations of instant access for cloud LSU (logical storage unit), refer to the [NetBackup Deduplication Guide](#).
 - The storage server must be an appliance or BYO that meets the earlier specified prerequisites.

Note: Instant access for incremental and transaction log backups depends on the instant access capability of its base backup image.

Hardware and configuration requirements of instant access

The following hardware requirements exist for the use of instant access.

Table 8-1 Hardware configuration requirement

CPU	Memory	Disk
■ Minimum 2.2-GHz clock rate ■ 64-bit processor ■ Minimum 4 cores; 8 cores recommended. For 64 TBs of storage, the Intel x86-64 architecture requires eight cores.	■ 16 GB (For 8 TB to 32 TB of storage) ■ 1 GB RAM for 1 TB of storage ■ 32 GB of RAM for more than 32 TB of storage ■ An additional 500MB of RAM for each live mount	Disk size depends on the size of your backup. Refer to the hardware requirements for NetBackup and Media Server Deduplication Pool (MSDP).

Additional configuration requirements exist for Windows clients that are in a domain. For example, SQL Server clients that use gMSA. For more information about storage server configuration requirements, see the following article:

https://isearch.veritas.com/internal-search/en_US/article.100051793.html

Things to consider before you configure an instant access database

Note the following about the instant access SQL Server feature:

- The SQL Server backup with the following backup options or scenarios does not support SQL Server instant access:
 - Application-aware backups (VMware)
 - Stream-based backups
 - NetBackup backup compression
 - Legacy SQL Server backups (with batch files)
 - File group or file backups
 - PFI backups (backup option: **Retain snapshot for Instant Recovery or SLP management**)
 - SQL Server database mirroring (only support is to create as a standalone IA database)
 - SQL Server clusters (only support is to create as a standalone IA database)
- Instant access on Flex WORM storage requires the following services:
 - NGINX, NFS, SAMBA, WINBIND (if Active directory is required), SPWS, VPFS

- Instant access does not support a restore of a filestream database. Restore the entire VM without instant access. Or restore the database without instant access. For details see the following article:
<https://www.veritas.com/docs/100048546>
- For instant access to work after an upgrade of the storage and the primary server from an earlier NetBackup version, restart NetBackup web service on the upgraded primary server with the following commands:
 - `/usr/opensv/netbackup/bin/nbwmc stop`
 - `/usr/opensv/netbackup/bin/nbwmc start`

Configure Samba users for SQL Server instant access

You must first configure Samba users for SQL Server instant access on the corresponding storage server and enter the credentials on the client.

If the Samba service on a storage server is part of Windows domain, the Windows domain users can be used for Samba share. In this scenario, credentials are not required to access the share.

For Azure Kubernetes Service (AKS) and Amazon Elastic Kubernetes Service (EKS) cloud platforms, only Samba local user can access Samba share. You must add Samba users to access the Samba share.

The following table describes how to add or manage Samba users if the Samba service is not part of Windows domain.

Table 8-2 Steps to add or manage Samba users

User	Steps
For NetBackup Appliance users	For NetBackup Appliance, local users are also Samba users. To manage local users, logon to CLISH and select Main > Settings > Security > Authentication > LocalUser. The Samba password is the same as the appliance local user's logon password.

Table 8-2 Steps to add or manage Samba users (*continued*)

User	Steps
For Flex Appliance users	For a Flex Appliance application instance, log in to the instance and add any local user to Samba, as follows: ◆ If you want, create a new local user with the following commands: ■ <code>#useradd <username></code> ■ <code>#passwd <username></code> You can also use an existing local user. ◆ Run the following commands to create user credentials for Samba and enable the user: ■ <code>smbpasswd -a <username></code> ■ <code>smbpasswd -e <username></code>
For Build Your Own (BYO) users	For new users: 1 Create a Linux user, then add the user to Samba. For example, the following commands create a <code>test_samba_user</code> for Samba service only. <pre># adduser --no-create-home -s /sbin/nologin test_samba_user</pre> <pre># smbpasswd -a test_samba_user</pre> 2 Enter a new SMB password. 3 Enter the new SMB password again. The new user is added. For existing users: If you want to add an existing user to the Samba service, run the following command: <code>smbpasswd -a test_samba_user</code>

Table 8-2 Steps to add or manage Samba users (*continued*)

User	Steps
For AKS and EKS platform users	For new users: 1 Log in to the MSDP engine pod in a cluster using <code>kubectl</code>. 2 Run the following command to log in to <code>rshell</code> in the MSDP engine. <code>su - msdpadm</code> 3 Run the following <code>rshell</code> command to add a Samba user. <code>setting samba add-user username=[samba user name] password=[samba password]</code> For example, <code>msdp-16.1] > setting samba add-user username=test_samba_user password=Te@Pss1fg0</code> You can use the same command to update the password for an existing user. In AKS and EKS cloud platforms, the Samba <code>rshell</code> command configures Samba servers in all MSDP engines in a cluster.

To automatically start the SQL Server database, ensure that you can access the share when you log on with the instance credentials from the web UI.

For the cloud platforms such as AKS and EKS, add the Samba user and each MSDP engine host name in Windows credential manager. This action allows the NetBackup client can connect to the Instant Access Samba share automatically.

Configure an instant access database

When you configure an instant access database, you can choose to add the database automatically to the instance. Or, you can export the database to a Samba share.

Configure an instant access database and then start the database

To configure an instant access database and automatically add the database to the instance, you can use a full, incremental, or transaction log backup.

To configure an instant access database and then start the database

- 1** On the left, click **Microsoft SQL Server**.
- 2** On the **Databases** tab, click the database for which you want to configure the instant access database.
- 3** Click the **Recovery points** tab, then click the date on which the backup occurred.

The available images appear in rows with the backup timestamp for each image.
- 4** Right-click on the backup image and click **Actions > Configure instant access**.
- 5** (Conditional) For a full backup, after the instant access database is created you can add the database to the instance and start the database. Click **Yes > Next** for this option.
- 6** (Conditional) For a transaction log, select a replay option and click **Next**.
- 7** Review the recovery target and host name, instance name and make any wanted changes.

To change the host and instance, click **Change instance**.
- 8** In the **Database name** field, enter the instant access database name that you want to create.
- 9** Enter the credentials for the recovery target. Or, click **Select existing credentials** to select the credential you want to use.

The user account must be a member of the Windows administrator group and a member of the local SQL Server sysadmin role.
- 10** Click **Next**.
- 11** Review the recovery options and make changes if needed and then click **Next**.

See [“Options for NetBackup for SQL Server instant access”](#) on page 78.
- 12** (Optional) To view a list of the backup images for the selected recovery point, click the link that displays the number of backup images.
- 13** Review the summary of the selected recovery target and recovery options. Then click **Start recovery**.
- 14** After the instant access job starts, you can click on the **Restore activity** tab to view the progress.

See [“View the livemount details of an instant access database”](#) on page 77.

Configure an instant access database, but do not start the database

To configure an instant database and export the database to Samba share, you must use a full backup.

To configure an instant access database, but not start the database

- 1 On the left, click **Microsoft SQL Server**.
- 2 On the **Databases** tab, click the database for which you want to configure the instant access database.
- 3 Click the **Recovery points** tab, then click the date on which the backup occurred.

The available images appear in rows with the backup timestamp for each image.
- 4 Right-click on the backup image and click **Actions > Configure instant access**.
- 5 If you want to add the database to the instance and start the database, choose **No > Next**.
- 6 Select one of the following options for the recovery target:
 - To enter the recovery target host name, click **Enter host name**.
 - To select from a list of hosts, click **Select host name**
- 7 (Optional) To view a list of the backup images for the selected recovery point, click the link that displays the number of backup images.
- 8 Click **Start recovery**.
- 9 After the instant access job starts, you can click on the **Restore activity** tab to view the progress.

See [“View the livemount details of an instant access database”](#) on page 77.

View the livemount details of an instant access database

You can view the livemount details of an instant access database.

To view the livemount details of an instant access database

- 1 On the left, click **Microsoft SQL Server**.
- 2 Click the **Instant access databases** tab.
- 3 On the **Instant Access databases** tab, click the database for which you want to see the livemount details.

Mount ID	Unique ID for an instant access livemount.
Export path	Exported instant access livemount path from the storage server.
Recovery point ID	Unique ID of a recovery point.
Livemount path	UNC path of the instant access livemount on the Microsoft SQL client.
Export server	Server where the livemount share is exported from.

Delete an instant access database

You can delete an instant access database that may or may not be added to an instance.

To delete an instant access database

- 1 On the left, click **Microsoft SQL Server**.
- 2 Click the **Instant access databases** tab.
The tab lists the names of the configured instant access databases.
- 3 Select **Actions > Delete**.
- 4 Perform one of the following:
 - Your instant access database is added to an instance and is started.
Enter the SQL Server instance credentials and then click **Delete**.
 - Your instant access database is not added to an instance and is not started.
If you are sure that you want to delete the database, then click **Delete**.

Options for NetBackup for SQL Server instant access

The table describes the recovery options that are available when you perform instant access.

Table 8-3 Recovery options

Option	Description
Database recovery state after restore	Select the state for the database after the restore. ■ Recover Restore the last image in a restore sequence and make the database ready for use. ■ Restoring Restore an intermediate backup image. The database is left in a loading state so you can restore and apply additional backup images. ■ Standby Create and maintain a standby during a transaction log and database restore. This option requires a standby undo log, which by default is placed in the same directory as the primary datafile. The account that runs the SQL Server service must have full access permission to the <code>SQLStandBy</code> folder.
Consistency check	The consistency check to perform after the restore. Output from the consistency check is written to the SQL Server client progress log. ■ Do not perform Do not perform consistency checking. ■ Full check, including indexes Include indexes in the consistency check. Any errors are logged. ■ Full check, excluding indexes Exclude indexes from the consistency check. If indexes are not selected, the consistency check runs significantly faster but is not as thorough. Only the data pages and clustered index pages for each user table are included in the consistency check. The consistency of the non-clustered index pages is not selected. ■ Check catalog Check for consistency in and between system tables in the specified database. ■ Physical check only Perform a low overhead check of the physical consistency of the SQL Server database. This option only verifies the integrity of the physical structure of the page and the record headers. It also verifies the consistency between the pages' object ID and index ID and the allocation structures.
VDI timeout	Determines the time-out interval for SQL Server Virtual Device Interface. The selected interval is applied to backups and restores of databases and of transaction logs. The default value for backups is 300. The default value for restore operations is 600. Range is 300 - 2147483647.

See [“Configure an instant access database”](#) on page 75.

NetBackup for SQL Server terms

The table describes the important terms that might be new to a SQL Server database administrator or a NetBackup administrator.

Table 8-4 NetBackup for SQL Server terms

Term	Definition
Full backup	A complete backup of the database that contains all of the data files and the log file. (Note that a full backup does not truncate the transaction log.)
Incremental backup	A backup of the changed blocks since the last full backup.
Transaction log	An ongoing record of updates that were made to a database.
Transaction log backup	Backs up the transactions that have occurred since the last transaction log backup. After a successful backup, the log is cleared so that new transactions can be written to the file. A transaction log backup can only be performed against a database that is configured to run in the full recovery model.
Restore	To copy data back to a SQL Server object.
Recovery	To bring a database online as a result of a restore.
SQL Server host	The host machine on which SQL Server resides. It may also refer to the virtual name of a cluster that supports a SQL Server installation.
SQL Server instance	A SQL Server installation. If an instance is not specified, it is considered the default SQL instance for the SQL host.

Frequently asked questions

Here are some frequently asked questions for Microsoft SQL instant access Build Your Own (BYO).

Table 8-5

Applicable for	Frequently asked question	Answer
BYO	How can I enable the Microsoft SQL instant access feature on BYO after storage is configured or upgraded without the nginx service installed?	Perform the steps in the following order: 1 Install the required nginx service version. 2 Ensure that the new BYO nginx configuration entry: <code>/etc/nginx/conf.d/byo.conf</code> is part of the HTTP section of the original: <code>/etc/nginx/nginx.conf</code> file. 3 Run the command: <code>/usr/openv/pdde/vpfs/bin/vpfs_config.sh --configure_byo</code>
BYO	How can I resolve the following issue in the <code>vpfs-config.log</code> file that is raised from: Verifying that the MSDP REST API is available via https on port 10087	Perform the steps in the following order: 1 Install the <code>polycoreutils</code> and <code>polycoreutils-python</code> packages through yum tool. 2 Add the following rules that SELinux for Nginx requires to bind on the 10087 port. ■ <code>semanage port -a -t http_port_t -p tcp 10087</code> ■ <code>setsebool -P httpd_can_network_connect 1</code> 3 Run the following command: <code>/usr/openv/pdde/vpfs/bin/vpfs_config.sh --configure_byo</code>

Table 8-5 (continued)

Applicable for	Frequently asked question	Answer
BYO	Instant Access for BYO uses a self-signed certificate by default and only supports *.pem external certificate. How do I replace it with a certificate signed by external CA (*.pem certificate), if required?	To configure the external certificate, perform the following steps. If the new certificate is already generated (the certificate must contain long and short host names for the media server), go to step 4. 1 Create the RSA public or private key pair. 2 Create a certificate signing request (CSR). The certificate must contain long and short host names for the media server. 3 The External Certificate Authority creates the certificate. 4 Replace <PDDE Storage Path>/spws/var/keys/spws.cert with the certificate and replace <PDDE Storage Path>/spws/var/keys/spws.key with the private key. 5 Run the following command to reload the certificate: /usr/openv/pdde/vpfs/bin/vpfs_config.sh --configure_byo
BYO	How can I disable media automount for the instant access livemount share in gnome? If the automount is enabled, the source folder is mounted from the livemount share in gnome and smaller disks appear. In this scenario, the instant access feature does not work properly. The mounted disk content source is from the .../meta_bdev_dir/... folder under livemount share, while the mount target is in the /run/media/... folder.	Follow the guideline to disable the gnome automount: https://access.redhat.com/solutions/20107

Table 8-5 (continued)

Applicable for	Frequently asked question	Answer
BYO	How can I resolve the following issue in the /var/log/vpfs/vpfs-config.log file? <pre>**** Asking the NetBackup Webservice to trust the MSDP webserver (spws) **** /usr/opensv/netbackup/bin/nblibcurlcmd failed (1):</pre>	Perform the steps in the following order: 1 Ensure that your NetBackup primary server is up and there is no firewall blocking the connection between the NetBackup primary server and storage server. 2 Run the following command on storage server to verify the connection status: <pre>/usr/opensv/netbackup/bin/bpclntcmd -pn</pre> 3 After the NetBackup primary server is up and connection between the NetBackup primary server and storage server is allowed, run the following command: <pre>/usr/opensv/pdde/vpfs/bin/vpfs_config.sh --configure_byo</pre>
BYO and Flex Appliance	How can I enable host-based authentication and secure logon for Samba share so that the MSSQL instant access works on specific windows clients? The clients windows version and the background are listed at the following link: https://support.microsoft.com/en-us/help/4046019/guest-access-in-smb2-disabled-by-default-in-windows-10-and-windows-ser	Perform the steps in the following order: 1 In the storage server (one time operation) where the Samba share is exported from: ■ Override the following Samba option to disable the guest logon: <pre>map to guest = Never</pre> ■ Create user credentials for Samba. ■ <code>smbpasswd -a spws</code> Set Samba password for Samba user spws ■ <code>smbpasswd -e spws</code> Enable Samba user spws 2 For each Windows client, where the Samba share is accessed using the earlier credentials, save the spws credentials in the credential manager.

Table 8-5 (continued)

Applicable for	Frequently asked question	Answer
NetBackup Appliance	How can I enable host-based authentication and secure logon for Samba share so that the MSSQL instant access works on NetBackup Appliance and windows clients?	Perform the steps in the following order: 1 In the storage server (one time operation) where the Samba share is exported from, create new local user credentials for Samba with the following Appliance CLISH path: Main_Menu > Settings > Security > Authentication > LocalUser 2 In each Windows client, where the Samba share is accessed using the earlier credentials, save the new local user credentials in the credential manager.

Protecting SQL Server with VMware backups

This chapter includes the following topics:

- [About protecting an application database with VMware backups](#)
- [Create a protection plan to protect SQL Server data with a VMware backup](#)
- [Protect SQL Server data with a VMware backup](#)
- [Restore SQL Server databases from a VMware backup](#)

About protecting an application database with VMware backups

With a VMware backup policy and the Veritas VSS provider, NetBackup can create consistent, full backups of an application database that resides on a virtual machine.

VMware application backups let you:

- Use the existing database restore process to restore and recover data from VMware backups.
- Restore and recover databases from VMware backups to alternate clients. The target destination client can be a physical computer or a virtual machine.

Supported environments and configuration

See the following information on virtual systems compatibility:

https://www.veritas.com/content/support/en_US/doc/NB_70_80_VE

Veritas VSS provider

Veritas recommends the Veritas VSS provider. VMware Tools calls the provider to quiesce the VSS writers for a file-level consistent backup. Without this VSS provider (or the VMware VSS Provider), database recovery may require manual steps and granular recovery is not supported.

See [“Installing the Veritas VSS provider for vSphere”](#) on page 14.

The Veritas VSS provider allows VMware backups that truncate the logs on SQL Server virtual machines. The Veritas VSS provider truncates the logs by means of full VSS backups. Note that the VMware VSS provider creates copy-only backups, which cannot be used as a basis to truncate logs.

Limitations of VMware application backups

Databases are cataloged and protected only for the configurations that are supported for VMware backups. Make sure to store databases and transaction logs on supported storage.

VMware application backups do not support the following policy options and configurations:

- Incremental backups. Instead, you can create a protection plan for SQL Server incremental backups.
- SQL Server clusters or SQL Server availability groups.
- Restores from a non-primary copy. You can only restore from the primary copy. Only the primary copy is displayed for restore, even if there are other copies. If you want to restore from another copy, promote that copy to the primary copy.
- SQL Server databases are not cataloged and backed up if they exist on the following:
 - Any virtual machines that use raw device mapping (RDM).
 - Virtual Machine Disk (vmdk) volumes that are marked as independent.
 - Mount points that use MBR disks. Mount points that contain SQL Server database files are only supported when the underlying disk is a GPT disk.
 - Virtual hard disks (VHDs).
 - RAID volumes.
 - ReFS file systems.
 - An excluded Windows boot disk.

Create a protection plan to protect SQL Server data with a VMware backup

A VMware protection plan can protect SQL Server data on a virtual machine. Users can then subscribe assets to that protection plan. Note that before you create a protection plan, you must perform additional configuration requirements:

- Configure all storage options.
- Configure the logon account for the NetBackup services.
 See [“Configuring the NetBackup services for SQL Server backups and restores”](#) on page 14.
 See [“Configure local security privileges for SQL Server”](#) on page 16.
- Review the auto-discovered mappings for the hosts in your environment.
 This action requires the Default Security Administrator role or a role with similar RBAC permissions.

To create a protection plan to protect SQL Server data with a VMware backup

- 1 Configure the storage for the backup.
- 2 On the left, select **Protection > Protection plans** and then click **Add**.
- 3 In **Basic properties**, enter a **Name**, **Description**.
- 4 From the **Workload** list, select **VMware**.
- 5 (Optional) Indicate a **Policy name prefix** to append to the policy name.
 NetBackup automatically creates a policy when users subscribe assets to this protection plan.
- 6 In **Schedules and retention**, click **Add schedule**.
 - In the **Attributes** tab, select the **Full** backup type.
 - In the **Start window** tab, define the window during which the backup can start.
 - Click **Save** after all options are selected in the **Attributes** and the **Start window** tabs.
 - Review the **Backup schedule preview** window and verify that all schedules are set correctly.

See [“Schedules”](#) on page 34.

7 In **Storage options**, select the storage to use for the backup.

A protection plan can only use the storage that a NetBackup 8.1.2 or newer media server can access.

Storage option	Requirements	Description
Backup storage	OpenStorage is required for this option. Tape, storage unit groups, and Replication Director not supported.	Click Edit . Select the storage target then click Use selected storage .

8 In **Backup options**, review the available options for the backup.

See [“Backup options and Advanced options”](#) on page 88.

9 Under **Allow restore of application data from virtual machine backups**, select **Microsoft SQL Server**

(Optional) Select **Truncate logs** if you want to truncate the transaction logs when the VMware snapshot of the virtual machine is complete.

10 In **Permissions**, review the roles that have access to protection plans.

To give another role access to this protection plan, click **Add**. Select the **Role** in the table and customize the role by adding or removing permissions in the **Select permissions** section.

See [Configure RBAC](#).

11 In **Review**, verify that the protection plan details are correct and click **Save**.

Backup options and Advanced options

The user can adjust the following settings when subscribing to a protection plan.

Backup options

Table 9-1 Backup options for protection plans

Option	Description
Select server or host to use for backups	The host that performs backups on behalf of the virtual machines. Users can choose Automatic to have NetBackup pick the media server, based on the storage unit. Or, the user can select another host from the list. These hosts are other media servers in the environment or hosts that are configured as an access host.

Table 9-1 Backup options for protection plans (*continued*)

Option	Description
If a snapshot exists, perform the following action	Specifies the action that NetBackup takes when a snapshot is discovered before NetBackup creates a new snapshot for the virtual machine backup. For example, users can choose to stop a backup if any snapshots exist. If snapshots are not automatically deleted, the performance of the virtual machine may eventually decline. Undeleted snapshots can cause restore failures due to lack of disk space.
Exclude selected virtual disks from backups	Specifies the virtual disks to exclude from backups. See “Exclude disks from backups” on page 89.

Advanced options

Table 9-2 Advanced options for protection plans

Option	Description
Enable virtual machine quiesce	By default, I/O on the virtual machine is quiesced before NetBackup creates the snapshot. In the majority of cases, you should use this default. Without quiescing file activity, data consistency in the snapshot cannot be guaranteed. If you disable the quiesce, you must analyze the backup data for consistency.
Allow the restore of application data from virtual machine backups	This option allows users to restore application data from full backups of the virtual machine. Note that in NetBackup 8.3 or earlier, application data for Microsoft Exchange Server or Microsoft SharePoint Server must be restored with the NetBackup Backup, Archive, and Restore interface. Data for Microsoft SQL Server must be restored with the NetBackup MS SQL Client. See the documentation for your NetBackup database agent for more details.
Transport mode	Specifies the transport mode to use for backups or how to read the data from the datastore. For more information on transport modes, see the vendor documentation for your virtualization environment.
Snapshot retry options	See “Snapshot retry options” on page 90.

Exclude disks from backups

Excluding virtual disks can reduce the size of the backup, but use these options carefully. They are intended only for the virtual machines that have multiple virtual disks.

Table 9-3 Options for excluding virtual disks

Exclude option	Description
All boot disks	Consider this option if you have another means of recreating the boot disk. The virtual machine's boot disk is not included in the backup. Any other disks are backed up. Note: Data files are available in the restored data disks. However, you cannot start a virtual machine that is restored from this backup.
All data disks	Consider this option only if you have a separate protection plan that backs up the data disks. The virtual machine's data disks are not included in the backup. Only the boot disk is backed up. Note: When the virtual machine is restored from the backup, the virtual machine data for the data disk may be missing or incomplete.
Exclude disks based on a custom attribute	Use this option to allow the VMware administrator to use a custom attribute to control which disks are excluded from backups. The attribute must have comma-separated values of device controllers for the disks to be excluded. For example: <code>scsi0-0, ide0-0, sata0-0, nvme0-0</code>. The default value for this attribute is <code>NB_DISK_EXCLUDE_DISK</code>. Or, you can choose your own value. If you add disks to the custom attribute value between any differential backups, those disks are excluded from the next backup. The VMware administrator must use a VMware interface to apply the attribute to the disks to exclude. See the NetBackup Plug-in for VMware vSphere Web Client Guide or the NetBackup Plug-in for VMware vSphere Client (HTML5) Guide.
Specific disks to be excluded	Use this option to exclude a specific disk by the disk type, controller, and LUN that represent the virtual device node of the disk. Click Add to specify additional disks. If you add controllers between any differential backups, their disks are excluded from the next backup.

Snapshot retry options

For most environments, the default values for the snapshot retry options are appropriate. It may be helpful to adjust these settings based on the size of the virtual machine and the processing load on the VMware server.

Table 9-4 Snapshot retry options

Option	Description
Maximum number of times to retry a snapshot	The number of times the snapshot is retried.

Table 9-4 Snapshot retry options (*continued*)

Option	Description
Maximum length of time to complete a snapshot	The time, in minutes, to allow the snapshot operation to complete. If snapshots do not complete, set this option to a specific period to force a time-out. Use the Maximum length of time to wait before a snapshot is retried setting to retry the snapshot at a later time.
Maximum length of time to wait before a snapshot is retried	The time to wait (in seconds) before the snapshot is retried.

Protect SQL Server data with a VMware backup

Use the following procedure to subscribe a VM that contains SQL Server data to a protection plan. When you subscribe an asset to a protection plan, you assign predefined backup settings to the asset.

Note: The RBAC role that is assigned to you must give you access to the assets that you want to manage and to the protection plans that you want to use.

To protect SQL Server data with a VMware backup

- 1 On the left, click **VMware**.
- 2 On the **Virtual machines** tab or **Intelligent VM groups** tab, click the box for the VM or the VM group and click **Add protection**.
- 3 Select a protection plan and click **Next**.
- 4 You can adjust one or more of the following settings:
 - **Schedules and retention**
Change when backups occur and the backup start window.
See “[Schedules](#)” on page 34.
 - **Backup options**
Adjust the server or host to use for backups, snapshot options, and exclude options.
See “[Backup options and Advanced options](#)” on page 88.
 - **Advanced options**
Change or enable any advanced options for the protection plan.
See “[Backup options and Advanced options](#)” on page 88.
The plan must allow for restores of SQL Server databases from a VMware image. **Microsoft SQL Server** must be enabled under **Allow restore of**

application data from virtual machine backups. If you also want the backup to truncate logs, select **Truncate logs**.

- 5 Click **Protect**.

The results of your choices appear under **Virtual machines** or **Intelligent VM groups**.

Restore SQL Server databases from a VMware backup

The following steps describe how to restore a SQL Server database from a full VMware backup.

To restore a SQL Server database from a VMware backup

- 1 On the right, click **Workloads > Microsoft SQL Server**.
- 2 Click the **Databases** tab.
- 3 Select the database that you want to restore and click **Recover**.
- 4 Select the date that the backup was performed.
- 5 On the right, locate the recovery point. Then click **Actions > Recover single recovery point**.

Note: Even if multiple copies exist, only the primary copy is available for restore. If you want to restore from another copy, you must first promote that copy to the primary copy.

Performance and troubleshooting

This chapter includes the following topics:

- [NetBackup for SQL Server performance factors](#)
- [Troubleshooting credential validation](#)
- [Troubleshooting VMware backups](#)
- [SQL Server log truncation failure during VMware backups of SQL Server](#)

NetBackup for SQL Server performance factors

Many factors can influence the backup performance, including your hardware environment and the settings in SQL Server and NetBackup.

Note: Some of the factors are only applicable to SQL Server stream-based operations and have no effect on snapshot backups or restores.

For a SQL Server Intelligent policy, set these parameters in the policy, on the **Microsoft SQL Server** tab. For a backup batch file (legacy SQL Server policy) or for a restore batch file, configure these parameters in the NetBackup MS SQL Client interface. The parameters in the NetBackup client properties are saved for the session.

SQL Server buffer space parameters

The **Maximum transfer size**, **Backup block size**, and **Client buffers per stripe** can increase buffer space in SQL Server. SQL Server must have the available

resources to support the increase of these values. Buffer space parameters are applicable for stream-based backups only.

The **Maximum transfer size** parameter can be set for each backup or restore operation. **Maximum transfer size** is the buffer size used by SQL Server for reading and writing backup images. Generally, you can get better SQL Server performance by using a larger value.

The **Backup block size** parameter can be set for each backup operation. For restore operations, NetBackup automatically chooses the same size that was used for the backup. **Backup block size** is the incremental size that SQL Server uses for reading and writing backup images.

The **Client buffers per stripe** determines how many buffers to allocate for reading or writing each data stream during a backup or restore operation. Setting this factor to a value greater than **1** enables multi-buffer during data transfer. By allocating a greater number of buffers, you can affect how quickly NetBackup can send data to the NetBackup media server. Multi-buffer prevents short-term producer-consumer imbalances during a backup or restore operation. Although you can set the number of buffers as high as **32**, normally a value of **2** or **3** is sufficient.

Stripes and parallel backup operations

You can improve performance and throughput by increasing the backup stripes or parallel backup operations, depending on the size and number of databases.

Multiple stripes (**Number of backup stripes**) are useful for larger databases when the performance gains outweigh the additional overhead for the SQL Server agent to configure them. For smaller databases, striping can decrease performance speed. In general, if the SQL Server instance only has a few large databases, the use of stripes improves performance. If the instance has numerous smaller databases, increasing the amount of **Parallel backup operations** is a better choice to improve performance. You can increase both stripes and parallel backup operations at the same time, but be careful not to overwhelm the system resources.

See [“Configure the number of jobs allowed for backup operations”](#) on page 18.

Caution: Do not enable multiplexing if the policy is also configured with multiple stripes. Restores fail when both multiplexing and multiple stripes are configured for a backup policy.

Shared memory usage

For optimal performance, install NetBackup server on the same host as NetBackup for SQL Server. Also use shared memory for data transfer instead of sockets. Shared memory is the default unless you create a `install_path\NetBackup\NOSHM` file.

Alternate buffer method

NetBackup for SQL Server supports an alternate buffer method. It optimizes CPU usage by allowing NetBackup and SQL Server to share the same memory buffers without transferring data between them.

The alternate buffer method for backup and restore typically does not improve data transfer rate, only CPU utilization. A situation may occur in which the transfer rate is significantly degraded when alternate buffer method is in use. To improve the transfer rate set the **Maximum transfer size** for the backup to the maximum allowed, which is 4 MB.

About alternate buffer method with backup operations

This method is chosen automatically for backups if all of the following conditions apply:

- NetBackup shared memory is in use.
- The backup is stream-based.
- The backup is not multiplexed.
- The backup policy does not specify either NetBackup compression or NetBackup encryption.
- The NetBackup buffer size equals the SQL Server block size.
The default NetBackup buffer size is 64 KB, but this value can be overridden in the following settings:

`install_path\NetBackup\db\config\SIZE_DATA_BUFFERS` (for tape backups),
or,
`install_path\NetBackup\db\config\SIZE_DATA_BUFFERS_DISK` (for disk backups)

- NetBackup for SQL Server agent is started with the same account as the NetBackup Client Service.

The backups that are initiated from an automatic backup policy are started with the NetBackup Client Service so the same account is already in use. However, you can start a SQL Server backup through NetBackup for SQL Server or through `dbbackupex`. In this case, your logon account must be the same as the NetBackup Client Service account. Then your backups can be candidates for the alternate buffer method.

About alternate buffer method with restore operations

Conditions for backups require that you use the alternate buffer method. Restores also require that backups have been made with the alternate buffer method. You can verify that the alternate buffer method was used. Look for the words `Using`

`alternate buffer method`, which appear in the `dbclient` log and the progress report.

SQL Server checksum

You can choose to perform a checksum before you perform a backup. When this option is enabled, it imposes a performance penalty on a backup or restore operation.

For legacy backup policies, set the **Page verification** value when you create the script. For restore scripts, choose **Verify backup image, but don't restore** option when you create the script.

Instant data file initialization

When you restore a database, filegroup, or database file, SQL Server zeroes the file space before it begins the restore operation. This action can slow the total recovery time by as much as a factor of 2. To eliminate file initialization, run the MSSQLSERVER service under a Windows account that has been assigned the `SE_MANAGE_VOLUME_NAME`. For more information, see the SQL Server and the Windows documentation.

Troubleshooting credential validation

[Table 10-1](#) describes the reasons that validation can fail for an instance, replica, or instance group.

Table 10-1 Reasons for credential validation failure

Status code or error	Description	Explanation
40	Could not validate credentials. Failed to connect to client: <client>.	The host name is invalid.
46	The validation operation timed out waiting for a response from the client	You cannot connect to the host because the host is down.

Table 10-1 Reasons for credential validation failure (*continued*)

Status code or error	Description	Explanation
41	Validation of operating system user/password failed for client: <client>.	- The host name is correct, but the user name or password is invalid. - The credentials use have the setting Use these specific credentials, but the user account does not have the required the local security privileges Impersonate a client after authentication and Replace a process level token. See “Configure local security privileges for SQL Server” on page 16.
1939	The specified user does not have SQL Server System Administrator privileges.	The credentials do not have the “sysadmin” role and the validation fails.
Invalid configuration detected.	Invalid configuration detected. The service user for the NetBackup Client and NetBackup Legacy Network services must be the same user. Change the service users in the Windows Service Manager and try again.	The NetBackup Client Service or the NetBackup Legacy Network Service requires but does not use the same user for the logon account. See “Configuring the NetBackup services for SQL Server backups and restores” on page 14.

Troubleshooting VMware backups

Note the following when you perform a VMware backup that protects an application:

- The Application State Capture (ASC) job contacts the NetBackup client on the guest virtual machine and catalogs the application data for recovery.
- One ASC job is created per VM, regardless of which applications are selected in policy.
- ASC messages are filtered based on the ASC job details in the Activity Monitor.
- Failure results in the discovery job or parent job exiting with status 1.
- If you enable recovery for a particular application but that application does not exist on the VM, the ASC job returns Status 0.
- `bpfis` is run and simulates a VSS snapshot backup. This simulation is required to gain logical information of the application.

Table 10-2 Issues with using a VMware policy to protect databases

Issue	Explanation
A database backup fails.	Databases are cataloged and protected only if the configuration is supported for VMware backups. See “Limitations of VMware application backups” on page 86. NetBackup is installed on an excluded Windows boot disk. The ASC job detects this type of disk and treats it like an independent disk. Do not select the Exclude boot disk option if NetBackup is installed on the boot drive (typically C:).
ASC job produces a status 1 (partially successful).	You selected databases for backup that exist on both supported and on unsupported disks. See “A database backup fails” for unsupported disk information. Full-text catalog files exist on the mounted folders. The databases are not cataloged.
The Application State Capture (ASC) job fails and the databases are not protected.	When the ASC job fails, the VMware snapshot or backup continues. Application-specific data cannot be restored. When you query the SQL Server Management Studio (SSMS), it may show that the database was backed up. In this case, though the database was skipped, the snapshot was still successful. You disabled the Virtual Machine quiesce option. Database objects are on a VHD disk. No objects in the backup are not cataloged, including those that do not exist on the VHD. You excluded any data disks from the VMware policy, on the Exclude Disks tab. Be sure that any disks that you exclude do not contain database data. The VMware disk layout has changed since the last discovery. In this situation, you must force NetBackup to rediscover virtual machines by lowering the value of the Reuse VM selection query results for option. See the NetBackup for VMware Administrator's Guide. You cannot use a VMware incremental policy to protect SQL Server. However, the VMware backup job is successful.
You can recover the entire virtual machine from the backup, but you cannot recover the databases individually.	You did not select Microsoft SQL Server, which allows recovery of the databases from the virtual machine backups.

Table 10-2 Issues with using a VMware policy to protect databases
(continued)

Issue	Explanation
The databases are not quiesced.	Neither the Veritas VSS provider nor the VMware VSS Provider were installed at the time of backup. In this case, the recovery of a database after it is restored may require manual steps.

SQL Server log truncation failure during VMware backups of SQL Server

SQL Server transaction log truncation may fail during VMware backups of SQL Server if a database name contains special characters or if the %TEMP% directory path is too long. During SQL Server log truncation, the NetBackup for SQL Server agent creates a temporary log backup. This backup specifies the current user's configured %TEMP% directory and database name as part of the destination backup device. SQL Server limits the path that can be used for backup devices to 259 characters. Under certain circumstances the SQL Server agent may generate a backup device that is longer than 259 character and cause log truncation to fail.

The following conditions cause failure:

- A configured %TEMP% directory that is longer than 259 characters.
- When the combined length of the database name and %TEMP% directory path is longer than 259 characters.

One workaround for this issue is to configure the %TEMP% directory so that the path is substantially less than 259 characters long.