

NetBackup™ for Microsoft SQL Server Administrator's Guide

Release 11.2

NetBackup™ Microsoft SQL Server Administrator's Guide

Last updated: 2026-09-07

Legal Notice

Copyright © 2026 Cohesity, Inc. All rights reserved.

© 2026 Cohesity, Inc. All Rights Reserved. Cohesity, the Cohesity Logo and other Cohesity Marks are trademarks of Cohesity, Inc. in the US and/or internationally. The information supplied herein is the confidential and proprietary information of Cohesity and may only be used (a) by the intended recipients and (b) in conjunction with validly licensed Cohesity software and services. Find the terms of Cohesity licenses at www.cohesity.com/agreements.

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. COHESITY 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.

Cohesity Support

Reach Cohesity Support

There are several ways to create a Cohesity support case.

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

Support/Service Assistance

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

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

Cohesity Software Running on Partner Hardware

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

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

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

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

Contents

Chapter 1	About NetBackup for SQL Server	13
	Overview of NetBackup for SQL Server	13
	NetBackup for SQL Server features	14
Chapter 2	Installation	17
	Planning the installation of NetBackup for SQL Server	17
	NetBackup server and client requirements	18
	Requirements for using NetBackup for SQL Server in a NetBackup cluster	19
	License for NetBackup for SQL Server	20
Chapter 3	Host configuration and job settings	21
	Configuring SQL Server hosts and user permissions	21
	Installing the VSS provider for vSphere	23
	Disable the SQL Server VSS Writer service	23
	Configure the NetBackup services for SQL Server backups and restores	24
	Configure local security privileges for SQL Server	25
	Reviewing the auto-discovered mappings	26
	Configuring mappings for restores of a distributed applications, clusters, or virtual machines	31
	Configure the ODBC connection	32
	Configure NetBackup for the SQL Server non-readable secondary instances that are hidden	34
	Configure the primary server host name for the SQL Server agent	35
	Configure the number of jobs allowed for backup operations	36
	Configure the Maximum jobs per client setting	37
Chapter 4	Configuring RBAC for SQL Server administrators	39
	RBAC roles for the SQL Server administrator	39
	Create a custom role for a non-SQL Server administrator	40

	RBAC permissions that are needed to view and manage SQL Server and VMware jobs	41
Chapter 5	Managing SQL Server assets and their credentials	42
	About the Workloads > Microsoft SQL Server utility	43
	About discovery of SQL Server objects	43
	Discover instances on demand	44
	Discover advanced or basic availability groups on demand	45
	Discover databases on demand	45
	Discover read-scale availability groups	45
	About registering SQL Server instances	46
	Authentication options for SQL Server credentials	47
	Register a SQL Server instance with an existing credential	48
	Register a SQL Server instance with a new credential	49
	Add a credential for SQL Server	50
	View the credential name that is applied to an asset	50
	Edit or delete a named credential	51
	About credential rules	51
	Add a credential rule	52
	Edit or delete a credential rule	53
	Deactivate or activate a credential rule	54
	Preview a credential rule	55
	Credential rule examples	56
	Query builder reference for credential rules	56
	Browse SQL Server assets	58
	View the protection status of SQL Server assets	61
	About intelligent groups	62
	Add an intelligent group	62
	Edit or delete an intelligent group	63
	Preview an intelligent group	64
	Intelligent group examples	65
	Query builder for intelligent groups reference	66
	Use Backup now to back up a SQL Server asset	69
	Remove SQL Server instances	70
	Remove SQL Server databases	71
	Manually add a SQL Server instance	71
	Deactivate or activate an instance	72
	Deactivate or activate a database	72

Chapter 6	Configuring backups with SQL Server Intelligent Policy	74
	About SQL Server Intelligent Policies	75
	Create a SQL Server Intelligent Policy	75
	About policy attributes	76
	Schedule properties for SQL Server Intelligent Policies	77
	Schedule backup types for SQL Server Intelligent Policies	78
	Add instances to a policy	80
	Add databases to a policy	81
	Add intelligent groups to a policy	82
	Add filegroups or files to the backup selections list	83
	Add instance groups to a backup policy	84
	Performance tuning and configuration options	84
	Configure multistriped backups of SQL Server	88
	Converting differential backups to full backups	89
	Converting log backups to full backups	89
	Back up read-only filegroups	90
	Back up read-write filegroups	91
	Perform a manual backup	92
 Chapter 7	 Protecting SQL Server availability groups	 93
	About protecting SQL Server availability groups	93
	Protecting SQL Server availability groups with intelligent policies	94
	Prerequisites for protecting SQL Server availability groups	95
	Configure a backup policy to protect a SQL Server availability group	96
	Protecting SQL Server availability groups with batch file-based policies	99
	About protecting the preferred replica in a SQL Server availability group (batch file-based policies)	99
	About protecting a specific node in a SQL Server availability group (batch file-based policies)	105
	Protect a SQL Server availability group that crosses NetBackup domains with a policy	108
 Chapter 8	 Protecting SQL Server with VMware backups	 109
	About protecting an application database with VMware backups	109
	Limitations of VMware application backups	111
	About configuring NetBackup for VMware backups that protect SQL Server	112

	Configuring a VMware backup policy to protect SQL Server	113
	Configuring a VMware policy to protect SQL Server using Replication	
	Director to manage snapshot replication	115
	Protect SQL Server data with a VMware backup	116
Chapter 9	Configuring backup policies with Snapshot Client	
		118
	About NetBackup Snapshot Client for SQL Server	118
	How SQL Server operations use Snapshot Client	119
	Snapshot methods	121
	Configuration requirements for SQL Server snapshot and Instant	
	Recovery backups	122
	Configure a snapshot policy for SQL Server	123
	Configure a policy for Instant Recovery backups of SQL Server	125
	Using copy-only snapshot backups to affect how differentials are based	
		127
	Create a copy-only backup	128
	Creating an Instant Recovery backup that is not copy-only (batch	
	file-based policies)	129
	About SQL Server agent grouped snapshots	129
	Restoring a database backed up in a group	130
Chapter 10	Protecting SQL Server in a cluster environment	
		131
	Configure backups of clustered SQL Server instances (SQL Server	
	Intelligent Policy)	131
	Configure backups of clustered SQL Server instances (batch file-based	
	policies)	133
Chapter 11	Managing protection plans for SQL Server	134
	Create a protection plan to protect SQL Server assets	134
	Schedules	137
	Add SQL Server assets to a protection plan	138
	Customize protection settings for a NetBackup for SQL Server asset	
		140
	Remove protection from SQL Server assets	141
	Create a protection plan to protect SQL Server data with a VMware	
	backup	142
	Backup options and Advanced options	144
	Exclude disks from backups	145
	Snapshot retry options	146

	Protect a SQL Server availability group that crosses NetBackup domains with a protection plan	147
Chapter 12	Restoring SQL Server with the NetBackup web UI	149
	Requirements for restores of SQL Server	149
	Perform a complete database recovery	150
	Recover a single recovery point	153
	Options for SQL Server restores	156
	Restore a database (non-administrator users)	157
	Select a different backup copy for recovery	158
	Configuring permissions for redirected restores	161
	Restore SQL Server databases from a VMware backup	163
	Restore a SQL Server availability database to a secondary replica	163
	Restore a SQL Server availability database to the primary and the secondary replicas	165
	Restoring an availability database when an availability group crosses NetBackup domains	168
Chapter 13	Using instant access with SQL Server	169
	Prerequisites when you configure an instant access SQL Server database	169
	Hardware and configuration requirements of instant access	170
	Things to consider before you configure an instant access database	171
	Configure Samba users for SQL Server instant access	172
	Configure an instant access database	175
	View the livemount details of an instant access database	177
	Delete an instant access database	178
	Options for NetBackup for SQL Server instant access	179
	NetBackup for SQL Server terms	180
	Frequently asked questions	181
Chapter 14	Configuring batch-file based policies for SQL Server backups	186
	About batch file-based policies for SQL Server backups	187
	Overview of configuring SQL Server backups with batch file-based policies	187
	Configure the NetBackup services for SQL Server backups and restores (batch file-based policies)	188

About SQL Server security with batch file-based policies	189
Requirements to use batch files with NetBackup for SQL Server	190
Keywords and values used in batch files	191
Create a batch file	199
Run batch files	200
Add a batch file-based policy	200
Schedule properties for SQL Server batch file-based policies	201
Schedule backup types for batch file-based policies	202
Configure an application backup schedule	203
Example application backup schedule	203
Configure automatic backup schedules	204
Example automatic backup schedule	205
Add clients to a policy	205
Add batch files to the backup selections list	206
Options for SQL Server backup operations	207
Create a script to backup a remote SQL Server installation	209
About automatic retry of unsuccessful SQL Server backups	210
Configure a batch file-based policy for a user-directed backup of read-only filegroups	211
View SQL Server read-only backup sets (NetBackup MS SQL Client)	211
Configure a batch file-based policy for a user-directed backup of read-write filegroups	212

Chapter 15

Performing backups and restores with the NetBackup MS SQL Client	214
About the NetBackup MS SQL Client	215
Start the NetBackup MS SQL Client for the first time	215
Select the SQL Server host and instance (NetBackup MS SQL Client)	216
About viewing the properties of the objects selected for backup	217
Perform a user-directed backup of SQL Server databases (NetBackup MS SQL Client)	218
Perform a user-directed backup of SQL Server transaction logs (NetBackup MS SQL Client)	218
Perform a user-directed backup of SQL Server database filegroups (NetBackup MS SQL Client)	219
Perform a user-directed backup of SQL Server database files (NetBackup MS SQL Client)	220
Perform a partial database backup (NetBackup MS SQL Client)	221
Options for NetBackup for SQL Server restores	222

Browsing for SQL Server backup images (NetBackup MS SQL Client)	225
Restore a SQL Server database backup (NetBackup MS SQL Client)	227
Stage a full SQL Server database recovery (NetBackup MS SQL Client)	227
Restore SQL Server filegroup backups (NetBackup MS SQL Client)	228
Recover a SQL Server database from read-write filegroup backups (NetBackup MS SQL Client)	229
Restore SQL Server read-only filegroups (NetBackup MS SQL Client)	230
Restore SQL Server database files (NetBackup MS SQL Client)	230
Restore a SQL Server transaction log image without staging a full recovery (NetBackup MS SQL Client)	231
Perform a SQL Server database move (NetBackup MS SQL Client)	232
About performing a SQL Server page-level restore (NetBackup MS SQL Client)	234
Redirect a SQL Server database to a different host (NetBackup MS SQL Client)	236
About selecting a primary server	237
Perform a restore of a remote SQL Server installation (NetBackup MS SQL Client)	238
Restoring multistreamed SQL Server backups	238
About conventional backups using multiple streams	239
About snapshot backup methods using multiple streams	240
Restoring a multistreamed SQL Server backup with fewer devices than it was backed up with	240
About using bplist to retrieve SQL Server backups	241
About NetBackup for SQL Server backup names	241

Chapter 16

Using NetBackup for SQL Server with multiple NICs	244
Configuration and requirements for SQL Server backups with multiple NICs	245
Configure the NetBackup client with the private interface name	246
Configure backups of SQL Server when you have multiple NICs (SQL Server Intelligent Policies)	247
Configure backups for SQL Server when you have multiple NICs (batch file-based policies)	248

Restore SQL Server when you have multiple NICs (NetBackup MS SQL Client)	249
Configure backups of a SQL Server cluster when you have multiple NICs (SQL Server Intelligent Policies)	250
Configure backups of a SQL Server cluster when you have multiple NICs (batch file-based policies)	251
Create a batch file for backups of a SQL Server cluster when you have multiple NICs (batch file-based policies)	252
Restore a SQL Server cluster when you have multiple NICs (NetBackup MS SQL Client)	253

Chapter 17	Performance and troubleshooting	256
	NetBackup for SQL Server performance factors	257
	About monitoring NetBackup for SQL Server operations	260
	About debug logging for SQL Server troubleshooting	261
	Set the debug level on a SQL Server client	263
	VSS provider logs	263
	Troubleshooting credential validation	264
	Troubleshooting VMware backups	265
	SQL Server log truncation failure during VMware backups of SQL Server	267
	Intelligent group backup completes with a status code 1 when database not found in the SQL Server environment	268
	Set the maximum trace level for NetBackup for SQL Server	268
	Reporting of unsuccessful filegroup or file backups	269
	About minimizing timeout failures on large SQL Server database restores	269
	SQL Server restore fails when you restore a SQL Server compressed backup image as a single stripe or with multiple stripes	270
	Incorrect backup images are displayed for availability group clusters	271
	A restore of a SQL Server database fails with Status Code 5, or Error (-1), when the host name of the SQL Server or the SQL Server database name has trailing spaces	271
	A move operation fails with Status Code 5, or Error (-1), when the SQL Server host name, the database name, or the database logical name has trailing spaces	272
	Unable to discover or browse availability group replicas	272
	SQL Server policy backups that use intelligent groups may fail with a status code 200 when the associated credentials in a credential rule are invalid	272
	About disaster recovery of SQL Server	274

	Preparing for disaster recovery of SQL Server	275
	Recovering SQL Server databases after disaster recovery	275
Appendix A	Other configurations	278
	Configuring multiplexed backups of SQL Server	278
	Restoring a multiplexed SQL Server backup	279
	About SQL Server backups and restores in an SAP environment	279
	Creating batch files for automatic backups in for SQL Server in an SAP environment	280
	Monitoring backups on SQL Server	281
	Restoring the R/3 database	281
	About policy configuration for SQL Server in an SAP environment	284
	Configure NetBackup to support database log-shipping	284
	Backing up SQL Server in an environment with log shipping	285
	About NetBackup for SQL Server with database mirroring	286
	Configure NetBackup to support database mirroring	286
	Perform simultaneous backups for mirrored partners	287
	Restore a mirrored database backup image	288
Appendix B	Register authorized locations	290
	Registering authorized locations used by a NetBackup database script-based policy	290

About NetBackup for SQL Server

This chapter includes the following topics:

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

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.
Instances are automatically discovered in the NetBackup environment.
- Back up SQL Server instances, databases, and availability groups.
NetBackup offers the following types of SQL Server backup methods:
 - SQL Server Intelligent Policies (SIP)
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.
 - Protection plans
A SQL Server administrator can select one or more protection plans that contain the wanted storage, backup, and tuning settings.
 - Policies, using clients and batch files
These policies include a list of SQL Server 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 or policies.
- 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.

NetBackup for SQL Server features

Table 1-1 NetBackup for SQL Server features

Feature	Description
NetBackup integration	Full integration with the NetBackup primary server and media manager. Job monitoring from the primary server.
Integration with NetBackup role-based access control (RBAC)	The NetBackup web UI provides RBAC roles to control which NetBackup users can manage 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.
SQL Server Intelligent Policy and protection plans	The following benefits are included: ■ Intelligent groups that automatically discover instance databases. (Intelligent groups are only available with policies.) ■ Use an intelligent policy or single protection plan 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 protection plan or 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 are registered, the SQL Server workload administrator can protect the SQL Server assets with a policy or a protection plan. Credential rules and intelligent groups let you automatically register and protect instances and databases.

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

Feature	Description
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. ■ The user can use credential rules to automatically register instances.
Backup and restore features	The following features are available for backups and restores with the NetBackup web UI: ■ Backups are managed entirely by the NetBackup server from a central location. Administrators can schedule automatic, unattended backups for instances on local or remote hosts across the network. ■ The web UI provides recovery points from which you can easily perform various recovery operations. If instant access is configured, you can also create an instant access database from a recovery point. ■ NetBackup supports the backup of databases, files, filegroups, transaction logs. ■ Backup schedules for full, differential, or transaction log backups. ■ Manual backups and copy-only backups. ■ Backups of only read-write filegroups. This feature is available with policy-based backups. ■ Support for high availability (HA) environments, including SQL Server clusters and availability groups. ■ Browse backups and select the ones to restore. ■ Restore SQL Server objects to different locations (redirected restores). ■ Ability to use multiple stripes during a backup. ■ Configure 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.

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

Feature	Description
Support for the 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.
NetBackup encryption	(Backups with policies) When encryption is enabled, NetBackup encrypts the backup for the instances or the clients that are listed in the policy.
Batch file-based SQL Server policies	Support for the backup policies that use batch files and a list of clients.
Additional restore options with the NetBackup MS SQL Client interface	The NetBackup MS SQL Client interface supports some additional restore operations that are not available in the web UI. These operations include filegroup restores, the restore of read-write and read-only filegroups, database file restores, and page-level restores.

Installation

This chapter includes the following topics:

- [Planning the installation of NetBackup for SQL Server](#)
- [NetBackup server and client requirements](#)
- [Requirements for using NetBackup for SQL Server in a NetBackup cluster](#)
- [License for NetBackup for SQL Server](#)

Planning the installation of NetBackup for SQL Server

To use the new features that are included in NetBackup for SQL Server in NetBackup 11.2, upgrade your NetBackup for SQL Server clients to NetBackup 11.2. 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 20.

Table 2-1 Installation steps for NetBackup for SQL Server (*continued*)

Step	Action	Description
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 18.
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 19.

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](#).
 Every NetBackup server includes the NetBackup client software by default. Therefore, you can use NetBackup for SQL Server on a NetBackup server or client (if NetBackup for SQL Server is supported on that platform).
- Make sure that you configure any backup media that the storage unit uses. The number of media volumes that are required depends on several things:
 - The devices that are used and the storage capacity of the media.
 - The sizes of the databases that you want to back up.
 - The amount of data that you want to archive.
 - The size of your backups.
 - The frequency of backups or archives.
 - The length of retention of the backup images.

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

NetBackup client requirements

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

- The NetBackup client software is installed on the computer that has the databases you want to back up.
- If the database is clustered, you must use the same version of NetBackup on each node in the cluster.
- For SQL Server availability groups, install the client on each replica in the availability group where you want backups to occur.
- In a SQL Server cluster environment, install the NetBackup client on each node in the cluster. Each node must have the same version of NetBackup.
- In a VMware environment, install the NetBackup client software on the virtual machines that have SQL Server running.
- If you have multiple NICs, install the NetBackup client using the private interface name.
- If the SQL Server client is on a different host than the primary server or media server, then install the NetBackup client on that host.

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.
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 Web UI Administrator's Guide](#).

For a NetBackup cluster, a valid license for NetBackup for SQL Server must exist on each node where NetBackup server resides.

Host configuration and job settings

This chapter includes the following topics:

- [Configuring SQL Server hosts and user permissions](#)
- [Installing the VSS provider for vSphere](#)
- [Configure the NetBackup services for SQL Server backups and restores](#)
- [Configure local security privileges for SQL Server](#)
- [Reviewing the auto-discovered mappings](#)
- [Configuring mappings for restores of a distributed applications, clusters, or virtual machines](#)
- [Configure the ODBC connection](#)
- [Configure NetBackup for the SQL Server non-readable secondary instances that are hidden](#)
- [Configure 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](#)

Configuring SQL Server hosts and user permissions

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

Table 3-1 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 Cohesity VSS provider. For VMware backups with T-SQL snapshots, you also need to disable the SQL Server VSS Writer service.	See "Installing the VSS provider for vSphere" on page 23. See "Disable the SQL Server VSS Writer service" on page 23.
Step 2	Assign users to the necessary RBAC roles.	See "RBAC roles for the SQL Server administrator" on page 39.
Step 3	(Conditional) To use SQL Server Intelligent Policies or protection plans, add the necessary SQL Server credentials.	Add the SQL Server credentials that you need for database discovery and the credentials to perform recovery. See "Register a SQL Server instance with an existing credential" on page 48. See "Register a SQL Server instance with a new credential" on page 49.
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 "Configure the NetBackup services for SQL Server backups and restores" on page 24.
Step 5	(Conditional) To use SQL Server Intelligent Policies or protection plans, 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 25.
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. See "Reviewing the auto-discovered mappings" on page 26.

Installing the VSS provider for vSphere

To use the 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 VSS provider

- 1 Browse to the following location:

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

- 2 Double-click on the **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 VSS provider

- 1 Open **Add or Remove Programs**.
- 2 Double-click on **VSS provider**.

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

Disable the SQL Server VSS Writer service

To perform VMware application backups with the option **Enable T-SQL snapshots**, you must disable the SQL Server VSS Writer service.

Note that when you disable this service, any jobs fail that do not use T-SQL snapshots. All policies that protect a VM should use the same snapshot method. Do not mix any policies that use and do not use the T-SQL snapshot method.

To disable the SQL Server VSS Writer service

- 1 On the SQL Server system where the NetBackup client is installed, log on as an Administrator.
- 2 Open the Windows Services application.
- 3 In the right pane, right-click on **SQL Server VSS Writer** service and select **Stop**.
- 4 In the right pane, right-click on **SQL Server VSS Writer** and select **Properties**.
- 5 From the **Startup type** list, click **Disabled**.
- 6 Select **OK**.

Configure 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 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**.
- For VMware backups with Replication Director, the account has access to the CIFS shares on the NetApp disk array.

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 If you use NetBackup MS SQL Client and if the SQL Server host and instance use standard or mixed security, perform the following steps:
 - Open the NetBackup MS SQL Client.
 - Select **File > Set SQL Server connection properties**.
 - Provide the SQL Server **Userid** and **Password**, click **Apply > Close**.
- 3 In the Windows Services application, open the **NetBackup Client Service**.
- 4 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**.

5 Open the **NetBackup Legacy Network Service**.

6 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.

7 If you selected a different logon account, restart the services.

8 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 25.

9 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:
 - **Act as part of the operating system**
 - **Create a token object**
 - **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.

Reviewing the auto-discovered mappings

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

See [the section called “Approve the auto-discovered mappings for a cluster”](#) on page 27.

See [the section called “Auto-discovered mappings for a SQL Server cluster in a multiple NIC environment”](#) on page 29.

See [the section called “Manually map host names”](#) on page 30.

Examples of the configurations that have multiple host names include:

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

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

Auto-discovered mappings for a cluster

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

- If the backup policy includes the cluster name (or virtual name)
- If the NetBackup client is installed on more than one node in the cluster, the virtual name must be mapped to each node.

If the NetBackup Client is only installed on one node, then no mapping is necessary.

Approve the auto-discovered mappings for a cluster

To approve the auto-discovered mappings for a cluster

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

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

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

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

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

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

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

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

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

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

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

In [Table 3-2](#), FCI is a SQL Server failover cluster instance. WSFC is Windows Server Failover Cluster.

Table 3-2 Example mapped host names for SQL Server environments

Environment	Host	Mapped host names
FCI (cluster with two nodes)	Physical name of <i>Node 1</i>	Instance cluster name
	Physical name of <i>Node 2</i>	Instance cluster name
Basic or advanced availability group (primary and secondary)	Primary name	WSFC name

Table 3-2 Example mapped host names for SQL Server environments
(continued)

Environment	Host	Mapped host names
	Secondary name	WSFC name
Basic or advanced availability group, with an FCI (primary FCI and secondary FCI)	Primary FCI name	WSFC name
	Secondary FCI name	WSFC name
	Physical name of <i>Node 1</i>	Virtual name of the SQL Server cluster
	Physical name of <i>Node 2</i>	Virtual name of the SQL Server cluster

Auto-discovered mappings for a SQL Server cluster in a multiple NIC environment

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

To approve the auto-discovered mappings for a SQL Server cluster in a multiple NIC environment

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

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

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

Host	Auto-discovered mapping
<code>client01-bk.lab04.com</code>	<code>clustername-bk.lab04.com</code>
<code>client02-bk.lab04.com</code>	<code>clustername-bk.lab04.com</code>

- 3 Click the name of the host.

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

For example, if following mapping is valid for `client01-bk.lab04.com`, then you approve it.

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

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

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

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

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

Example mapped host names for a SQL Server cluster in a multi-NIC environment

Table 3-3 Example mapped host names for a SQL Server cluster in a multi-NIC environment

Host	Mapped host names
Private name of <i>Node 1</i>	Virtual name of the SQL Server cluster on the private network
Private name of <i>Node 2</i>	Virtual name of the SQL Server cluster on the private network

Manually map host names

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

To manually map host names

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

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

Configuring mappings for restores of a distributed applications, clusters, or virtual machines

This configuration is required for restores of a SQL Server cluster or a SQL Server availability group.

To configure mappings for restores of a distributed application, cluster, or virtual machine

- 1 On the left, click **Hosts > Host properties**.
- 2 Select the primary server.
- 3 If necessary, click **Connect**. Then click **Edit primary server**.
- 4 Select **Distributed application restore mapping**.
- 5 Click **Add**.
- 6 Provide the name of the application host and the name of the component host.
See [Example entries for SQL Server](#).

Example entries for SharePoint hosts for a VMware backup

Table 3-4 Example entries for SharePoint hosts for a VMware backup

Application hosts	Component hosts
VM Display name of node	Physical name of SharePoint central Administration Server

Example entries for SQL Server

Table 3-5 Example entries for SQL Server

Environment	Application host	Component host
FCI (cluster with two nodes)	Instance cluster name	Physical name of <i>Node 1</i>
	Instance cluster name	Physical name of <i>Node 2</i>
Advanced or basic availability group (primary and secondary)	WSFC name	Primary replica name
	WSFC name	Secondary replica name
Advanced or basic availability group with an FCI (primary FCI and secondary FCI)	WSFC name	Primary replica FCI name
	WSFC name	Secondary replica FCI name
	Instance cluster name	Physical name of <i>Node 1</i>
	Instance cluster name	Physical name of <i>Node 2</i>
VMware	VM display name, VM BIOS UUID, or VM DNS name (Primary VM identifier other than VM hostname)	Host name of the VM

Configure the ODBC connection

NetBackup handles the encryption settings of an ODBC connection from a NetBackup client to a target SQL Server instance. These settings are configured in the host properties for the connecting client and can only be configured with the `hostProperties` API endpoint or the `nbsetconfig` command.

Note: The RBAC role **Default Microsoft SQL Server Administrator** does not have permissions to edit the host properties. Alternatively, workload administrators can log on locally to the host and use the `nbsetconfig` command to make the host property changes.

ODBC connections are created using an ODBC connection string. This string is made up of a list of key-value pairs that changes the connection's behavior depending on the key-value pair.

Table 3-6 `hostProperties` API endpoint parameters for ODBC connections

Parameter	Description
<code>encrypt: true false</code>	Whether to encrypt the connection using TLS. For NetBackup clients that are updated to 10.4 or later, the SQL Server ODBC connections from the client to a target SQL Server instance are encrypted by default.
<code>trustServerCertificate: true false</code>	Whether to trust the target SQL Server instance's certificate. For 10.4 and later clients, <code>TrustServerCertificate</code> is set to true by default to prevent unexpected connection failures on upgrade.
<code>preferredODBCDriver: {driver 1, driver 2, ...}</code>	The name of the supported SQL Server ODBC driver to use during the connection. The value can be one or many individual driver names. List driver names in the order of preference. Or, set the value to "NEWEST" or to "OLDEST". ("NEWEST" or "OLDEST" must be specified alone.) The driver <code>SQL Server</code> does not support encryption. Customers that have strict security policies and concerns should use the <code>NEWEST</code> driver or whatever version their company has certified. The available driver values are as follows: "SQL Native Client" "SQL Server Native Client 10.0" "SQL Server Native Client 11.0" "SQL Server" "ODBC Driver 11 for SQL Server" "ODBC Driver 13 for SQL Server" "ODBC Driver 17 for SQL Server" "ODBC Driver 18 for SQL Server" "OLDEST" "NEWEST"

Example using host properties API endpoint

The following example uses the host properties API endpoint to enable encryption, trust the target client certificate, and indicates NetBackup should use the driver "SQL Server". If that driver is not available, then NetBackup should use "ODBC Driver 17 for SQL Server".

```
PATCH https://{primary-server}/netbackup/config/hosts/{client-host-id}/
host-properties?fieldset%5BhostProperties%5D=clientMssql
Body:
{
  "data": {
    "type": "hostProperties",
```

```
"id": "{{client-host-id}}",
"attributes": {
  "clientMssql": {
    "trustServerCertificate": true,
    "preferredODBCDriver": [
      "SQL Server",
      "ODBC Driver 17 for SQL Server"
    ],
    "encrypt": true
  }
}
```

Configure NetBackup for the SQL Server non-readable secondary instances that are hidden

To support the non-readable secondary instances that are hidden, additional configuration is required. No configuration on the secondary provides the port number of the primary to NetBackup, so the NetBackup user must provide it. These settings are configured in the host properties for the client and can only be configured with the `hostProperties` API endpoint or the `bpsetconfig` command.

Note: The RBAC role **Default Microsoft SQL Server Administrator** does not have permissions to edit the host properties. Alternatively, workload administrators can log on locally to the host and use the `bpsetconfig` command to make the host property changes.

This entry is a multi-string entry. Each entry is in the form that is normally used to connect to the instances with SQL Server utilities: `host\instance,port`

host-properties API example

`https://{{host}}/netbackup/config/hosts/{{host-uuid}}/host-properties`

```
"clientMssql": {
  "trustServerCertificate": true,
  "preferredODBCDriver": [
    "OLDEST"
  ],
  "configList": [
```

```
        "host16vm5\\SQL2K22,1633",  
        "host16vm6\\SQL2K22,1634"  
    ],  
    "encrypt": true  
}
```

host-properties patch API example

```
patch: https://{host}/netbackup/config/hosts/{host-uuid}/host-properties  
{  
  "data": {  
    "type": "hostProperties",  
    "id": "{host-uuid}",  
    "attributes": {  
      "clientMssql": {  
        "configList": [  
          "host16vm5\\SQL2K22,1633",  
          "host16vm6\\SQL2K22,1634"  
        ]  
      }  
    }  
  }  
}
```

Configure 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 intelligent policies and protection plans, this name is the host name by which the primary server identifies itself. For other operation types, this name is the host name of the primary server that is configured in the batch file. 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

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**. (For batch file-based policies, the equivalent settings are the **Stripes** setting and the `BATCHSIZE` keyword.)

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

Limit jobs per policy	Sets the maximum number of instances that NetBackup can back up concurrently in each policy. This setting is configured in the policy attributes.
Maximum jobs per client	In a policy, the maximum number of jobs per client that you want to allow. This setting applies to all clients in all policies. You can configure this property in the web UI in the Global attributes host properties for the primary server. See “Configure the Maximum jobs per client setting” on page 37.
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
Maximum concurrent write drives	The number of tape drives that NetBackup can use at one time for jobs to this storage unit. This setting is configured in the storage unit properties. 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 primary server.
- 3 If necessary, click **Connect**. Then click **Edit primary server**.
- 4 Click **Global attributes**.
- 5 Change the **Maximum jobs per client** value to the wanted value.
The default is 1.

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

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

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

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

Refer to the following definitions:

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

Configuring RBAC for SQL Server administrators

This chapter includes the following topics:

- [RBAC roles for the SQL Server administrator](#)
- [Create a custom role for a non-SQL Server administrator](#)
- [RBAC permissions that are needed to view and manage SQL Server and VMware jobs](#)

RBAC roles for the SQL Server administrator

NetBackup enables control over which users can access which SQL Server assets using Role Based Access Control (RBAC). You can grant RBAC access globally (to all SQL Server assets) or to specific databases, instances, or availability groups.

The Default Microsoft SQL Server Administrator role has access to all SQL Server assets (global). With this role the administrator can also manage credentials for instances and availability group replicas. (These credentials are managed on the **Instances** tab in **Workloads > Microsoft SQL Server**.)

In addition, you may need other custom roles to give additional access to your SQL Server administrators.

- A role that is restricted to an individual instance, database, or availability group.
- A role that gives specific privileges to manage the host properties for SQL Server clients.
- A role for a non-SQL Server administrator.

Note the following:

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

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

Create a custom role for a non-SQL Server administrator

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

To create a custom role for a non-SQL Server administrator

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

RBAC permissions that are needed to view and manage SQL Server and VMware jobs

To view or manage jobs for SQL Server operations (including the VMware backups that protect SQL Server), the proper asset permissions are needed, as follows:

- To restart a failed availability database backup job that was started from an intelligent policy or a batch file-based policy, the RBAC user needs to have the necessary RBAC permissions for the availability group asset. The restart option is only available from the parent backup job that has the availability group asset ID. RBAC permissions for the SQL Server instance or the database asset are not sufficient to view the parent backup job.
- To view the job details for a VMware backup or restore of SQL Server the user must have the necessary RBAC permissions for the instance and the database. To view the corresponding backup job for the database, the user must have the proper RBAC permissions for the VMware asset.
- To view the snapshot jobs that are associated with the database assets, the user must have the necessary RBAC permissions for the database assets.

Managing SQL Server assets and their credentials

This chapter includes the following topics:

- [About the Workloads > Microsoft SQL Server utility](#)
- [About discovery of SQL Server objects](#)
- [About registering SQL Server instances](#)
- [About credential rules](#)
- [Browse SQL Server assets](#)
- [View the protection status of SQL Server assets](#)
- [About intelligent groups](#)
- [Use Backup now to back up a SQL Server asset](#)
- [Remove SQL Server instances](#)
- [Remove SQL Server databases](#)
- [Manually add a SQL Server instance](#)
- [Deactivate or activate an instance](#)
- [Deactivate or activate a database](#)

About the Workloads > Microsoft SQL Server utility

NetBackup displays the instances, databases, and availability groups that it discovers in the **Workloads > Microsoft SQL Server** node of the NetBackup web UI. It also displays any instances you added manually and any instant access databases that you created. The properties for an instance, replica, or database indicate the name of any Intelligent Policies that protect those objects.

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, 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.

The **Microsoft SQL Server** node contains the following tabs:

- **Instances**
Contains all the SQL Server instances that NetBackup discovers or that you manually added. Instances that belong to an availability group are also included in this list.
- **Availability groups**
Contains all the SQL Server availability groups that NetBackup discovers.
- **Databases**
Contains all the SQL Server databases that NetBackup discovers.
- **Instant access databases**
Contains the instant access databases that you created.
See the chapter *Using instant access with SQL Server*.
- **Credential rules**
Contains the credential rules that you have created.
- **Intelligent groups**
Contains the intelligent groups that you have created.

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

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 instances on demand

You can manually start the NetBackup discovery process if you want to immediately discover new SQL Server instances or availability group instances in your environment.

To discover new SQL Server instances

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

Discover 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 must have credentials before you can perform on-demand discovery.

To discover advanced or basic availability groups

- 1 Click on the **Availability groups** tab.
- 2 Click **Discover availability groups**.
- 3 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.

- 4 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.

Note that after you recover a database that was deleted from the SQL Server environment NetBackup does not immediately reflect the updated database state. To refresh the database state you must run discovery manually or wait until the next autodiscovery.

To discover databases

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

Note that only registered instances are shown in this list.

- 4 Select **Discover**.

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 Click on the **Instances** tab.
- 2 Select one of the replicas that is part of the availability group and click **Manage credentials**.

- 3 Follow the prompts to provide the credentials for the replica.
- 4 Click on the **Availability groups** tab.
- 5 Click **Discover availability groups**.
- 6 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.
- 7 Click **Discover**.

About registering SQL Server instances

To allow for full discovery of SQL Server assets and to protect those assets, you must register instances with a credential. Credentials are not supported at the database or the availability group level.

Instances can be registered in one of the following ways:

- Manually, by registering individual instances.
See [“Register a SQL Server instance with an existing credential”](#) on page 48.
See [“Register a SQL Server instance with a new credential”](#) on page 49.
- Automatically, with a credential rule. A rule registers any unregistered or any newly discovered instances that meet the rule criteria.
See [“About credential rules”](#) on page 51.

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.

See [“Configure the NetBackup client with the private interface name”](#) on page 246.

Registering Microsoft SQL Server failover cluster instances (FCIs)

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

Validation of credentials

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

For a SQL Server cluster or if an availability group instance is part of a 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.

Authentication options for SQL Server credentials

The following authentication options are available to register SQL Server instances. The NetBackup web UI supports Windows authentication and Windows Active Directory authentication. It does not support Mixed Mode or SQL Server authentication.

Table 5-1 Authentication options for SQL Server 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. See “Configure the NetBackup services for SQL Server backups and restores” on page 24. See “Configure local security privileges for SQL Server” on page 25.
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. ■ 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. See “Configure the NetBackup services for SQL Server backups and restores” on page 24.

Register a SQL Server instance with an existing credential

For unregistered or newly discovered instances, you can also use a credential rule to automatically register these assets.

See [“About credential rules”](#) on page 51.

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.

To select an existing credential for a SQL Server instance

- 1 Review the recommendations and requirements for the type of credentials that you want to use for authentication.
See [“Authentication options for SQL Server credentials”](#) on page 47.
- 2 Select the **Instances** tab.
- 3 Select the check box for each instance that you want to register. Then select **Manage credentials**.
For availability groups, each replica must be registered with credentials.
- 4 Select **Select from existing credentials**.
- 5 Select **Next**.
- 6 Select the credential that you want to use for the selected assets and select **Next**.

Register a SQL Server instance with a new credential

For unregistered and newly discovered instances, you can also use a credential rule to automatically register these assets.

See [“About credential rules”](#) on page 51.

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.

To add a new credential for a SQL Server instance

- 1 Review the recommendations and requirements for the type of credentials that you want to use for authentication.
See [“Authentication options for SQL Server credentials”](#) on page 47.
- 2 Select the **Instances** tab.
- 3 Select the check box for each instance that you want to register. Then select **Manage credentials**.
For availability groups, each replica in the group must be registered with credentials.

- 4 Select **Add credentials** and select **Next**.
- 5 Select the authentication option.
See [“Authentication options for SQL Server credentials”](#) on page 47.

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, select **Credential management**.
- 2 On the **Named credentials** tab, select the **Add** button.
- 3 From the **Credential store** options, select **NetBackup**.
- 4 Select the **Start** button.
- 5 Provide the following properties:
 - Credential name (for example: *server_credential1*)
 - Tag (for example: *workload name*)
 - Description (for example: This credential is used to access *workload name*)
- 6 Select the **Next** button.
- 7 From the **Category** list, select **Microsoft SQL Server**.
- 8 Provide the authentication details that are needed to connect to the SQL Server.
See [“Register a SQL Server instance with an existing credential”](#) on page 48.
See [“Register a SQL Server instance with a new credential”](#) on page 49.
- 9 Select the **Next** button.
- 10 Add one or more RBAC roles that you want to have access to the credential.
 - Select the **Add** button.
 - Select the role name.
 - Select the credential permissions that you want the role to have.
- 11 Select the **Next** button and follow the prompts to complete the wizard.

View the credential name that is applied to an asset

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

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.

If a credential is registered for an instance the credential name is displayed.

Edit or delete a named credential

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

Edit a named credential

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

To edit a named credential

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

Delete a named credential

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

To delete a named credential

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

About credential rules

With a credential rule, NetBackup registers any instances that are unregistered or newly discovered with the credentials that are configured in the credential rule. You can configure multiple credential rules for automatic registration. If the criteria in

two or more rules match an instance, NetBackup automatically applies the rule that was added first.

Credential rules provide the following benefits for the management of SQL Server assets:

- A rule can automatically register any newly discovered or any unregistered instances.
- All the instances in the credential rule use the same credentials. Any changes to the associated credential are automatically applied to the instances in the rule.
- On the **Credential rules** tab, you can easily see the credential that each rule uses.

Add a credential rule

You can add a credential rule so that unregistered or newly discovered instances are automatically registered with a credential.

To add a credential rule

- 1 Select the **Credential rules** tab.
- 2 Select **Add** or **Add credential rule**.
- 3 Provide the **Name** and **Description** for the rule.
- 4 Select **Add condition**.
- 5 To select the criteria for the rule, make the following selections from the menus.
Start by making a selection for the **Field** list.
Then make a selection for **Operator**.
For the **Value** field, select or enter the value.
For details on the query criteria, see the following topics.
See [“Credential rule examples”](#) on page 56.
See [“Query builder reference for credential rules”](#) on page 56.
- 6 Add any additional conditions or sub-queries for the rule.
- 7 (Optional) To see the instances that match the rule criteria, select the **Preview** button. See the following topic for details.
See [“Preview a credential rule”](#) on page 55.
- 8 Select **Next**.
- 9 Select from the following options:

- **Select from existing credentials**
Select the credential that you want to apply to the credential rule.
Only the credentials with the type **Microsoft SQL Server** display.
- **Add credentials**
Enter the **Credential name** and select the authentication option. Then select **Next**.
See [“Authentication options for SQL Server credentials”](#) on page 47.
Add any RBAC roles or permissions that you want to apply to the new credential.

10 Select **Add**.

Any unregistered or any newly discovered instances that meet the rule criteria are registered with the associated credential in the rule. If you have the necessary RBAC permissions, you can view the instances that were registered in the audit messages. (Go to **Security > Security events**. Then select the **Audit events** tab.)

Edit or delete a credential rule

You can edit a credential rule if you want to change the criteria by which NetBackup applies a credential rule to instances. You can delete a credential rule if it is no longer needed.

Edit a credential rule

To edit a credential rule

- 1 Select the **Credential rules** tab.
- 2 Locate and select the credential rule.
- 3 Select **Edit**.
- 4 Make any wanted changes to the rule.
- 5 (Optional) To see the instances that match the rule criteria, select the **Preview** button. See the following topic for details.
See [“Preview a credential rule”](#) on page 55.
- 6 Select **Next**.

- 7 Make any wanted changes to the credentials that you want to use for the rule.
- 8 Follow the prompts in the wizard to finish editing the rule.

After you save the changes, any unregistered or any newly discovered instances that meet the rule criteria are registered with the associated credential in the rule. If you have the necessary RBAC permissions, you can view the instances that were registered in the audit messages. (Go to **Security > Security events**. Then select the **Audit events** tab.)

Delete a credential rule

Note: If you delete a rule, any newly discovered instances that match the rule criteria are not automatically registered.

To delete a credential rule

- 1 Select the **Credential rules** tab.
- 2 Locate and select the credential rule.
- 3 Select **Delete > Delete**.

Deactivate or activate a credential rule

After you create a credential rule, you can deactivate the rule temporarily or make a rule active again. When a rule is deactivated, NetBackup does not automatically register any instances that meet the rule criteria.

Deactivate a credential rule

To deactivate a credential rule

- 1 Select the **Credential rules** tab.
- 2 Locate and select the credential rule.
- 3 Select **Deactivate**.

Activate a credential rule

To activate a credential rule

- 1 Select the **Credential rules** tab.
- 2 Locate and select the credential rule.
- 3 Select **Activate > Activate**.

Any unregistered instances that match the rule criteria are automatically registered. Any instances that are discovered after you activate the rule are automatically registered. The credential rule is not applied to any instances that are already registered.

If you have the necessary RBAC permissions, you can view the instances that were registered in the audit messages. (Go to **Security > Security events**. Then select the **Audit events** tab.)

Preview a credential rule

You can preview the instances that a credential rule discovers in the current SQL Server environment.

To preview a credential rule

- 1 Select the **Credential rules** tab.
- 2 Locate and select the credential rule and select **Preview**.

The following information is displayed:

Query	The query criteria for the rule.
All	Displays the number of unregistered instances in the NetBackup environment.
Included	Displays the number of unregistered instances in the NetBackup environment that match the query criteria.
Excluded	Displays the number of unregistered instances in the NetBackup environment that do not match the query criteria.
List of instances	The names of the unregistered instances in the NetBackup environment, based on the tab that you selected.

Credential rule examples

This topic provides examples of the kinds of credential rules that you can create to automatically register SQL Server instances.

In the following rule, the credential is applied to any instances where the host name ends with the string `domain.com`. The query string with OData keywords and operators is: `endswith(tolower(clientName), tolower('domain.com'))`

Field	Operator	Value
Host name	Ends with	domain.com

In the following rule, the credential is applied to any instances where the host has NetBackup version 10.5 installed. The query string with OData keywords and operators is: `tolower(nbuVersion)eq tolower('10.5')`

Field	Operator	Value
NetBackup version	Equal to	10.5

Query builder reference for credential rules

You can use the Query builder to create rules to automatically apply credentials to SQL Server instances that match a credential rule.

OData keywords are indicated for when you build queries for credential rules with the NetBackup APIs. The **Preview** feature displays the query string that contains OData keywords and operators. For example, the query to search for the host names that end with "domain.com" is: `endswith(tolower(clientName), tolower('domain.com'))`.

[Table 5-2](#) describes the fields and options for creating credential rules with the Query builder.

Table 5-2 Query builder options for credential rules

Query Builder fields	Description
AND OR	When you add two or more conditions, you can select a connector to join the rules.

Table 5-2 Query builder options for credential rules (*continued*)

Query Builder fields	Description
Field	Select a parameter on which to build the rule. See Table 5-3 on page 57.
Operator	Select an operator. The available operators depend on the parameter that is selected for the Field . See Table 5-4 on page 58.
Value	Specifies a Value for the Field parameter. The Value field can be a list of possible values or a manual entry, depending on the selections that are made in the other fields.

Field (keywords)

[Table 5-3](#) describes the keywords that are available in the **Field** dropdown. The values for each keyword (in the **Value** field) are case-insensitive.

Note: Use OData **Field** keywords are indicated for when you build queries for credential rules with the NetBackup APIs.

Table 5-3 Keywords in the Field dropdown

Field keyword	OData field keyword	Description
Cluster type	clusterType	The type of cluster that the SQL Server is configured in.
Edition	sqlEdition	The edition that is associated with the SQL Server version.
Host name	clientName	The name of the SQL Server host.
Instance name	displayName	The name of the SQL Server instance.
Instance SP	sqlServicePack	The service pack of the SQL Server instance.
Instance state	instanceState	The state of the SQL Server instance.
Microsoft SQL Server release	sqlRelease	The release version of SQL Server that is installed on the host.
Microsoft SQL Server version	sqlVersion	The version of SQL Server that is installed on the host.

Table 5-3 Keywords in the Field dropdown (*continued*)

Field keyword	OData field keyword	Description
NetBackup version	nbuVersion	The version of NetBackup that is installed on the host.

Operators

[Table 5-4](#) describes the operators available in the **Operator** list. The values for each keyword (in the **Value** field) are case-insensitive.

Table 5-4 Operators in the **Operator** list

Operator	OData operator	Description
Starts with	startswith	Matches the value in the Value field when it occurs at the start of a string. For example: If the Value entry is "box", Starts with matches the string "box_car" but not "flatbox".
Ends with	endswith	Matches the value in the Value field when it occurs at the end of a string. For example: If the Value entry is "dev", Ends with matches the string "01dev" but not "01dev99", "devOP", or "Development_machine".
Contains	contains	Matches the value in the Values field wherever that value occurs in the string. For example: If the Value entry is "dev", Contains matches strings such as "01dev", "01dev99", "devOP", and "Development_machine".
Equal to	eq	Matches only the value that is specified in the Value field. For example: If the host name to search for is "SQLtest27", Equal to matches the virtual machine names such as SQLTest27 or sqltest27 or sqlTEST27, and so forth. The name SQLtest28 is not matched.
Not equal	ne	Matches any value that is not equal to the value in the Value field.

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, 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 Select the **Instances** tab.
- 2 To view the available actions for one or more instances, select the check box for the instances. Note that **Backup now** is only available when you select one instance.
- 3 To view the details for an instance, select the link for the instance. You can perform the following tasks.
 - To perform an immediate backup of the instance, select **Backup now**.
 - Select **Add protection** to add the instance to a protection plan.
 - Select **Remove protection** to remove an instance from a protection plan.
 - Select the **Databases** tab to see the databases that are discovered for the instance and their protection information and status.
 - To view the roles that have access to the instance, select 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 To view the available actions for one or more availability groups, select the check box for the availability groups. Note that the **Backup now** option is only available when you select one availability group.
- 2 Select the link for an availability group to view its details. You can perform the following tasks.
 - Select **Backup now** to perform an immediate backup of the availability group.

- Select **Add protection** to add the availability group to a protection plan.
- Select **Remove protection** to remove an availability group from a protection plan.
- Select the **Databases** tab to see the databases that are discovered for the availability group and their protection information and status.
- Select the **Replicas** tab to see the replicas for the availability group, their protection information and status, and any credential that is registered.
- Select the **Permissions** tab to view the roles that have access to the availability group.

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 Select the **Databases** tab.
- 2 To view the available actions for one or more databases, select the check box for each database. Note that the **Backup now** option is only available when you select one database.
- 3 Select the link for the database to view its details. You can perform the following tasks.
 - Review the **Database state**, which indicates the state of the database in the NetBackup environment:
 - Active - The database is active in the NetBackup environment and available for backups.
 - Inactive - This state can indicate that the database is marked as inactive in the NetBackup environment. NetBackup then ignores the database during backups. This state also indicates a database that was deleted from the SQL Server environment. You can recover a database that was deleted from the SQL Server environment. If NetBackup successfully locates the database during discovery or a backup, then it changes the state from **Deleted** to **Active**.
 - Deleted - The database exists in the NetBackup asset database but not in the live SQL Server environment.

- Select **Backup now** to perform an immediate backup of the database.
- Select **Add protection** to add the database to a protection plan.
- Select **Remove protection** to remove a database from a protection plan.
- To see the available recovery points for the database, select **Recovery points**.
- To view the restore jobs for the database, select **Restore activity**.
- To view the roles that have access to the database, select the **Permissions** tab.

View the protection status of SQL Server assets

You can view the policies or protection plans that are used to protect instances, availability groups, or intelligent groups.

To view the protection status of SQL Server assets

- 1 Select one of the following tabs: **Instances**, **Availability groups**, **Databases**, or **Intelligent groups**.
- 2 For instances, availability groups, or databases, the **Protected by** column indicates how the asset is protected. See [Table 5-5](#). For intelligent groups, the **Protected by policy** column indicates if an intelligent group is included in a policy.

Table 5-5 Protection status of SQL Server databases, instances, or availability groups

Protection type or status	"Protected by" value	
	Database	Instance or availability group
A classic policy protects the asset.	Classic policy	Not protected Go to Protection > Policies to see how classic policies are used to protect instances or availability groups.
A protection plan protects the asset.	Protected	Protected
Neither a plan nor a policy protects the asset.	Not protected	Not protected
A policy or a 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

About intelligent groups

Intelligent groups let you create dynamic groups of the SQL Server assets that meet one or more conditions. You can configure multiple intelligent groups. Note that if the criteria in two or more groups match an asset, NetBackup backs up the asset for each policy that contains each of those intelligent groups.

Intelligent groups provide the following benefits for the protection of SQL Server assets:

- Better management of SQL Server assets with the ability to group them together for automatic protection. For example, you can group stand-alone instances, failover cluster instances, or availability groups. You can also group databases by their type (user or system).
- NetBackup dynamically creates the list of assets to back up when the backup policy runs. You do not need to manually add or remove assets from the list.
- Changes to the NetBackup environment are reflected in the intelligent groups. For example, if a SQL Server asset is added or removed from the NetBackup environment, that change is automatically reflected in any intelligent groups that are associated with that asset.

Add an intelligent group

You can create an intelligent server group based on a set of filters called queries. NetBackup automatically selects instances or databases based on the queries and adds them to the group. You can then apply protection to the group. Note that an intelligent group automatically reflects changes in the SQL Server environment and eliminates the need to manually revise the list of assets in the group.

Note: The instances or databases in the intelligent group must be in an active state in the NetBackup environment. If you change the state of an instance or database to inactive, then NetBackup excludes it from the backup.

To add an intelligent server group

- 1 Select the **Intelligent groups** tab and then select **Add**.
- 2 Enter a name and description for the group.

- 3 To select the criteria for the rule, make the following selections from the menus.
Start by making a selection for the **Field** list.
Then make a selection for **Operator**.
For the **Value** field, select or enter the value.
For details on the query criteria, see the following topics.
See [“Intelligent group examples”](#) on page 65.
See [“Query builder for intelligent groups reference”](#) on page 66.
- 4 Add any additional conditions or sub-queries for the rule.
- 5 (Optional) To see the databases that match the rule criteria, select the **Preview** button. See the following topic for details.
See [“Preview an intelligent group”](#) on page 64.
Note: The query-based selection process is dynamic. Changes in the environment can affect which servers the query selects when the policy runs. As a result, the databases that the query selects later when the policy runs may not be identical to those currently listed in the preview.
- 6 Select **Add**.

Edit or delete an intelligent group

You can edit the query for an intelligent group. You can also delete an intelligent group if it is no longer needed.

To edit an intelligent group

- 1 On the **Intelligent groups** tab, select the link for the group that you want to edit.
The table displays the databases that match the query criteria for the intelligent group.
- 2 On the right, select **Edit**.
- 3 Make any wanted changes to the query.

- 4 (Optional) To see the databases that match the rule criteria, select the **Preview** button. See the following topic for details.

See [“Preview an intelligent group”](#) on page 64.

Note: The query-based selection process is dynamic. Changes in the environment can affect which servers the query selects when the policy runs. As a result, the databases that the query selects later when the policy runs may not be identical to those currently listed in the preview.

- 5 Select **Save**.

To delete an intelligent group

- 1 Select the **Intelligent groups** tab.
- 2 Locate the intelligent group that you want to delete. Review the **Protected by policy** column.
 - If a policy does not protect the group, select the check box for the group and then select **Delete**.
 - If a policy does protect the group, edit the policy to remove the intelligent group. Then you can delete the group.

Preview an intelligent group

You can preview the SQL Server assets in the current SQL Server environment that match the query criteria of an intelligent group.

To preview an intelligent group

- 1 Select the **Intelligent groups** tab.
- 2 Locate and select the intelligent group and select **Preview**.

The following information is displayed:

Query	The query criteria for the group.
All	Displays the number of SQL Server databases that are discovered in the NetBackup environment.
Included	Displays the number of discovered databases in the NetBackup environment that match the query criteria.
Excluded	Displays the number of discovered databases in the NetBackup environment that don't match the query criteria. Any databases that are inactive in the NetBackup environment are also excluded.
List of discovered databases	The names of the discovered databases, based on the tab that you selected.

Intelligent group examples

This topic provides examples of the kinds of intelligent groups that you can create to back up SQL Server assets.

Example 1

In the following example, the intelligent group is applied to any instances where one of the following conditions applies.

Cluster type	Equal to	Windows Server Failover Cluster (WSFC)
Host name	Equal to	instance-cluster-name

The query string with OData keywords and operators is: `(tolower(clusterType) eq tolower('WSFC') or tolower(clientName) eq tolower('instance-cluster-name'))`

Example 2

In the following example, the intelligent group is applied to any availability groups one of the following criteria applies.

Cluster type	Equal to	Windows Server Failover Cluster (WSFC)
Host name	Equal to	cluster-name
Availability group ID	Equal to	01234567-89ab-0123-4567-89abcdef0123
Availability group name	Equal to	ag-name

The query string with OData keywords and operators is: `(tolower(clusterType) eq tolower('WSFC') or tolower(clientName) eq tolower('cluster-name') or availabilityGroup/groupId eq 01234567-89ab-0123-4567-89abcdef0123 or tolower(agName) eq tolower('ag-name'))`

Example 3

In the following example, the intelligent group locates any databases that are not part of an availability group and adds them to the group.

Availability group name	Equal to	<empty>
-------------------------	----------	---------

The query string with OData keywords and operators is: `Availability Group name Equal to <empty> (tolower(agName) eq tolower(""))`.

Example 4

In the following example, the intelligent group finds any non-AG databases that are not part of an availability group and adds them to the group.

Availability group name	Not equal to	<empty>
-------------------------	--------------	---------

The query string with OData keywords and operators is: `Availability Group name Not equal <empty> (tolower(agName) eq tolower(""))`.

Query builder for intelligent groups reference

You can use the Query builder to create rules to automatically add SQL Server instances to an intelligent group.

OData keywords are indicated for when you build queries for intelligent groups with the NetBackup APIs. The **Preview** feature displays the query string that contains OData keywords and operators. For example, the query to search for the host names that end with "domain.com" is: `endswith(tolower(clientName), tolower('domain.com'))`.

[Table 5-6](#) describes the fields and options for creating intelligent groups with the Query builder.

Table 5-6 Query builder options for intelligent groups

Query Builder fields	Description
AND OR	When you add two or more conditions, you can select a connector to join the rules.
Field	Select a parameter on which to build the rule. See Table 5-7 on page 67.
Operator	Select an operator. The available operators depend on the parameter that is selected for the Field . See Table 5-8 on page 69.
Value	Specifies a Value for the Field parameter. The Value field can be a list of possible values or a manual entry, depending on the selections that are made in the other fields.

Field (keywords)

[Table 5-7](#) describes the keywords that are available in the **Field** dropdown. The values for each keyword (in the **Value** field) are case-insensitive.

Note: Use OData **Field** keywords are indicated for when you build queries for intelligent groups with the NetBackup APIs.

Table 5-7 Keywords in the Field dropdown

Field keyword	OData field keyword	Description
Availability group ID	availabilityGroup/groupId	The ID of the availability group. Use the format: 01234567-89ab-0123-4567-89abcdef0123
Availability group name	agName	The name of the availability group.
Cluster name	clusterName	The name of the cluster that the SQL Server is configured in.
Cluster type	clusterType	The type of the cluster that the SQL Server is configured in.
Credential name	credentialName	The name of the NetBackup credential.

Table 5-7 Keywords in the Field dropdown (*continued*)

Field keyword	OData field keyword	Description
Database name	<code>displayName</code>	The name of the SQL Server database.
Database size (MB)	<code>dbSize</code>	The size of the SQL Server database. Enter a whole numeric value, in MB.
Edition	<code>sqlEdition</code>	The edition that is associated with the SQL Server version.
Host name	<code>clientName</code>	The name of the host. This name depends on the SQL Server environment, as follows: ■ For a standalone environment, the host name is the client name or the SQL Server name. ■ For availability groups, the host name is the replica name. ■ For a failover cluster instance, the host name is the instance cluster name.
Instance SP	<code>sqlServicePack</code>	The service pack of the SQL Server instance.
Instance cluster name	<code>sqlClusterName</code>	The name of the SQL Server cluster.
Instance name	<code>instanceName</code>	The name of the SQL Server instance.
Instance state	<code>instanceState</code>	The state of the SQL Server instance.
Microsoft SQL Server database state	<code>dbState</code>	The state of the SQL Server database as indicated by SQL Server.
Microsoft SQL Server host name	<code>clientName</code>	The name of the SQL Server host.
Microsoft SQL Server release	<code>sqlRelease</code>	The release version of SQL Server that is installed on the host.
Microsoft SQL Server version	<code>sqlVersion</code>	The version of SQL Server that is installed on the host.
NetBackup version	<code>nbuVersion</code>	The version of NetBackup that is installed on the host.

Operators

[Table 5-8](#) describes the operators available in the **Operator** list. The values for each keyword (in the **Value** field) are case-insensitive.

Table 5-8 Operators in the **Operator** list

Operator	OData operator	Description
Starts with	startswith	Matches the value in the Value field when it occurs at the start of a string. For example: If the Value entry is "box", Starts with matches the string "box_car" but not "flatbox".
Ends with	endswith	Matches the value in the Value field when it occurs at the end of a string. For example: If the Value entry is "dev", Ends with matches the string "01dev" but not "01dev99", "devOP", or "Development_machine".
Contains	contains	Matches the value in the Values field wherever that value occurs in the string. For example: If the Value entry is "dev", Contains matches strings such as "01dev", "01dev99", "devOP", and "Development_machine".
Equal to	eq	Matches only the value that is specified in the Value field. For example: If the host name to search for is "SQLtest27", Equal to matches the virtual machine names such as SQLTest27 or sqltest27 or sqlTEST27, and so forth. The name SQLtest28 is not matched.
Not equal	ne	Matches any value that is not equal to the value in the Value field.
Less than	lt	Applies only to the field Database size . Matches any value that is less than the specified Value , according to the UTF-8 collating sequence.
Less than or equal to	le	Applies only to the field Database size . Matches any value that is less than or equal to the specified Value , according to the UTF-8 collating sequence.
Greater than	gt	Applies only to the field Database size . Matches any value that is greater than the specified Value , according to the UTF-8 collating sequence.
Greater than or equal to	ge	Applies only to the field Database size . Matches any value that is greater than or equal to the specified Value , according to the UTF-8 collating sequence.

Use Backup now to back up a SQL Server asset

With Backup now, SQL Server administrators can back up an asset immediately. For example, you can use Backup now to prepare for the upcoming events that are

outside scheduled backups, such as system maintenance. This type of backup is independent of scheduled backups and does not affect future backups. You can manage and monitor a Backup now job in the same way you manage and monitor other NetBackup jobs. Note that Backup now jobs cannot be restarted.

To use Backup now to back up a SQL Server asset

- 1** Select from the following options:
 - To back up an instance, click the **Instances** tab.
 - To back up a database, click the **Databases** tab.
 - To back up an availability group, click the **Availability groups** tab.
- 2** Select the single instance, database, or availability group that you want to back up.
- 3** Click **Backup now**.
- 4** Choose a protection plan for the backup.

All protection plans to which the asset is subscribed are listed.

To back up an asset that is not subscribed to any protection plan, select **Backup now** and choose from the existing protection plans. You can also create a new protection plan and then use it with a **Backup now** operation.
- 5** Select the type of backup you want to perform from the list. The list only contains the backup types that are available in the protection plan.
- 6** Click **Start backup**.

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

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

Remove SQL Server databases

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

To remove a SQL Server database

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

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

Manually add 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 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 instance cluster name. NetBackup enumerates the FCI under the physical node names and the instance 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 [“Register a SQL Server instance with an existing credential”](#) on page 48.
See [“Register a SQL Server instance with a new credential”](#) on page 49.
You may omit credentials at this time. The instance is marked as unregistered and the **Registered** column is empty.
- 7 Click **Finish**.

Deactivate or activate an instance

You can make an instance inactive in NetBackup so it is excluded from a backup. For example, if the instance is under maintenance.

To deactivate an instance

- 1 Select the **Instances** tab.
- 2 Select the instance that you want to deactivate.
- 3 Select **Deactivate**.

To activate an instance

- 1 Select the **Instances** tab.
- 2 Select the instance that you want to activate.
- 3 Select **Activate**.

Deactivate or activate a database

You can make a database that is part of an intelligent group inactive in NetBackup so it is excluded from a backup. For example, if the database is under maintenance. Note that when you use legacy (non-intelligent group) policies or batch files, NetBackup ignores the **Database state** and backs up both **Active** and **Inactive** databases.

To deactivate a database

- 1 Select the **Databases** tab.
- 2 Select the database that you want to deactivate.
- 3 Select **Deactivate**.

To activate a database

- 1** Select the **Databases** tab.
- 2** Select the database that you want to activate.
- 3** Select **Activate**.

Configuring backups with SQL Server Intelligent Policy

This chapter includes the following topics:

- [About SQL Server Intelligent Policies](#)
- [Create a SQL Server Intelligent Policy](#)
- [About policy attributes](#)
- [Schedule properties for SQL Server Intelligent Policies](#)
- [Schedule backup types for SQL Server Intelligent Policies](#)
- [Add instances to a policy](#)
- [Add databases to a policy](#)
- [Add intelligent groups to a policy](#)
- [Add filegroups or files to the backup selections list](#)
- [Add instance groups to a backup policy](#)
- [Performance tuning and configuration options](#)
- [Back up read-only filegroups](#)
- [Back up read-write filegroups](#)
- [Perform a manual backup](#)

About SQL Server Intelligent Policies

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

The SQL Server Intelligent Policy includes the following criteria:

- Storage unit and media to use
- Policy attributes
- Backup schedules: Full, differential-incremental, transaction log
- The SQL Server objects to back up.

Different policies are required to back up instances, availability groups, instance groups, or intelligent groups. You cannot mix instances, availability groups, instance groups, or intelligent groups. For example, if you create a policy with instances or databases and later select the **Protect intelligent groups** option, the instances or databases are deleted from the policy.
- Backup selections: Whole database, filegroups, or files

Create a SQL Server Intelligent Policy

This topic describes how to create an intelligent policy to protect SQL Server instances, databases, intelligent groups, or instance groups. Other topics cover how to protect availability groups and clusters.

See [“About protecting SQL Server availability groups”](#) on page 93.

See [“Configure backups of clustered SQL Server instances \(SQL Server Intelligent Policy\)”](#) on page 131.

To create a SQL Server Intelligent Policy

- 1 Before you configure an intelligent policy ensure that you have registered the SQL Server instances that you want to protect.

See [“Register a SQL Server instance with an existing credential”](#) on page 48.

See [“Register a SQL Server instance with a new credential”](#) on page 49.

See [“Add a credential rule”](#) on page 52.
- 2 Sign in to the primary server as a user that has the RBAC Administrator role or a role that can manage policies.
- 3 On the left, select **Protection > Policies**.

- 4 Select the **Add** button.
- 5 Type a unique name for the new policy.
- 6 In the **Policy type** list, select **MS-SQL-Server**.
- 7 Complete the entries on the **Attributes** tab.
See [“About policy attributes”](#) on page 76.
- 8 Add other policy information as follows:
 - On the **Instances and databases** tab, choose how you want to protect SQL Server.
If you choose the **Protect instances and databases** option, you can select either individual instances or databases.
See [“Add instances to a policy”](#) on page 80.
See [“Add databases to a policy”](#) on page 81.
See [“Add intelligent groups to a policy”](#) on page 82.
See [“Add instance groups to a backup policy”](#) on page 84.
 - Add schedules.
See [“Schedule properties for SQL Server Intelligent Policies”](#) on page 77.
 - (Optional) Select the specific filegroups or files that you want to back up.
By default, NetBackup backs up an entire database. You cannot use intelligent groups to perform filegroup or file backups.
See [“Add filegroups or files to the backup selections list”](#) on page 83.
 - (Optional) Make changes to any tuning parameters.
See [“Performance tuning and configuration options ”](#) on page 84.
- 9 When you have completed the policy configuration, select **Create**.
In the **Workloads > Microsoft SQL Server** utility, the properties for an instance or database indicate the name of any intelligent policies that protect those objects.

About policy attributes

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

For more information on policy attributes, see the [NetBackup Administrator's Guide, Volume I](#).

Table 6-1 Policy attributes for NetBackup for SQL Server policies

Attribute	Description
Policy type	Determines the types of clients that can be backed up with the policy. For SQL Server databases, select the policy type MS-SQL-Server .
Limit jobs per policy	Sets the maximum number of instances that NetBackup can back up concurrently with this policy.
Compress	Enables the compression of backups by NetBackup. If you enable NetBackup compression, do not enable SQL Server compression. For more information on the advantages and disadvantages of compression, see the NetBackup Administrator's Guide, Volume I .
Snapshot Client and Replication Director	This group contains the options that enable backups with Snapshot Client and Replication Director. See "About NetBackup Snapshot Client for SQL Server" on page 118. See "Configuring a VMware policy to protect SQL Server using Replication Director to manage snapshot replication" on page 115.

Schedule properties for SQL Server Intelligent Policies

This topic describes how to configure certain schedule properties for SQL Server Intelligent Policies. Other schedule properties vary according to your specific backup strategy and system configuration. Additional information about other schedule properties is available in the [NetBackup Administrator's Guide, Volume I](#).

[Table 6-2](#) describes how the schedule properties affect a SQL Server Intelligent Policy.

Table 6-2 Description of schedule properties

Property	Description
Type of backup	Specifies the type of backup that this schedule can control. The selection list shows only the backup types that apply to the policy you want to configure. See "Schedule backup types for SQL Server Intelligent Policies" on page 78.

Table 6-2 Description of schedule properties (*continued*)

Property	Description
Schedule type	You can schedule a backup in one of the following ways: ■ Frequency Frequency specifies the period of time that can elapse until the next backup operation begins on this schedule. For example, assume that the frequency is 7 days and a successful backup occurs on Wednesday. The next full backup does not occur until the following Wednesday. Typically, incremental backups have a shorter frequency than full backups. The frequency can be hours, days, or weeks. For transaction log backups, the frequency can also be minutes. ■ Calendar The Calendar option lets you schedule the backup operations that are based on specific dates, recurring week days, or recurring days of the month.
Retention	Specifies a retention period to keep backup copies before they are deleted. The retention period for a schedule controls how long NetBackup keeps records of when scheduled backups occurred. Set the time period to retain at least two full backups of your database. In this way, if one full backup is lost, you have another full backup to restore. The type of schedule you select affects the retention period as follows: ■ Frequency-based scheduling Set a retention period that is longer than the frequency setting for the schedule. For example, if the frequency setting is set to one week, set the retention period to be more than one week. When NetBackup expires a backup image it does not notify SQL Server. Use SQL Server to periodically delete expired backup sets from the SQL Server repository. ■ Calendar-based scheduling The retention period setting is not significant for calendar-based scheduling.
Media multiplexing	Multiplexing is useful if you have many simultaneous backups using the same tape drive. However, it can interfere with SQL Server recovery due to how SQL Server requests streams during restore. In most cases, Cohesity does not recommend multiplexing multiple SQL Server streams from the same backup to a single tape. See “Configuring multiplexed backups of SQL Server” on page 278.

Schedule backup types for SQL Server Intelligent Policies

The **Type of backup** attribute specifies the type of backup that the schedule controls. Refer to the following guidelines when you configure schedules:

- The backup operation is skipped for a specific database if the database recovery model is not supported for the selected backup type.
 See [the section called “Schedules and unsupported recovery models”](#) on page 80.
- If a differential backup runs and a full backup do not already exist for the database or filegroup, NetBackup can convert the backup to a full backup. Similarly, NetBackup can convert transaction log backups if a full backup does not already exist for the database.
 See [“Performance tuning and configuration options ”](#) on page 84.

[Table 6-3](#) shows the backup types that you can specify.

Table 6-3 Schedule backup types for SQL Server Intelligent Policies

Backup type	Description
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.)
Differential incremental backup	A backup of the changed blocks since the last full backup. If you configure a differential incremental backup, you must also configure a full backup.
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. You can choose to turn off truncation in the Microsoft SQL Server tab. See the section called “Configuring high-frequency transaction log backups” on page 79. If you want to configure transaction log backups to run at a high-frequency, review the recommendations. See “Configure the number of jobs allowed for backup operations” on page 36.

Configuring high-frequency transaction log backups

Consider the following when you configure transaction log backups:

- Create a dedicated storage unit for transaction log backup images.

- If a policy includes transaction log backups and full or differential backups, the transaction log backups are run at the scheduled time and frequency. These backups occur even when full or differential backups are active.
- Configure the number of jobs that are allowed for backup operations.
See [“Configure the number of jobs allowed for backup operations”](#) on page 36.

Schedules and unsupported recovery models

NetBackup skips database backups in certain situations. The first case is if the database recovery model for a database does not support the selected backup type. For example, the simple recovery model does not allow transaction log backups. The second case is for the master database, which is skipped for any backups other than full database backups. To back up the master database, you must have a full backup schedule and select **Whole database** in the backup selections. Specifically, the master database is skipped for the following types of backups: differential, filegroup, filegroup differential, file, and transaction log.

In these cases, NetBackup skips the backup of the database, but continues with the backup of the other databases that the policy protects. The backup completes with a status 0 and the job details indicate that the database was skipped.

Example backup schedules for a policy

[Table 6-4](#) shows an example of the schedules you can create for a single SQL Server Intelligent Policy.

Table 6-4 Examples of backup schedules

Schedule	Frequency	Backup window
Full backup	Weekly	Sunday 12 hours
Differential incremental backup	Daily	Monday - Saturday 2 hours in the evening
Transaction log backup	Per your RTO and RPO	Sunday - Saturday 24 hours

Add instances to a policy

This topic describes how to add instances to a policy when you choose the **Protect instances** option. You can also add individual databases to the same policy.

See [“Add databases to a policy”](#) on page 81.

To add instances to a policy

- 1** In the policy, click on the **Instances and databases** tab.
- 2** Click **Protect instances and databases**.
- 3** Click **Add**.
All instances that you registered are displayed.
- 4** In the left pane, select the **Instances** node.
- 5** In the right pane, select the check box next to each instance that you want to add to the list.

Note: Note that for a SQL Server cluster, there is only one entry that is displayed for the cluster. This entry represents all nodes in the cluster; the host is the virtual name of the SQL Server cluster.

- 6** Click **Select**.

The objects you select in the backup selections list apply only to the instances or databases that you add to the list on this tab.

Add databases to a policy

This topic describes how to add databases to a policy when you choose the **Protect instances and databases** option. You can also add instances to the same policy.

Note: Note that when you use legacy (non-intelligent group) policies, NetBackup ignores the **Database state** and backs up any databases that you marked as **Inactive**.

See [“Add instances to a policy”](#) on page 80.

To add databases to a policy

- 1** In the policy, select on the **Instances and databases** tab.
- 2** Select **Protect instances and databases**.
- 3** Select **Add**.
All instances that you registered are displayed.
- 4** In the left pane, expand the **Databases** node and select the instance that contains the databases that you want to protect.

- 5 In the right pane, select the check box next to each database that you want to add to the list.

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

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

- 6 Select the **Select** button.

The objects you select in the backup selections list apply only to the instances or databases that you add to the list on this tab.

Add intelligent groups to a policy

This topic describes how to add intelligent groups to a SQL Server Intelligent Policy.

Note: Intelligent groups replace the legacy method of using instance groups to organize instances in a group.

You can create and add a new intelligent group to a policy or add an existing intelligent group to a policy.

See [“To add an existing intelligent group to a SQL Server Intelligent policy”](#) on page 82.

See [“To add a new intelligent group to a SQL Server Intelligent policy”](#) on page 82.

To add an existing intelligent group to a SQL Server Intelligent policy

- 1 In the policy, select the **Instances and databases** tab.
- 2 Select **Protect intelligent groups**.
- 3 Select the **Add** button.
- 4 Select **Select existing intelligent group**. Then select **Next**.
The intelligent groups that you added are displayed.
- 5 Select the intelligent groups that you want to add to the policy. Then select the **Select** button.

To add a new intelligent group to a SQL Server Intelligent policy

- 1 In the policy, select the **Instances and databases** tab.
- 2 Select **Protect intelligent groups**.

- 3 Select the **Add** button.
- 4 Select **Add new intelligent group**.
- 5 Select **Next**.
- 6 Enter a name for the group and create a query.
- 7 Select **Add**.
- 8 On the **Instances and databases** tab, select **Protect intelligent groups**.
- 9 Select the **Add** button.
- 10 Select **Select existing intelligent group**. Then select **Next**.
 The intelligent groups that you added are displayed.
- 11 Select the intelligent groups that you want to add to the policy. Then select the **Select** button.

Add filegroups or files to the backup selections list

This topic describes how to browse for the filegroups or the files that you want to add to the backup selections list.

To add filegroups or files to the backup selections list

- 1 In the policy, click the **Backup selections** tab.
- 2 Select **Filegroups** or **Files**.
- 3 Click **Add**.
- 4 You can add filegroups or files in the following ways:
 - To manually add the name of the filegroup or file, type the name and then click **Add to list**.
 - To **Browse** the environment for the filegroup or file, click **Browse**.
 In the left pane, select an instance to view the filegroups or files that it contains. The instances, databases, or instance groups that you selected

on the **Instances and databases** tab determine the list of instances that displays here.

- 5 Click **Add** to add the filegroups or files that you selected to the backup selections list.

Note: When you add a filegroup or file to the backup selections list, NetBackup backs up that object for all databases in the policy that contain a filegroup or file with that name.

Add instance groups to a backup policy

This topic describes how to add instance groups to a SQL Server Intelligent Policy.

Note: Instance groups are the legacy method to organize instances in a group and back up that group. Consider using intelligent groups to discover and back up instances.

See [“Add intelligent groups to a policy”](#) on page 82.

To add instance groups to a SQL Server Intelligent policy

- 1 In the policy, select the **Instances and databases** tab.
- 2 Select **Protect instance groups**.
- 3 Select the **Add** button.
- 4 Select the instance groups that you want to add and select the **Select** button.

The list of instance groups that is displayed here controls the instances you can browse and select from when you create the backup selections list.

To see a list of all the instances in the group, select the instance group and select the **Preview instances** button.

Performance tuning and configuration options

The **Microsoft SQL Server** tab in a MS-SQL-Server backup policy contains the tuning parameters that can improve the performance of your backups. These settings, and other factors that affect performance, are discussed in this topic. When you configure a SQL Server protection plan, these options are available on the **Backup options** page. For 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.

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.

Table 6-5 Tuning parameters for SQL Server backups

Field	Description
Availability database backup preference	Note: For policies, the selection on the Select instances and databases tab determines the options you can select in this list. None and Skip availability databases are only available for instances, intelligent groups, and instance groups. The following preference settings are available: ■ None This setting only applies to backup policies and not to protection plans. Perform the backup on the specified instance. ■ 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 to protect only the databases that are not part of an availability group when the policy includes any instances that contain both standalone databases and availability databases.
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 bytes * 2 ^{BLOCK_SIZE} . The value for this option ranges from 0.5 KB to 64 KB. The default is 64 KB.
Backup stripes	See <i>Number of backup stripes</i> .

Table 6-5 Tuning parameters for SQL Server backups (*continued*)

Field	Description
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. The range is 1–32.
Convert differential backups to full (when no full exists)	(Databases only) If no previous full backup exists for the database or filegroup, then NetBackup converts a differential backup to a full backup.
Convert log backups to full (when no full exists)	(Transaction logs only) If no previous full backup exists for the database, then NetBackup converts a transaction backup to a full backup.
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. The default value is cleared except for full database Instant Recovery backups.
Group size for snapshots	(Databases only) This option quiesces a group of databases together and creates a snapshot to back them up as a group. NetBackup automatically discovers and groups any databases that can be grouped, up to this value. A separate snapshot is created for the master database. If both availability databases and standard databases are included in the same policy, separate snapshots are created for each availability group and for the standard databases. If the limit is reached, NetBackup creates additional snapshots. The default value is 1. The range is 1 to 64.
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.

Table 6-5 Tuning parameters for SQL Server backups (*continued*)

Field	Description
SQL Server checksum	Verifies the checksums before the backup. Note that this verification imposes a performance penalty on a backup or restore operation. To verify the checksums before the backup, choose one of the following options: ■ None. This setting only applies to backup policies and not to protection plans. Disables the backup checksums. ■ 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.
Number of backup stripes	This option divides the backup operation into multiple concurrent stripes. 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. The range is 1–32.
Parallel backup operations	This option is the number of backup operations to start simultaneously, per database instance. The range is 1–32. The default is 1. You may need to configure other options when you configure two or more parallel backup operations.
Skip read-only file groups	(Databases only and policies only) This option excludes any filegroups that are read-only from the backup. The resulting backup is a partial image because the image does not contain all filegroups. The partial image contains data from the read-write filegroups and data from the primary filegroup. This option applies only to the Whole database backup selection.
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, and so on NetBackup skips the backup of the unavailable database, but continues with the backup of the other databases that the policy includes. The backup completes with a status 0 and the job details indicate that the database was skipped.
Use 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.
Truncate logs after backup	(Transaction logs only) This option backs up the transaction log and removes the inactive part of the transaction log. This option is enabled by default.

Table 6-5 Tuning parameters for SQL Server backups (*continued*)

Field	Description
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 is 600 for restores. The range is 300–2147483647.

Configure multistriped backups of SQL Server

SQL Server supports backups of databases through multiple data streams, which are called stripes. NetBackup stores each stripe as a separate image. The purpose of this feature is to speed up the rate of data transmission with the use of multiple tape devices.

Backup images can be written to more tapes than available drives. When you restore this type of backup image, in the restore batch file indicate the number of drives that are available.

See [“Restoring multistreamed SQL Server backups”](#) on page 238.

Caution: Do not enable multiplexing for a schedule that is also configured to backup with multiple stripes. Restores fail when multiplexing is enabled for a schedule that uses more than one stripe.

Configure the following to create a multistriped backup:

- In the backup policy, select the number of **Stripes** you want to use.
For a SQL Server Intelligent policy, configure this setting on the **Microsoft SQL Server** tab. For batch file-based SQL Server policies, configure the **Stripes** setting when you create the backup batch file.
- In the schedules for your policy, set **Media multiplexing** to **1** to disable multiplexing.
For batch file-based policies, disable multiplexing in the “Application Backup” schedule. When you disable multiplexing, during a restore all streams are made available simultaneously so the restore operations are successful.
- Ensure that the storage unit has as many drives as you want to have stripes.
- Configure backup schedules so that enough drives are available at the time you want to perform striped backups.

Converting differential backups to full backups

The agent checks to determine if a full backup exists for each database. If no previous full backup exists, a differential backup is converted to a full. For batch file-based policies, you can enable this behavior with the keyword `CONVERTBACKUP`.

The differential backup is converted as follows:

- If you select a database for a differential backup, the backup is converted to a full database backup.
(Intelligent policies) If the **Skip read-only file groups** option is selected the backup is converted to a full read/write filegroup backup.
(Batch file-based policies) If you select **Read-write filegroups** for the **Type of Backup**, the backup is converted to a full read/write filegroup backup.
- (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, you must create a **Full backup** schedule for NetBackup 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 or filegroup. 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 web UI or 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 web UI or the NetBackup MS SQL Client.

Converting log backups to full backups

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. For batch file-based policies, you can enable this behavior with the keyword `CONVERTBACKUP`.

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 web UI or the NetBackup MS SQL Client. Or, if the backup was expired by NetBackup, you can import the full backups into the NetBackup catalog. Then you can restore the full, differential, and log backups with the NetBackup web UI or the NetBackup MS SQL Client.

Back up read-only filegroups

When you separate read-only and read-write filegroups in your backup strategy, you can reduce total media usage and the total time you spend on backup operations. To back up read-only filegroups you must first create a separate policy for this type of backup. You can also verify that all read-only filegroups are backed up.

See [“View SQL Server read-only backup sets \(NetBackup MS SQL Client\)”](#) on page 211.

To back up read-only filegroups

- 1 Create a new policy to protect read-only filegroups.
- 2 Select the policy attributes.
See [“About policy attributes”](#) on page 76.
- 3 Create a **Full** backup schedule and set the **Retention** level to **Infinite**.
All read-only filegroups must be included in some combination of full, or individual filegroup and file backups. You only need to perform this backup one time.
See [“Schedule properties for SQL Server Intelligent Policies”](#) on page 77.
- 4 Choose to protect instances and databases or instance groups.
See [“Add instances to a policy”](#) on page 80.
See [“Add instance groups to a backup policy”](#) on page 84.
- 5 On the **Backup selections** tab, select **Filegroups**.
See [“Add filegroups or files to the backup selections list”](#) on page 83.
- 6 Select the filegroups you want to back up.

- 7 When you complete the policy configuration, click **Create**.
- 8 Back up the read-only filegroups.
- 9 If necessary, confirm that all read-only groups are backed up by viewing the read-only backup set.
See [“View SQL Server read-only backup sets \(NetBackup MS SQL Client\)”](#) on page 211.

Back up read-write filegroups

When you separate read-only and read-write filegroups in your backup strategy, you can reduce total media usage and the total time you spend on backup operations. More information is available on backing up read-only filegroups.

See [“Back up read-only filegroups”](#) on page 90.

Note: Immediately back up any filegroup when you change it from read-write to read-only.

To back up read-write filegroups

- 1 Create a new policy or open the policy you want to configure.
- 2 Select the policy attributes.
See [“About policy attributes”](#) on page 76.
- 3 Create a **Full backup**, **Differential incremental backup**, and **Transaction log backup** schedule.
See [“Schedule properties for SQL Server Intelligent Policies”](#) on page 77.
- 4 On the **Instances and databases** tab, select **Protect instances**.
- 5 Add the instances or the databases that contain the read-write filegroups.
See [“Add instances to a policy”](#) on page 80.
- 6 On the **Backup selections** tab, select **Whole database**.
- 7 Click the **Microsoft SQL Server** tab.
- 8 Select **Skip read-only file groups**.
See [“Performance tuning and configuration options ”](#) on page 84.
- 9 When you have completed the policy configuration, click **Create** or **Save**.

Perform a manual backup

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

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

See [“Use Backup now to back up a SQL Server asset”](#) on page 69.

To perform a manual backup from a policy

- 1 On the left, select **Protection > Policies**.
- 2 Select the policy you want to test.
- 3 Select **Manual backup**.
- 4 Select the schedule that you want to use for the manual backup.
- 5 For SQL Server Intelligent Policies, select the databases or instances that you want to include for the manual backup. For batch file-based policies, select the clients that you want to include for the manual backup.

Protecting SQL Server availability groups

This chapter includes the following topics:

- [About protecting SQL Server availability groups](#)
- [Protecting SQL Server availability groups with intelligent policies](#)
- [Protecting SQL Server availability groups with batch file-based policies](#)
- [Protect a SQL Server availability group that crosses NetBackup domains with a policy](#)

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.
- With a policy 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 with a policy”](#) on page 108.

Note the following before you configure the policy or 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

NetBackup does not support the following types of backups for availability databases:

- Snapshot backups of filegroups or files
- Instant Recovery backups (configured with policies)
- VMware backups

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.
- Certain differential backups that are performed with a SQL Server Intelligent policy.
 - A differential backup is skipped when an availability group is added as a backup selection.
 - A differential backup fails when an availability database is added as a backup selection.
- A differential backup that is performed with a batch file-based policy.
Backups of this type result in a failed backup.
- Copy-only transaction log backups
Backups of this type result in a failed backup.
- (SQL Server 2022) An issue exists with filegroup backups for the replicas that participate in an Always On availability group that is marked as non-readable. They are not supported at this time due to a Microsoft limitation.

Protecting SQL Server availability groups with intelligent policies

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

- With an intelligent policy 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 with a policy”](#) on page 108.
- NetBackup supports backups of availability groups in multi-NIC environments. For more information, see the following topic:
See [“Configuration and requirements for SQL Server backups with multiple NICs”](#) on page 245.

Prerequisites for protecting SQL Server availability groups

Before you configure protection for availability groups, review and complete the following prerequisites. Perform the steps after you create the SQL Server availability group.

See [Table 7-1](#)

Table 7-1 Prerequisites for protecting the preferred or the primary replica in an availability group

Step	Action	Description
Step 1	Register the credentials for the availability replicas.	See “Register a SQL Server instance with an existing credential” on page 48. See “Register a SQL Server instance with a new credential” on page 49. See “Add a credential rule” on page 52.
Step 2	Review the mappings for the hosts in your environment.	Approve each valid auto-discovered mapping that NetBackup discovers in your environment. Perform this configuration in the Security > Host mappings properties on the primary server. See “Reviewing the auto-discovered mappings” on page 26.
Step 3	Configure the mappings for distributed application restores.	For basic and advanced availability groups, map the WSFC (Windows Server Failover Cluster) name to each availability group node. If you have an availability group with an FCI, you must configure additional mappings. Configure these mappings in the Distributed application restore mapping host property on the primary server. See “Configuring mappings for restores of a distributed applications, clusters, or virtual machines ” on page 31.

Configure a backup policy to protect a SQL Server availability group

You can create a backup policy to perform scheduled backups of a SQL Server availability group. By default, NetBackup performs backups on the primary replica. Alternatively, you can protect the preferred replica.

To configure a backup policy for the preferred or the primary replica of a SQL Server availability group

- 1 On the left, select **Protection > Policies**.
- 2 Select **Add**.
- 3 On the **Attributes** tab, configure the following:
 - Select the **MS-SQL-Server** policy type.
 - Specify a storage unit.

See [“About policy attributes”](#) on page 76.
- 4 On the **Instances and databases** tab, select **Protect availability groups** or **Protect intelligent groups**.
- 5 Select **Add**.
- 6 Select the availability groups, availability databases, or intelligent groups that you want to protect.

See [“Add an availability group to a policy”](#) on page 97.

See [“Add availability databases to a policy”](#) on page 97.

See [“Add intelligent groups to a policy”](#) on page 82.
- 7 Add schedules.

See [“Schedule properties for SQL Server Intelligent Policies”](#) on page 77.
- 8 Select the **Microsoft SQL Server** tab.
- 9 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 84.
- 10 (Optional) Make any other changes to the tuning parameters.

See [“Performance tuning and configuration options ”](#) on page 84.
- 11 Select **Create** to create the policy.

Add an availability group to a policy

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

To add an availability group to a policy

- 1 On the **Instances and databases** tab, click **Protect availability groups**.
- 2 Click **Add**.
All availability groups that you registered are displayed.
- 3 In the left pane, select the **Availability groups** node.
- 4 In the right pane, select the check box next to each availability group that you want to add to the list.
- 5 Click **Select**.

When you select an availability group, all databases in the availability group are included in the backup.

The objects you select in the backup selections list apply only to the availability groups or availability databases that you add to the list on this tab.

Add availability databases to a policy

This topic describes how to add availability databases to a policy when you choose the **Protect availability groups** option. You can also add availability groups to the same policy. If you want to back up databases outside the availability group, you must create a different policy for those databases.

To add availability databases to a policy

- 1 On the **Instances and databases** tab, select **Protect availability groups**.
- 2 Select **Add**.
All availability groups that you registered are displayed.
- 3 In the left pane, expand the node for the availability group that contains the databases that you want to protect.
- 4 In the left pane, select a replica.
- 5 In the right pane, select the check box next to each database that you want to add to the list.

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

- 6 Select the **Select** button.

The objects you select in the backup selections list apply only to the availability groups or availability databases that you add to the list on the **Instances and databases** tab.

Add filegroups or files in an availability database to the backup selections list

This topic describes how to browse for filegroups or files, which are part of an availability group, that you want to add to the backup selections list.

To add filegroups or files in an availability group to the backup selections list

- 1 On the **Instances and databases** tab, select **Protect availability groups**.
- 2 Select **Add**.

All availability groups that you registered are displayed.
- 3 In the left pane, expand the node for the availability group that contains the databases that you want to protect.
- 4 In the left pane, select a replica.
- 5 In the right pane, select the check box next to each database that you want to add to the list.

If you choose to select individual databases, you must manually add any new databases in your environment to a policy. In this case, NetBackup does *not* dynamically create a list of databases at run-time.
- 6 Select the **Select** button.

The objects you select in the backup selections list apply only to the availability groups or availability databases that you add to the list on the **Instances and databases** tab.
- 7 On the **Backup selections** tab, select **Filegroups** or **Files**.
- 8 Select **Add**.
- 9 Select **Browse**.
- 10 In the left pane, expand the availability group and select the replica.
- 11 In the right pane select the filegroups or files.

- 12 Select **Add**.
- 13 To add the filegroups or files that you selected to the backup selections list, select **Add**.

Note: When you add a filegroup or file to the backup selections list, NetBackup backs up that object for all databases in the policy that contain a filegroup or file with that name.

Protecting SQL Server availability groups with batch file-based policies

You can use batch file-based policies to protect an availability group environment in the following ways:

- With a policy that protects the preferred replica.
- With a policy that protects a specific node in the availability group.
- 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 with a policy”](#) on page 108.
- NetBackup supports backups of availability groups in multi-NIC environments. For more information, see the following topic:
See [“Configuration and requirements for SQL Server backups with multiple NICs”](#) on page 245.

About protecting the preferred replica in a SQL Server availability group (batch file-based policies)

You can use a SQL Server Intelligent Policy to protect the preferred or primary replica in a SQL Server availability group. Note the following before you configure the policy:

- To protect the preferred replica, use the `PREFERREDREPLICA PREFERRED` keyword. NetBackup honors your SQL Server backup preferences. These preferences include the preferred replica, backup priority, and excluded replicas. NetBackup backs up the preferred replica, as determined by SQL Server.
- To protect the primary replica, use the `PREFERREDREPLICA PRIMARY` keyword.

- NetBackup can only fully protect the availability group environment if the backup policy includes each replica on which backups occur in the **Clients** list. Also, all batch files in the **Backup selections** list must exist on each replicas on which backups occur.
- Note that a backup job runs on each replica in the availability group. On replicas which are not the backup source, the job skips the backup.
- Review the information on support and limitations for availability groups. See [“About protecting SQL Server availability groups”](#) on page 93.
- Review the prerequisites for protecting the availability group. See [“Prerequisites for protecting SQL Server availability groups”](#) on page 95.

Prerequisites for protecting SQL Server availability groups

Before you configure policies to protect availability groups with batch file-based policies, review and complete the following prerequisites. Perform the steps after you create the SQL Server availability group.

See [Table 7-2](#)

Table 7-2 Prerequisites for protecting the preferred replica in an availability group

Step	Action	Description
Step 1	On each replica where you want backups to occur, configure the NetBackup services.	See “Configure the NetBackup services for SQL Server backups and restores (batch file-based policies)” on page 188.
Step 2	Configure the mappings for distributed application restores.	For basic and advanced availability groups, map the WSFC (Windows Server Failover Cluster) name to each availability group node. If you have an availability group with an FCI, you must configure additional mappings. Configure these mappings in the Distributed application restore mapping host property on the primary server. See “Configuring mappings for restores of a distributed applications, clusters, or virtual machines” on page 31.

Table 7-2 Prerequisites for protecting the preferred replica in an availability group *(continued)*

Step	Action	Description
Step 3	Review the auto-discovered mappings for the hosts in your environment.	Approve each valid auto-discovered mapping that NetBackup discovers in your environment. Perform this configuration in the Security > Host mappings properties on the primary server. See “Reviewing the auto-discovered mappings” on page 26.

Configure an automatic backup policy for the preferred or the primary replica of a SQL Server availability group

This topic describes how to create a backup policy that uses batch files for automatic (scheduled) backups of the preferred or the primary replica in a SQL Server availability group. Create a policy for each type of backup that you want to perform. For example:

- Policy A

Schedules: Full backup, run weekly
Backup Selections: Batch file for full backups
Clients: Node A, Node B, Node C
- Policy B

Schedules: Differential backup, run daily
Backup Selections: Batch file for differential backups
Clients: Node A, Node B, Node C
- Policy C

Schedules: Full backup, run per your RTO and RPO
Backup Selections: Batch file for transaction log backups
Clients: Node A, Node B, Node C

To configure an automatic backup policy for the preferred or the primary replica of a SQL Server availability group

- 1

Sign in to the primary server as a user that has the RBAC Administrator role or a role that can manage policies.
- 2

On the left, select **Protection > Policies**.
- 3

Select **Add**.
- 4

Type a unique name for the new policy.
- 5

On the **Attributes** tab, configure the following:

- Select the **MS-SQL-Server** policy type.
- Specify a storage unit.

See [“About policy attributes”](#) on page 76.

6 On the **Instances and databases** tab, select **Clients for use with batch files**.

The tab name changes to the name **Clients**. The **Backup selections** tab now lets you specify and browse for scripts.

7 On the **Schedules** tab, add a **Full backup** schedule.

NetBackup also creates a Default-Application-Backup schedule. Use this schedule to set the retention level for the policy.

See [“Schedule backup types for batch file-based policies”](#) on page 202.

8 On the **Clients** tab, add the name of each replica on which you want backups to occur.

Use the NetBackup client name for each replica. If a replica is hosted on a failover cluster instance (FCI), use the instance cluster name.

9 Repeat step **3** to **8** in this procedure to create a policy for each type of backup (full, differential, transaction log) that you want to perform.

Each type of backup requires a separate policy.

10 On each replica where you want to perform backups, create a batch file for each type of backup that you want to perform.

See [“Create batch files for the policy that protects the preferred or the primary replica”](#) on page 102.

Create batch files for the policy that protects the preferred or the primary replica

This topic describes how to create batch files for the backup policies that protect the availability group. These batch files can use either the `PREFERREDREPLICA` or `PREFERREDREPLICA PRIMARY` to protect either the preferred or the primary replica.

To create the batch files for an availability group, you must log on to each replica separately. Then use the NetBackup MS SQL Client to create the batch files on each replica.

To create batch files for the policy that protects the preferred replica

- 1** This procedure assumes that you already created a separate policy for each type of backup that you want to perform.

See [“Configure an automatic backup policy for the preferred or the primary replica of a SQL Server availability group”](#) on page 101.

- 2** Perform steps 3 to 14 in this procedure on each replica in the availability group.

You must log on to each replica separately and create the batch files from that replica. This way the batch files have the correct settings for each node. Backups may fail if you create a batch file on one replica and copy it to the other replicas in the availability group.

- 3** Log on to one of the replicas in the availability group.
- 4** Open the NetBackup MS SQL Client.
- 5** Select **File > Set SQL Server connection properties**.
- 6** From the **Instance** drop-down list, select the instance that hosts the availability group.
- 7** Select **File > Backup SQL Server objects**.
- 8** Select the objects you want to backup in one of the following ways:
 - Select one or more databases, filegroups, or files.
 - To back up all databases, including the system databases (`DATABASE $ALL`), select the instance. From the **Back up** group, select **All**.
- 9** Select the **Type of Backup** and any other settings.
- 10** In the **NetBackup Policy** field, enter the name of the MS-SQL-Server policy that you created.
- 11** From the **Backup script** group, select **Save**.
- 12** Select **Backup** and open the batch file.
- 13** For each operation in the batch file configure one of the following options:
 - To protect the preferred replica, add the keyword `PREFERREDREPLICA PREFERRED`.
 - To protect the primary replica, add the keyword `PREFERREDREPLICA PRIMARY`.
- 14** Save and close the batch file.

Note the location of the batch file. Save the batch file for each replica to the same file location. This way you only need to enter one file location for the batch file in the **Backup Selections** list.

- 15 Repeat steps 7 to 14 for any other types of backups that you want to perform. For example, full, differential, or transaction log.

More information is available on how to create batch files.

See [“Requirements to use batch files with NetBackup for SQL Server”](#) on page 190.

- 16 Repeat the steps in this procedure (steps 3 to 15) to create batch files for the other availability group replicas.

- 17 When you have created batch files for all the replicas on which you want backups to occur, add the batch files to the policies that you created previously.

See [“Add the batch files to the policy that protects the preferred or the primary replica”](#) on page 104.

Add the batch files to the policy that protects the preferred or the primary replica

This topic describes how to add the batch files that you created to the backup policy that protects the preferred or the primary replica in the availability group.

To add the batch files to the policy that protects the preferred or the primary replica

- 1 This procedure assumes that you already created a policy. It also assumes that you created batch files on each replica on which you want backups to occur.

See [“Configure an automatic backup policy for the preferred or the primary replica of a SQL Server availability group”](#) on page 101.

See [“Add the batch files to the policy that protects the preferred or the primary replica”](#) on page 104.
- 2 Open the policy that you created.
- 3 On the **Backup selections** tab, add the batch files that you created. If you saved the batch files to the same location on each replica, you need only one entry in the **Backup selections** list.

Include batch files for only one type of backup in this policy. (For example, full, differential, or transaction log.)
- 4 Select **Save** to save the policy.
- 5 Repeat the step 2 through 4 in this procedure for each policy that you created.

About protecting a specific node in a SQL Server availability group (batch file-based policies)

This topic describes how to protect a specific node in a SQL Server availability group using a batch file-based policy.

Note the following when you configure a NetBackup policy to protect a specific node in an availability group:

- For this backup scenario, do not use the `PREFERREDREPLICA TRUE, PRIMARY, or PREFERRED` keyword in your batch files. Backups are skipped if the backup policy does not include the node that hosts the preferred replica.
- Review the information on support and limitations for availability groups. See [“About protecting SQL Server availability groups”](#) on page 93.

Configure an automatic backup policy for a specific replica of a SQL Server availability group

This topic describes how to create a backup policy for automatic (scheduled) backups of a specific replica in a SQL Server availability group. Create a policy for each type of backup that you want to perform. For example:

Policy A	Schedules: Full backup, run weekly Backup Selections: Batch file for full backups Clients: Node A
Policy B	Schedules: Full backup, run daily Backup Selections: Batch file for full differential backups Clients: Node A
Policy C	Schedules: Full backup, run per your RTO and RPO Backup Selections: Batch file for transaction log backups Clients: Node A

To configure an automatic backup policy for a specific replica of a SQL Server availability group

- 1 Sign in to the primary server as a user that has the RBAC Administrator role or a role that can manage policies.
- 2 On the left, select **Protection > Policies**.
- 3 Select **Add**.
- 4 On the **Attributes** tab, configure the following:

- Select the **MS-SQL-Server** policy type.
- Specify a storage unit.

See [“About policy attributes”](#) on page 76.

5 On the **Instances and databases** tab, select **Clients for use with batch files**.

The tab name changes to the name **Clients**. The **Backup selections** tab now lets you specify and browse for scripts.

6 On the **Schedules** tab, add a **Full backup** schedule.

NetBackup also creates a Default-Application-Backup schedule. Use this schedule to set the retention level for the policy. See the [NetBackup Administrator's Guide, Volume I](#) for more information.

See [“Schedule properties for SQL Server batch file-based policies”](#) on page 201.

7 On the **Clients** tab, add the name of the replica that you want to protect.

Use the NetBackup client name for the replica. If a replica is hosted on a failover cluster instance (FCI), use the instance cluster name.

8 Select **Create** to save the policy.

9 Repeat the step [3](#) through step [7](#) in this procedure to create a policy for each type of backup (full, full differential, transaction log) that you want to perform.

Each type of backup requires a separate policy.

10 Create a batch file for each type of backup that you want to perform with each policy.

See [“Create a batch file for the policy that protects a specific availability replica in an availability group”](#) on page 106.

Create a batch file for the policy that protects a specific availability replica in an availability group

This topic describes how to create batch files for the backup policies that protect a specific availability replica in the availability group.

To create batch files for the policy that protects a specific replica

1 This procedure assumes that you already created a policy.

See [“Configure an automatic backup policy for a specific replica of a SQL Server availability group”](#) on page 105.

2 Log on to the availability replica you want to protect.

3 Open the NetBackup MS SQL Client.

4 Select **File > Set SQL Server connection properties**.

- 5 From the **Instance** drop-down list, select the instance that hosts the availability group.
- 6 Select **File > Backup SQL Server objects**.
- 7 Select the objects you want to backup in one of the following ways:
 - Select one or more databases, filegroups, or files.
 - To back up all databases, including the system databases (`DATABASE $ALL`), select the instance. From the **Back up** group, select **All**.
- 8 Select the **Type of Backup** and any other settings.
- 9 In the **NetBackup Policy** field, enter the name of the MS-SQL Server policy that you created.
- 10 From the **Backup script** group, select **Save**.
- 11 Select **Backup** and save the batch file.

Do not use the `PREFERREDREPLICA TRUE`, `PRIMARY`, or `PREFERRED` keyword in your batch files. Backups are skipped if the backup policy does not include the node that hosts the preferred replica.
- 12 Repeat steps 6 to 11 for any other the types of backups that you want to perform. For example, full, full differential, or transaction log.

More information is available on how to create batch files.

See [“Requirements to use batch files with NetBackup for SQL Server”](#) on page 190.
- 13 When you have created all the batch files, add these files to the policies that you created previously.

See [“Add the batch files to the policy that protects a specific replica in the availability group”](#) on page 108.

Add the batch files to the policy that protects a specific replica in the availability group

To add the batch files to the policy that protects a specific replica in the availability group

- 1 This procedure assumes that you already created a policy and created batch files for a specific replica in the availability group.

 See [“Configure an automatic backup policy for a specific replica of a SQL Server availability group”](#) on page 105.

 See [“Create a batch file for the policy that protects a specific availability replica in an availability group”](#) on page 106.
- 2 Open the policy that you created.
- 3 On the **Backup selections** tab, add the batch files that you created.

 Include batch files for only one type of backup in this policy. (For example, full, full differential, or transaction log.)
- 4 Select **Save** to save the policy.
- 5 Repeat the steps 2 through 4 in this procedure for each policy that you created.

Protect a SQL Server availability group that crosses NetBackup domains with a policy

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.

Protecting SQL Server with VMware backups

This chapter includes the following topics:

- [About protecting an application database with VMware backups](#)
- [About configuring NetBackup for VMware backups that protect SQL Server](#)
- [Configuring a VMware backup policy to protect SQL Server](#)
- [Configuring a VMware policy to protect SQL Server using Replication Director to manage snapshot replication](#)
- [Protect SQL Server data with 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:

- Choose whether or not to truncate logs.
- Use the existing database restore process to restore and recover data from VMware backups.
- From one VMware backup, choose from these restore options: Volume-level restore, file-level recovery, or database restore.
- With the **Enable T-SQL snapshots** option, NetBackup creates a full SQL Server backup (not copy-only) that can be used as a basis for SQL Server incremental and transaction log backups. T-SQL snapshots were added with SQL Server

2022 and are strongly recommended for systems with SQL Server 2022 or later. This type of snapshot is not supported with SQL Server 2019 and earlier.

- 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://support.cohesity.com/s/>

Veritas VSS Provider

Cohesity 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 VSS provider for vSphere”](#) on page 23.

The Veritas VSS Provider allows the 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.

Disabling the SQL Server VSS Writer service

For VMware backups with T-SQL snapshots, you must also disable the SQL Server VSS Writer service.

See [“Disable the SQL Server VSS Writer service ”](#) on page 23.

Using NetBackup Accelerator to increase the speed of full VMware backups

Select the **Use Accelerator** policy option to use NetBackup Accelerator to potentially increase the speed of full VMware backups. (This option is not available in the settings for a protection plan.) By reducing the backup time, it is easier to perform the VMware backup within the backup window. To use this feature, you must first perform an initial backup with **Use Accelerator** enabled. Subsequent backup times can then be significantly reduced. Accelerator support for database agents currently restricts backups to the full schedule type.

To periodically establish a new baseline of change detection on the client, create a separate policy schedule with the **Accelerator forced rescan** option enabled.

For more details on Accelerator with VMware backups, see the [NetBackup for VMware Administrator's Guide](#).

For Accelerator to work with a VMware policy with application protection for SQL Server, a successful and complete previous backup is required to gain acceleration optimization. A backup job or child job (including the ASC) can have an incomplete status (Status 1). In that case, the backup image is not considered as a base backup for a subsequent job and acceleration optimization for that job is zero.

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 or policy for SQL Server incremental backups.
- SQL Server clusters or SQL Server availability groups.
- (NetBackup web UI only) 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.
- For VMware backups with the T-SQL snapshots, the following limitations apply:
 - T-SQL snapshots require SQL Server 2022. If there are multiple SQL Server instances of different versions (for example, 2019 and 2022) on the guest virtual machine and T-SQL snapshots are enabled, the policy only protects the SQL Server 2022 instances or databases.

- NetBackup limits the number of databases that can be processed simultaneously to 62.
- Only user databases are protected. System databases cannot be protected with this method. This limitation is from Microsoft. (A policy can contain system databases, but NetBackup skips these databases.)
- Log truncation is not supported with T-SQL snapshots.

About configuring NetBackup for VMware backups that protect SQL Server

Table 8-1 Steps to configure VMware backups that protect SQL Server

Step	Action	Description
Step 1	Configure the logon account for the NetBackup services.	The logon account for the NetBackup Client Service and the NetBackup Legacy Network Service must meet certain requirements. See “Configure the NetBackup services for SQL Server backups and restores” on page 24. See “Configure local security privileges for SQL Server” on page 25.
Step 2	If you want to use Replication Director to manage your VMware snapshots and snapshot replicas, create a storage lifecycle policy (SLP).	See the NetBackup Replication Director Solutions Guide .
Step 3	Configure a VMware policy.	See “Configuring a VMware backup policy to protect SQL Server” on page 113. See “Configuring a VMware policy to protect SQL Server using Replication Director to manage snapshot replication” on page 115.
Step 4	If you use a Primary VM identifier other than VM hostname , you need to map that identifier to the host name of the VM.	Configure this mapping in the Distributed Application Restore Mapping host property on the primary server. See “Configuring mappings for restores of a distributed applications, clusters, or virtual machines” on page 31.
Step 5	Review the auto-discovered mappings for the hosts in your environment.	Approve each valid Auto-Discovered Mapping that NetBackup discovers in your environment. Perform this configuration in the Host Management properties on the primary server. See “Reviewing the auto-discovered mappings” on page 26.

Configuring a VMware backup policy to protect SQL Server

Through a VMware backup policy, NetBackup can create full application-consistent backups of the SQL Server databases that reside on a virtual machine. Optionally you can use NetBackup Accelerator. VMware policies let you exclude certain virtual disks from the VMware backup. If you want to exclude specific SQL Server components, use a MS-SQL-Server policy.

To truncate logs, you must first perform a full VMware backup without log truncation. When this backup is complete, then enable log truncation in the policy.

Note that before you create a policy, you must perform additional configuration requirements.

See [“About configuring NetBackup for VMware backups that protect SQL Server”](#) on page 112.

More information on Accelerator is available:

See [“About policy attributes”](#) on page 76.

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

To configure a VMware backup policy to protect SQL Server

- 1 On the left, select **Protection > Policies**.
- 2 Add a new policy or open the policy that you want to edit.
- 3 Select the **Attributes** tab.
 - From the **Policy type** list, select **VMware**.
 - In the **Policy storage** list, select a disk storage unit.
If you want to use NetBackup Accelerator, select a supported storage unit type. The NetBackup device mapping files list all supported storage types.
 - If you want to use NetBackup Accelerator, select **Use Accelerator**.
Accelerator uses the initial full backup to establish a baseline. Any subsequent backups that are performed with Accelerator can run significantly faster. You may want to create an additional policy schedule that enables the **Accelerator forced rescan** option. This option establishes a new baseline for the next Accelerator backup.
Perform block-level incremental backups is automatically selected and grayed out. On the **VMware** tab, the **Enable block-level incremental backup** option is also selected and grayed out.
- 4 On the **Schedules** tab, create a schedule for full backups.
- 5 On the **Clients** tab, do the following:

- Select **Select automatically through VMware intelligent policy query**.
 - From the **NetBackup host to perform automatic virtual machine selection** list, select the host you want to use.
 - Use the Query builder to create the rules that select the virtual machines you want to back up.
- 6** On the **VMware** tab:
- Select the **Primary VM identifier** to use to catalog the backups.
 - Select **Enable file recovery from VM backup**.
 - Locate **Application protection** and select **Microsoft SQL Server**.
This option allows recovery of the databases from the virtual machine backups. If this option is disabled, you can recover the entire virtual machine from the backup, but you cannot recover the databases individually.
 - Do not enable **Truncate logs** at this time. You must first perform a full backup without log truncation, described later in this procedure.
 - (Conditional) Select **Enable T-SQL snapshots**.
This snapshot type creates a full SQL Server backup (not copy-only) that can be used as a basis for SQL Server incremental and transaction log backups. For this snapshot type, the only valid option for snapshot handling is **Stop the backup if any snapshots exist**, which is automatically selected.
Note: T-SQL snapshots were added with SQL Server 2022 and are strongly recommended for systems with SQL Server 2022 or later. This type of snapshot is not supported with SQL Server 2019 and earlier.
- 7** If you want to exclude certain disks from the VMware backup, select the **Exclude disks** tab.
- NetBackup excludes those disks from the VMware backup that protects SQL Server. Be sure that any disks that you exclude do not contain database data.
- 8** Select **Save** to save the policy.
- If you do not want to truncate transaction logs, no further action is necessary.
- If you want to truncate transaction logs, continue with step [9](#).
- 9** Perform a full backup without log truncation.
- When the backup completes, open the policy that you created in step [2](#).
- 10** Select the **VMware** tab.

- 11 Locate **Application protection** and select **Microsoft SQL Server**. Then select **Truncate logs**.

For SQL Server, this option truncates the transaction logs when the VMware snapshot of the virtual machine is complete.

- 12 Select **Save** to save the policy.
- 13 Perform a full VMware backup.

Configuring a VMware policy to protect SQL Server using Replication Director to manage snapshot replication

This topic describes how to configure a VMware policy to back up SQL Server using Replication Director to manage snapshot replication. Note that NetBackup must have access to the CIFS share on the NetApp disk array. For more details on VMware policies, see the [NetBackup for VMware Administrator's Guide](#).

For complete details on how to configure Replication Director with VMware backups, see the [NetBackup Replication Director Solutions Guide](#).

To configure a VMware policy to back up SQL Server using Replication Director to manage snapshot replication

- 1 On the left, click **Protection > Policies**.
- 2 Create a policy or open the policy you want to configure.
- 3 Click the **Attributes** tab.
 - From the **Policy type** list, select **VMware**.
 - In the **Policy storage** list select the storage lifecycle policy (SLP) that you want to use. This SLP must be configured for snapshot replication.
 - In the **Snapshot Client and Replication Director** group, click **Use Replication Director**.
- 4 On the **Schedules** tab, create a schedule for full backups.
- 5 On the **Clients** tab, do the following:
 - Click **Select automatically through query**.
 - Use the Query Builder to create the rules that select the virtual machines you want to back up.
 - Select **NetBackup host to perform automatic virtual machine selection** and the host you want to use.

- 6 On the **VMware** tab, enable the following options:
 - **Primary VM identifier** to use to catalog the backups.
 - **Enable file recovery from VM backup.**
This option allows for application protection of SQL Server.
 - Locate **Application protection** and click **Microsoft SQL Server**.
This option allows recovery of the SQL databases from the virtual machine backups. If this option is disabled, you can recover the entire virtual machine from the backup, but you cannot recover the databases individually.
 - Do not enable **Truncate logs** at this time. You must first perform a full backup without log truncation, described later in this procedure.
 - (Conditional) Select **Enable T-SQL snapshots**.
This snapshot type creates a full SQL Server backup (not copy-only) that can be used as a basis for SQL Server incremental and transaction log backups. For this snapshot type, the only valid option for snapshot handling is **Stop the backup if any snapshots exist**, which is automatically selected.
Note: T-SQL snapshots were added with SQL Server 2022 and are strongly recommended for systems with SQL Server 2022 or later. This type of snapshot is not supported with SQL Server 2019 and earlier.
- 7 Click **Save** to save the policy.
If you do not want to truncate transaction logs, no further action is necessary.
If you want to truncate transaction logs, continue with step 8.
- 8 Perform a full backup without log truncation.
When the backup completes, open the policy that you created in step 2.
- 9 Click the **VMware** tab and under **Microsoft SQL Server**, select **Truncate logs**.
- 10 Click **Save** to save the policy.
- 11 Perform a full VMware backup.

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 137.
 - **Backup options**
 Adjust the server or host to use for backups, snapshot options, and exclude options.
 See “[Backup options and Advanced options](#)” on page 144.
 - **Advanced options**
 Change or enable any advanced options for the protection plan.
 See “[Backup options and Advanced options](#)” on page 144.
 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**.

Configuring backup policies with Snapshot Client

This chapter includes the following topics:

- [About NetBackup Snapshot Client for SQL Server](#)
- [How SQL Server operations use Snapshot Client](#)
- [Snapshot methods](#)
- [Configuration requirements for SQL Server snapshot and Instant Recovery backups](#)
- [Configure a snapshot policy for SQL Server](#)
- [Configure a policy for Instant Recovery backups of SQL Server](#)
- [Using copy-only snapshot backups to affect how differentials are based](#)
- [About SQL Server agent grouped snapshots](#)

About NetBackup Snapshot Client for SQL Server

NetBackup for SQL Server includes support for snapshot backups. The snapshot technology uses SQL Server VDI (virtual device interface) quiescence to affect a momentary freeze on database activity. Then the agent can back up and restore SQL Server objects by taking snapshots of the component files. Data is captured at a particular instant. The resulting snapshot can be backed up without affecting the availability of the database. These snapshots are backed up to the storage unit.

A separate Snapshot Client license provides additional features for snapshot backups. You can configure the snapshot image for Instant Recovery and you can configure an alternate client to perform the snapshot backup.

The following NetBackup Snapshot Client features are available for use with NetBackup for SQL Server:

Snapshot backup	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.
Instant Recovery	Makes the backups available for recovery from the local disk. The snapshot can also be the source for an additional backup copy to tape or other storage.
Off-host backup	Shifts the burden of backup processing onto a separate backup agent, reducing the backup impact on the client's computing resources. The backup agent sends the client's data to the storage device.

How SQL Server operations use Snapshot Client

This topic describes how SQL Server operations use the Snapshot Client.

The following topics describe how NetBackup for SQL Server works with the Snapshot Client option:

- [About selection of backup method](#)
- [About SQL Server limitations with snapshots](#)
- [About Snapshot Client and SQL Server performance considerations](#)
- [About SQL Server snapshot backups](#)
- [About SQL Server snapshot restores](#)

About selection of backup method

The selection of a backup methodology, whether standard or Snapshot Client, is dependent on what policy is used. If a policy configured for Snapshot Client is selected, then additional attributes of policy determine the Snapshot Client features. It also determines the specific snapshot methods that are used.

About SQL Server limitations with snapshots

Due to SQL Server limitations certain objects cannot be backed up by snapshots. These are database differentials, filegroup differentials, and transaction logs. If a Snapshot Client policy is selected to back up one of these object types, then

NetBackup performs a stream-based backup. NetBackup uses the storage unit that is provided in the policy configuration. If a storage unit is not provided, then NetBackup uses the default storage unit for the server.

What is backed up by NetBackup for SQL Server

The database administrator works exclusively with logical objects, such as databases and filegroups. However, it is useful to understand the differences between file- and stream-based backups in terms of the data content that is archived. For stream-based backups, NetBackup captures the data stream content that is provided by SQL Server. If the user has specified multiple streams, then SQL Server opens multiple streams that NetBackup catalogs as separate images.

For file-based backups, NetBackup creates a file list that consists of all the physical files that constitute the object. This file list is supplied to the Snapshot Client, which is responsible for snapshot creation. If multiple streams are specified, then NetBackup divides the file list into sub-lists. Each sub-list is backed up separately and constitutes a separate image. Users may notice that if multiple streams are specified for a file-based backup and if the number of streams exceeds the number of component files, then the number of file-based streams does not exceed the number of files. With stream-based SQL Server backups, SQL Server always creates exactly the number of streams that the end user specifies.

The file list that is used to back up a SQL Server database consists of the physical files that constitute the primary filegroup. The file list also consists of any secondary filegroups, and the transaction log. Typically, these can be identified respectively by their name extensions, which are `.mdf`, `.ndf`, and `.ldf`. The file list for a filegroup backup consists of the physical files that belong to the filegroup. And, finally, the file list for a file object backup consists of a single physical file. This file is the file that maps to the SQL Server file object.

About Snapshot Client and SQL Server performance considerations

When a physical file is backed up with the Snapshot Client, the backup consists of the entire extent. This backup contrasts with stream-based SQL Server backups where only the actual data content of the objects are archived. If you intend to use snapshot technology to back up SQL Server, you may want to use the SQL Server dynamic file allocation. This configuration reduces the likelihood that any of the component files contain large areas of empty space.

Also review the other considerations for SQL Server disk initialization.

See [“NetBackup for SQL Server performance factors”](#) on page 257.

About SQL Server snapshot backups

No special interfacing considerations exist when you perform Snapshot Client backups of SQL Server. A snapshot backup is performed if the backup object is: a database, a filegroup, or a file and a policy is selected and configured for Snapshot Client. If a differential backup or transaction log backup is tried with a Snapshot Client backup, then the operation uses the selected policy. But a standard database backup is performed with the configured storage unit.

About SQL Server snapshot restores

Any backup images that were created from snapshots display along with standard backup images. That is, all backup items—without regard to method—display in a time-sequenced ordering that respects the composition of the database hierarchy. In addition, no weighting is given in to determine an optimal recovery that is based on the backup method.

Snapshot methods

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

Although all of these features are provided through Snapshot Client support for SQL Server, not all snapshot methods are supported. For a description of snapshot methods available for use with NetBackup for SQL Server, see the [Hardware and Cloud Storage Compatibility List \(HCL\)](#).

Table 9-1

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 9-1 (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. This option is not supported for MS-SQL-Server policies. - 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.

Configuration requirements for SQL Server snapshot and Instant Recovery backups

Review the following requirements before you configure NetBackup for SQL Server with snapshot backups:

- See the [NetBackup Compatibility Lists](#) for details on the hardware requirements and software requirements for the snapshot method that you want to use.
- Go to the Cohesity Support website for details on the snapshot methods and platforms that are supported for NetBackup for SQL Server.

- The volumes which contains the SQL Server databases and log files should be dedicated to SQL Server only. Other types of databases (for example, Exchange) should not reside on the volumes.
- For information on using Snapshot Manager, see the [NetBackup Snapshot Manager for Data Center Administrator's Guide](#).
- Only one snapshot method can be configured per policy. If you want to use a different snapshot method different clients, then create a separate policy for each group of clients and the snapshot method you want to use. Then select one method for each policy.
- NetBackup does not support Instant Recovery with availability groups.

Configure a snapshot policy for SQL Server

These instructions describe how to configure a Snapshot Client policy. Optionally you can choose to perform an off-host backup. This topic only covers what is necessary to configure snapshot backups for a MS-SQL-Server policy.

See [“About SQL Server Intelligent Policies”](#) on page 75.

See [“Add a batch file-based policy”](#) on page 200.

You can also configure a protection plan with snapshot backups.

See [“Create a protection plan to protect SQL Server assets”](#) on page 134.

To configure a snapshot policy for SQL Server

- 1 For batch file-based policies, create a backup script (.bch file) using the NetBackup MS SQL Client.
- 2 Create a policy.
- 3 Select the **Attributes** tab.
- 4 From the **Policy type** list, select **MS-SQL-Server**.
- 5 Select the **Policy storage**.

If database differentials, filegroup differentials, or transaction logs are included in the **Backup selections list** of a policy that uses Snapshot Client, then NetBackup performs a stream-based backup. The selected storage unit is used. If a storage unit is not provided, then NetBackup uses the default storage unit for the server.

- 6 Select **Perform snapshot backups**.
- 7 Choose to have NetBackup select the snapshot method or select the snapshot method manually.

Perform one of the following:

- By default, NetBackup chooses a snapshot method for you. If you have changed this setting and want NetBackup to choose the method automatically, select **Snapshot options**. Then from the **Snapshot method** list, select **auto**.
 - To use a specific snapshot method, select **Snapshot options**. From the **Snapshot method** list, select the method you want to use for this policy.
- 8** (Optional) To use an alternate client to reduce the processing load on the client, perform the following steps:
- The alternate client must be the client that shares the disk array. This option may require additional configuration.
 - Select **Perform off-host backup**.
 - In the **Use** list, select **Alternate client**. Then in the **Machine** list, select the client name.

Note: **Use data mover** is not a supported option for NetBackup for SQL Server.

- 9** On the **Instances and databases** tab, choose how you want to protect SQL Server:
- (SQL Server Intelligent Policy) Choose if you want to protect instances and databases, intelligent groups, or instance groups.
 See [“Add instances to a policy”](#) on page 80.
 See [“Add databases to a policy”](#) on page 81.
 See [“Add intelligent groups to a policy”](#) on page 82.
 See [“Add instance groups to a backup policy”](#) on page 84.
 - (Batch-file based policies) Choose **Clients for use with batch files**.
- 10** (SQL Server Intelligent Policy) Add other policy information as follows:
- Add schedules.
 See [“Schedule properties for SQL Server Intelligent Policies”](#) on page 77.
 - (Optional) Select the specific filegroups or files that you want to back up. By default, NetBackup backs up an entire database.
 See [“Add filegroups or files to the backup selections list”](#) on page 83.
 - (Optional) Make changes to any tuning parameters.
 See [“Performance tuning and configuration options”](#) on page 84.
- 11** (Batch-file based policies) Add other policy information as follows:

- Add schedules.
See [“Schedule properties for SQL Server batch file-based policies”](#) on page 201.
- Add clients.
See [“Add clients to a policy”](#) on page 205.
- Add batch files to the backup selections list.
See [“Add batch files to the backup selections list”](#) on page 206.

12 Select **Create** to save the policy.

Configure a policy for Instant Recovery backups of SQL Server

Note: NetBackup does not support Instant Recovery backups of availability databases.

These instructions describe how to configure a policy for Instant Recovery. Optionally you can choose to back up to disk only. This topic only covers what is necessary to configure Instant Recovery backups for a MS-SQL-Server policy.

See [“About SQL Server Intelligent Policies”](#) on page 75.

See [“Add a batch file-based policy”](#) on page 200.

To configure a policy for Instant Recovery

- 1** For batch file-based policies, create a backup script using the NetBackup MS SQL Client interface.
- 2** Create a policy.
- 3** Select the **Attributes** tab.
- 4** From the **Policy type** list, select **MS-SQL-Server**.
- 5** Select the **Policy storage**.

If you select an Instant Recovery option on the **Schedules** tab (see step **10**), the storage unit is not used. NetBackup creates only a disk snapshot.

If database differentials, filegroup differentials, or transaction logs are included in the policy, then NetBackup performs a stream-based backup. This backup uses the selected storage unit. If a storage unit is not provided, then NetBackup uses the default storage unit for the server.

- 6** Select **Perform snapshot backups**.

- 7 Choose to have NetBackup select the snapshot method or select the snapshot method manually.
 - By default, NetBackup chooses a snapshot method for you.
 - To use a specific snapshot method, select **Snapshot options**. Then select the method from the **Snapshot method** list.

- 8 Select **Retain snapshots for Instant Recovery**.

NetBackup retains the snapshot on disk, so that Instant Recovery can be performed from the snapshot.

A normal backup to storage is also performed, if you do not choose to create a snapshot only (see step 10).

- 9 On the **Instances and databases** tab, choose how you want to protect SQL Server:
 - (SQL Server Intelligent Policy) Choose if you want to protect instances and databases, intelligent groups, or instance groups.
See [“Add instances to a policy”](#) on page 80.
See [“Add databases to a policy”](#) on page 81.
See [“Add intelligent groups to a policy”](#) on page 82.
See [“Add instance groups to a backup policy”](#) on page 84.
 - (Batch-file based policies) Choose **Clients for use with batch files**.

- 10 To configure schedules, select the **Schedules** tab.

- (SQL Server Intelligent Policies) Configure a full backup schedule.
See [“Schedule properties for SQL Server Intelligent Policies”](#) on page 77.
- (Legacy policies) Follow the instructions to configure an Application and a full backup schedule.
See [“Schedule properties for SQL Server batch file-based policies”](#) on page 201.

For snapshot backup policies, a full backup schedule must exist for NetBackup to successfully convert differential backups to full backups.

- 11 (Optional) To create a disk image only, open the Full backup schedule (Intelligent Policies) or the Application schedule (batch file-based policies) and select an Instant Recovery option.

Select one of the following options:

- If **Snapshots and copy snapshots to a storage unit** is selected, NetBackup creates a disk snapshot. NetBackup also backs up the client's data to the storage unit that is specified for the policy.


- If **Snapshots only** is selected, the image is not backed up to tape or to other storage. NetBackup creates a disk snapshot only. Note that this disk snapshot is not considered a replacement for traditional backup.
- 12** (SQL Server Intelligent Policy) Add other policy information as follows:
- (Optional) Select the specific filegroups or files that you want to back up. By default, NetBackup backs up an entire database.
 See [“Add filegroups or files to the backup selections list”](#) on page 83.
 - (Optional) Make changes to any tuning parameters.
 See [“Performance tuning and configuration options ”](#) on page 84.
- 13** (Batch-file-based policies) Add other policy information as follows:
- Add clients.
 See [“Add clients to a policy”](#) on page 205.
 - Add batch files to the backup selections list.
 See [“Add batch files to the backup selections list ”](#) on page 206.
- 14** Select **Create** to save the policy.

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 9-1 Recovering after an error when using full and copy-only backups



The copy-only attribute appears in the properties for the snapshot backup image. Differential backups are automatically associated with the correct full backup. The SQL Agent recognizes these backups when it selects the recovery set for the full database restore.

Create a copy-only backup

Any backup can be created as copy-only. An Instant Recovery backup is automatically created as copy-only. For SQL Server Intelligent Policies, enable **Copy-only backup** on the **Microsoft SQL Server** tab. For batch file-based policies, set the `COPYONLY TRUE` setting in the backup batch file.

See [“Performance tuning and configuration options”](#) on page 84.

Create a copy-only backup (SQL Server Intelligent policy)

To create a copy-only backup (SQL Server Intelligent policy)

- 1 Follow the instructions for creating a snapshot policy
 See [“Configure a snapshot policy for SQL Server”](#) on page 123.
- 2 On the **Microsoft SQL Server** tab, enable **Copy-only backup**.

Create a copy-only backup (batch file-based policy)

To create a copy-only backup (batch file-based policy)

- 1 Open an existing batch file in a text editor.
- 2 Insert the following:


```
COPYONLY TRUE
```
- 3 Save the batch file.

Creating an Instant Recovery backup that is not copy-only (batch file-based policies)

For Instant Recovery backups, NetBackup automatically creates the backup image as copy-only. You can choose *not* to create the backup as copy-only.

To create an Instant Recovery backup that is not copy-only

- 1 Open an existing batch file in a text editor.
- 2 Insert the following:


```
COPYONLY FALSE
```
- 3 Save the batch file.

About SQL Server agent grouped snapshots

Note: This feature is only available with SQL Server intelligent policies and batch file-based policies. It is not available with protection plans.

Grouped snapshots quiesce a group of databases together and snapshot them at the same time to back them up as a group. NetBackup automatically discovers and groups the databases, up to specified group size. The constituent files of all databases are backed up to a single storage image under the same backup ID. This means that an "import and copy" procedure would use only one image to export all of the database backups in the group.

Requirements for a grouped snapshot

Certain requirements must be met for a grouped snapshot to be performed. If any of the following requirements are not met, a standard backup is performed:

- All backup operations must be full backups. Differential backups and transaction log backups are not supported.
- (Batch-file based policies) The same policy must be specified for each backup operation in the group.
- (Batch-file based policies) The same NetBackup server must be specified for each backup operation in the group.
- The group size is limited to 64. NetBackup automatically creates additional snapshots if there are more than 64 databases in an instance.

Restoring a database backed up in a group

A database that is backed up in a group can be restored like any other database.

See [“Perform a complete database recovery ”](#) on page 150.

See [“Recover a single recovery point ”](#) on page 153.

See [“Restore a SQL Server availability database to a secondary replica”](#) on page 163.

See [“Restore a SQL Server availability database to the primary and the secondary replicas”](#) on page 165.

See [“Restore a SQL Server database backup \(NetBackup MS SQL Client\)”](#) on page 227.

Protecting SQL Server in a cluster environment

This chapter includes the following topics:

- [Configure backups of clustered SQL Server instances \(SQL Server Intelligent Policy\)](#)
- [Configure backups of clustered SQL Server instances \(batch file-based policies\)](#)

Configure backups of clustered SQL Server instances (SQL Server Intelligent Policy)

This procedure describes how to protect SQL Server instances that are clustered with a SQL Server Intelligent Policy. Perform these steps after you create the virtual SQL Server (VIRTUALSERVER). The following actions must be performed on the primary server or on a NetBackup remote client console that acts for the primary server.

If you have a SQL Server cluster with multiple NICs, you must follow a different procedure.

See [“Configure backups of a SQL Server cluster when you have multiple NICs \(SQL Server Intelligent Policies\)”](#) on page 250.

To configure backups of clustered SQL Server instances (SQL Server Intelligent Policy)

- 1 Create a policy (for example, VIRTSQLPOLICY).
- 2 On the **Attributes** tab, configure the following:
 - Select the **MS-SQL-Server** policy type.

Configure backups of clustered SQL Server instances (SQL Server Intelligent Policy)

- Specify a storage unit. If you use a virtual media server, then specify a storage unit that belongs to the virtual media server.
- 3** On the **Instances and databases** tab, choose if you want to protect instances and databases, intelligent groups, or instance groups.
- 4** Add the assets that you want to protect.
 - See [“Add instances to a policy”](#) on page 80.
 - See [“Add databases to a policy”](#) on page 81.
 - See [“Add intelligent groups to a policy”](#) on page 82.
 - See [“Add instance groups to a backup policy”](#) on page 84.

For a clustered instance, the host name is the virtual name of the SQL Server cluster.
- 5** Add other policy information as follows:
 - Add schedules.
 - See [“Schedule properties for SQL Server Intelligent Policies”](#) on page 77.
 - (Optional) Select the specific filegroups or files that you want to back up. By default, NetBackup backs up an entire database. Filegroup and file backups are not available for any policies that use intelligent groups.
 - See [“Add filegroups or files to the backup selections list”](#) on page 83.
 - (Optional) Make changes to any tuning parameters.
 - See [“Performance tuning and configuration options ”](#) on page 84.
- 6** Map the virtual name of the SQL Server cluster to each node in the cluster.

Configure these mappings in the **Distributed application restore mapping** host property on the primary server.

See [“Configuring mappings for restores of a distributed applications, clusters, or virtual machines ”](#) on page 31.
- 7** Configure the mapped host names for the SQL Server hosts in your environment.

Configure these mappings in **Security > Host mappings** on the primary server.

See [“Reviewing the auto-discovered mappings”](#) on page 26.

Configure backups of clustered SQL Server instances (batch file-based policies)

This procedure describes how to protect SQL Server clustered instances with a batch file-based policy. Perform these steps after you create the virtual SQL Server (VIRTUALSERVER). The following actions must be performed on the primary server or on a NetBackup remote client console that acts for the primary server.

If you have a SQL Server cluster with multiple NICs, you must follow a different procedure.

See [“Configure backups of a SQL Server cluster when you have multiple NICs \(batch file-based policies\)”](#) on page 251.

To configure backups of clustered SQL Server instances

- 1 Create a policy (for example, VIRTSQLPOLICY).
- 2 On the **Attributes** tab, configure the following:
 - Select the **MS-SQL-Server** policy type.
 - Specify a storage unit. If you use a virtual media server, then specify a storage unit that belongs to the virtual media server.
- 3 On the **Instances and databases** tab, select **Clients for use with batch files**.
- 4 On the **Schedules** tab, add an automatic backup schedule.
- 5 On the **Clients** tab, add the virtual SQL Server name (VIRTUALSERVER).
- 6 On the **Backup selections** tab, add one or more script names (batch files).
- 7 Map the virtual name of the SQL Server cluster to each node in the cluster.

Configure these mappings in the **Distributed application restore mapping** host property on the primary server.

See [“Configuring mappings for restores of a distributed applications, clusters, or virtual machines”](#) on page 31.

- 8 Configure the mapped host names for the SQL Server hosts in your environment.

Configure these mappings in **Security > Host mappings** on the primary server.

See [“Reviewing the auto-discovered mappings”](#) on page 26.

Managing protection plans for SQL Server

This chapter includes the following topics:

- [Create a protection plan to protect SQL Server assets](#)
- [Add SQL Server assets to a protection plan](#)
- [Customize protection settings for a NetBackup for SQL Server asset](#)
- [Remove protection from SQL Server assets](#)
- [Create a protection plan to protect SQL Server data with a VMware backup](#)
- [Protect a SQL Server availability group that crosses NetBackup domains with a protection plan](#)

Create a protection plan to protect SQL Server assets

You can create a protection plan to perform scheduled backups of SQL Server assets. First create the protection plan for SQL Server. Then select this protection plan for your 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.


- 4 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 121. 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. A tape, storage unit groups, and Replication Director are 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 a unique storage for transaction logs.

- 5 In **Backup options**, configure the options that you want.
- See ["Performance tuning and configuration options "](#) on page 84.
- 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**.
- 8 You can now select the protection plan for your SQL Server assets.
 - On the left, click **Workloads > Microsoft SQL Server**.
 - Select the instances, availability groups, or databases that you want to protect.
 - Click **Add protection**.
 - Select the protection plan that you created.

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 11-1 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.

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 must have a valid credential configured for it in **Instances** tab.
 See [“Register a SQL Server instance with an existing credential”](#) on page 48.
 See [“Register a SQL Server instance with a new credential”](#) on page 49.
 See [“Add a credential rule”](#) on page 52.
- 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 21.
- 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.
- Refer to the additional information and limitations for availability groups:
 See [“About protecting SQL Server availability groups”](#) on page 93.

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 | <ul style="list-style-type: none"> ■ 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 | <ul style="list-style-type: none"> ■ 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) | <p>On the Instances tab, select the instance name depending on if you want to protect the cluster or a node in the cluster:</p> <ul style="list-style-type: none"> ■ The instance name, where the Host name is the instance 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 | <p>On the Instances tab, select the instance:</p> <ul style="list-style-type: none"> ■ 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 121.
- 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 137.
- **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 84.

6 Click **Protect.**

The results of your choices appear under **Instances** or **Databases**.

Customize protection settings for a NetBackup for SQL Server asset

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

- See [“Schedules”](#) on page 137.
- See [“Performance tuning and configuration options”](#) on page 84.

To customize protection settings for a NetBackup for SQL Server asset

1 On the left, click **Workloads > NetBackup for 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 137.
- For transaction log schedules, you can edit the frequency and the retention.
See [“Schedules”](#) on page 137.
- **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 84.

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 > NetBackup for 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**.

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

Through a VMware backup policy, NetBackup can create full application-consistent backups of the SQL Server databases that reside on a virtual machine. Optionally you can use NetBackup Accelerator. VMware policies let you exclude certain virtual disks from the VMware backup. If you want to exclude specific SQL Server components, use a MS-SQL-Server policy.

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 [“Configure the NetBackup services for SQL Server backups and restores”](#) on page 24.
 See [“Configure local security privileges for SQL Server”](#) on page 25.
- 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 137.

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

Storage option	Requirements	Description
Backup storage	OpenStorage is required for this option. Tape, storage unit groups, and Replication Director are 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 144.
- 9 Under **Allow restore of application data from virtual machine backups**, select **Microsoft SQL Server**.

Do not enable **Truncate logs** at this time. You must first perform a full backup without log truncation, described later in this procedure.
- 10 (Conditional) Select **Enable T-SQL snapshots**.

This snapshot type creates a full SQL Server backup (not copy-only) that can be used as a basis for SQL Server incremental and transaction log backups. For this snapshot type, the only valid option for snapshot handling is **Stop the backup if any snapshots exist**, which is automatically selected.

Note: T-SQL snapshots were added with SQL Server 2022 and are strongly recommended for systems with SQL Server 2022 or later. This type of snapshot is not supported with SQL Server 2019 and earlier.
- 11 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](#).
- 12 In **Review**, verify that the protection plan details are correct and click **Save**.
- 13 If you do not want to truncate transaction logs, no further action is necessary.

If you want to truncate transaction logs, continue with step [14](#).
- 14 Perform a full backup without log truncation.

When the backup completes, open the policy that you created in step [2](#).
- 15 Click the **VMware** tab.

- 16** Locate **Allow restore of application data from virtual machine backups** and click **Microsoft SQL Server**. Then click **Truncate logs**.

For SQL Server, this option truncates the transaction logs when the VMware snapshot of the virtual machine is complete.

- 17** Click **Save** to save the protection plan.
- 18** Perform a full VMware backup.

Backup options and Advanced options

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

Backup options

Table 11-2 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.
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 145.

Advanced options

Table 11-3 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.

Table 11-3 Advanced options for protection plans (*continued*)

Option	Description
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. If this option is disabled, you can recover the entire virtual machine from the backup, but you cannot recover the databases individually. See “Create a protection plan to protect SQL Server data with a VMware backup” on page 142. See “Configuring a VMware policy to protect SQL Server using Replication Director to manage snapshot replication” on page 115.
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 146.

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 11-4 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.

Table 11-4 Options for excluding virtual disks (*continued*)

Exclude option	Description
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 11-5 Snapshot retry options

Option	Description
Maximum number of times to retry a snapshot	The number of times the snapshot is retried.
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 a SQL Server availability group that crosses NetBackup domains with a protection plan

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, select **Protection > Protection plans** and then select **Add**.
- 2 In **Basic properties**, enter a **Name** and **Description**.
- 3 From the **Workload** list, select **Microsoft SQL Server**.
- 4 In **Schedules and retention**, select **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 the schedule that you configured in step 4.

Storage option	Requirements	Description
Backup storage	OpenStorage is required for this option. Tape, storage unit groups, and Replication Director not supported.	Select Edit to select the storage target. Select Use selected storage after selecting the storage target.
Replication target	The backup storage must be a source in a targeted A.I.R. environment.	Select Edit to select the replication target primary server. Select a primary server and then select a storage lifecycle policy. Select 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 84.
- (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 select **Finish**.

Additional resources

- [NetBackup Web UI Administrator's Guide](#)
- [NetBackup Deduplication Guide](#)
- [NetBackup OpenStorage Solutions Guide](#)
- <http://www.netbackup.com/compatibility>

Restoring SQL Server with the NetBackup web UI

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](#)
- [Configuring permissions for redirected restores](#)
- [Restore SQL Server databases from a VMware backup](#)
- [Restore a SQL Server availability database to a secondary replica](#)
- [Restore a SQL Server availability database to the primary and the secondary replicas](#)
- [Restoring an availability database when an availability group crosses NetBackup domains](#)

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 21.

- 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 157.
- 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.
- The security administrator has configured the necessary mappings for the hosts in **Security > Host mappings**. 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 instance 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 private interface name of the virtual SQL Server

- 3 Select **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, select **Copies**.

See [“Select a different backup copy for recovery”](#) on page 158.
- 5 Select **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.
 - Select **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,

Select **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. Select Edit file paths and select any directory path to edit the restore path for that file.

Example of a restore to different paths:

Edit file paths

Search...

Restore default settings

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

Cancel

Save

- 9 Enter the credentials for the recovery target. Or, select **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 Select **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 156.
- 12 Select **Next**.
- 13 On the **Review** page, review the restore options that you selected.
- At the top, select the link that follows **Recovery set** to view the backup images that are required for the restore.

- Select **Edit** to change the **Recovery target** settings or **Recovery options**.
- Select **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 instance 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 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 158.

- 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 156.
- 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 12-1 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 12-1 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 a consistency check. ■ 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	This option 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. Disconnect users Disconnects any active users before the restore operation runs.
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 [“Register a SQL Server instance with an existing credential”](#) on page 48.
See [“Register a SQL Server instance with a new credential”](#) on page 49.
See [“Add a credential rule”](#) on page 52.
 - Perform a new backup of the database.
Locate the database and click **Actions > 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 ▼ ⋮
Storage	Storage server	Storage type
storage1 (Primary copy)	storageserver1	MSDP ⋮
storage2	storageserver1	AdvancedDisk ⋮
E:\storage3	storageserver1	Perform complete database recovery
/storage4	storageserver2	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 12-1 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 12-2). Otherwise, you can click **Edit** to select the specific storage that you want to use (see Figure 12-3).

Figure 12-2 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 12-3 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

Configuring permissions for redirected restores

Certain restore procedures or environments require that you configure permissions to allow a client to restore a backup that another client performed. See the [NetBackup Web UI Administrator's Guide](#) for complete details on redirected restores.

You must configure the primary server for redirected restores if you want to redirect the restore of *ClientA* to *ClientB*.

If you use a non-root service user account, specific access must be allowed for that user when you add files to the `/usr/opensv/netbackup/db/altnames` directory. The service user account must have full access to these files through the ownership or group and the permissions. For example, if the service user is `svcname` and its group is `svrgrp`, the file can have permissions of `400`. If the file owner is for a different user and group, the file permissions must allow access to the service user. For example, `777`. Equivalent permission settings must be used in a Windows environment.

You do not need to configure redirected restores for the following configurations:

- Restore databases in a SQL Server cluster to any of the nodes in the cluster

- Restore databases in an availability group to any of the nodes in the availability group
- Restore clustered databases in a multi-NIC environment across the private interface

Instead these environments require that you review the auto-discovered mappings for the hosts in your environment.

See [“Reviewing the auto-discovered mappings”](#) on page 26.

See [“Configuring mappings for restores of a distributed applications, clusters, or virtual machines”](#) on page 31.

To allow a specific client or host to perform a redirected restore

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

For example, to give `HostB` permissions to redirect a restore, create the following file:

On Windows:

```
install_path\NetBackup\db\altnames\HostB
```

On UNIX:

```
/usr/opensv/netbackup/db/altnames/HostB
```

- 2 In the `altnames` file, add the names of the hosts whose files the requesting client wants to restore.

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

To give a SQL Server host the permissions to restore backups in a multi-NIC environment

- 1 Create an `altnames` file with the private name of the host, for example `SQLHOST1-NB`.

On Windows:

```
install_path\NetBackup\db\altnames\SQLHOST1-NB
```

On UNIX:

```
/usr/openv/netbackup/db/altnames/SQLHOST1-NB
```

- 2 In the `altnames` file, add the names of the hosts whose files the requesting client wants to restore.

For example, assume that you want `SQLHOST1-NB` to have permissions to redirect restores from `SQLHOST2-NB`. Then add `SQLHOST2-NB` to the `SQLHOST1-NB` file.

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 Select the **Databases** tab.
- 2 Select the database that you want to restore and select **Recover**.
- 3 Select the date that the backup was performed.
- 4 On the right, locate the recovery point. Then select **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.

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

- 1 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.
- 2 On the left, select **Workloads > Microsoft SQL Server**.
- 3 Select 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 secondary replica.
- 5 On the **Databases** tab, click on the database that you want to restore.
- 6 Select the **Recovery points** tab and locate the latest transaction log backup. By default NetBackup uses the primary copy. To select a different copy, select **Copies**.
 See [“Select a different backup copy for recovery”](#) on page 158.
- 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.

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

- 10 Enter the credentials of the instance that you want to restore to and select **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**
- (Optional) Select **Disconnect users**. This option disconnects any active users of the database.

See [“Options for SQL Server restores”](#) on page 156.

- 12 Select **Next**. Then select **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 158.
- 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**
 - (Optional) Select **Disconnect users**. This option disconnects any active users of the database.

See [“Options for SQL Server restores”](#) on page 156.

- 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**
 - (Optional) Select **Disconnect users**. This option disconnects any active users of the database.

See [“Options for SQL Server restores”](#) on page 156.
- 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.

Restoring an availability database when an availability group crosses NetBackup domains

To restore an availability group database that was backed up by an availability group node in another NetBackup domain, you must first configure NetBackup for Auto Image Replication (A.I.R.). The backup must complete and be replicated to the target replicas. Once the backup is replicated, you can perform a restore on a target replica in the same way as you perform any other restore of availability group databases.

Note: Replication may not occur immediately to the target availability group replicas. The time it takes for replication to occur is dependent on the settings for each primary server.

See [“Protect a SQL Server availability group that crosses NetBackup domains with a protection plan”](#) on page 147.

See [“Restore a SQL Server availability database to a secondary replica”](#) on page 163.

See [“Restore a SQL Server availability database to the primary and the secondary replicas”](#) on page 165.

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).
 - If SE Linux is configured, ensure that the `polycoreutils` and `polycoreutils-python` (for RHEL 7) or `polycoreutils-python-utils` (for RHEL 8) 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.
 - 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 13-1 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://support.cohesity.com/s/article/article-100051793>

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

A NetBackup client may need Samba user credentials to access Samba shares. You can configure Samba local users for SQL Server instant access on the corresponding storage server.

If the Samba service on a storage server is part of Windows domain, the Windows domain users can be used as Samba users.

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.

During SQL Server instant access, the SQL Server service needs to access the Samba share. If a Windows user is specified to start the instant access database, that Windows user also needs to access the Samba share.

How to make Samba shares available to a Windows user

- If the Windows user is a domain user and is in the same domain as the storage server:
The Windows user can directly access the Samba share and no configuration is required.
- If the Windows user is not a domain user, or is not in the same domain as the storage server:
Save the Samba user credentials for the Windows user by running the following command and enter the Samba account password:
`cmdkey /add:<Samba hostname> /user:<Samba account username> /pass`
The Windows user accesses the Samba share using the credentials.

How to make Samba shares available to a SQL Server service

- If the SQL Server service logs on as a Windows user, refer to the following topic:
See [the section called “How to make Samba shares available to a Windows user”](#) on page 172.
- If the SQL Server service logs on as a service account (for example, NT Service\MSSQLSERVER)
 - If the SQL Server Windows host is in the same domain as the storage server: SQL Server service is authenticated as the domain host and no configuration is required.
 - If the SQL Server Windows host is not in any domain and Samba guest access is not disabled, the SQL Server service can access the share as guest and no configuration is required.
 - For all other scenarios, create a Samba session for the SQL Server service account by running following SQL statement:

```
xp_cmdshell 'net use \\<Samba hostname>\<sharename> <Samba  
account password> /user:<Samba account username>'
```

The SQL Server service accesses the Samba share using the provided Samba user credentials. The share name must be a share that is available on the storage server. If there is no share at the time, you must create one. The Samba session is valid until the next restart. You must run the command again after restart to get Samba access.

If `xp_cmdshell` is not enabled for the SQL Server, use the following commands to enable or disable `xp_cmdshell`.

```
-- enable xp_cmdshell
EXEC sp_configure 'show advanced options', '1'
RECONFIGURE
EXEC sp_configure 'xp_cmdshell', '1'
RECONFIGURE

-- disable xp_cmdshell
EXEC sp_configure 'show advanced options', '1'
RECONFIGURE
EXEC sp_configure 'xp_cmdshell', '0'
RECONFIGURE
```

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

Table 13-2 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.
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: ■ #useradd <username> ■ #passwd <username> You can also use an existing local user. ◆ Run the following commands to create user credentials for Samba and enable the user: ■ smbpasswd -a <username> ■ smbpasswd -e <username>
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 # 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 13-2 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 **Workloads > 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** Locate 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 179.
- 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 177.

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 **Workloads > 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 Locate 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 177.

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 **Workloads > 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 **Workloads > Microsoft SQL Server**.
- 2 Click the **Instant access databases** tab.
The tab lists the names of the configured instant access databases.
- 3 Select the instant access database.
- 4 Click **Delete**.
- 5 Based on one of the following scenarios, perform the steps that apply:

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 13-3 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.

Table 13-3 Recovery options (*continued*)

Option	Description
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 175.

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 13-4 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 13-5

Applicable for	Frequently asked question	Answer
BYO	How can I enable the Microsoft SQL Server instant access feature on BYO after storage is configured or upgraded without the <code>nginx</code> service installed?	Perform the steps in the following order: 1 Install the required <code>nginx</code> service version. 2 Ensure that the new BYO <code>nginx</code> 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 <code>https</code> on port 10087	Perform the steps in the following order: 1 Install the <code>polycoreutils</code> and <code>polycoreutils-python</code> packages through the 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 13-5 (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 live mount share in gnome? If the automount is enabled, the source folder is mounted from the live mount 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 live mount 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 13-5 (continued)

Applicable for	Frequently asked question	Answer
BYO	How can I resolve the following issue in the <code>/var/log/vpfs/vpfs-config.log</code> 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: <code>/usr/opensv/netbackup/bin/bpclntcmd -pn</code> 3 After the NetBackup primary server is up and connection between the NetBackup primary server and storage server is allowed, run the following command: <code>/usr/opensv/pdde/vpfs/bin/vpfs_config.sh --configure_byo</code>

Table 13-5 (continued)

Applicable for	Frequently asked question	Answer
BYO and Flex Appliance	How can I enable host-based authentication and secure logon for Samba share so that the SQL Server instant access works on specific windows clients?	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: <code>map to guest = Never</code> ■ 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. 3 To save the Samba credentials on a Windows client, open go to Control Panel > User Accounts > Credential Manager > Add a Windows Credential. 4 In Internet or network address, enter the storage server domain name. 5 Enter the Samba username and password. Ensure that the username is same as the user credentials that you created for Samba. 6 Click OK and ensure that you can access <storage server domain name> without a logon prompt.

Table 13-5 (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. For Appliance, the <code>smb.conf</code> file configuration already contains <code>map to guest = Never</code>. The local users are added to Samba database automatically and the Samba password is the same as the logon password. The Windows clients can access the appliance's Samba share using credentials of the appliance's local users. Perform the steps in the following order: 1 To manage appliance local users, go to the following CLISH path: Main_Menu > Settings > Security > Authentication > LocalUser 2 To save the Samba credentials on a Windows client, open go to Control Panel > User Accounts > Credential Manager > Add a Windows Credential. 3 In Internet or network address, enter the storage server domain name. 4 Enter the Samba username and password. 5 Click OK and ensure that you can access <storage server domain name> without a logon prompt.

Configuring batch-file based policies for SQL Server backups

This chapter includes the following topics:

- [About batch file-based policies for SQL Server backups](#)
- [Overview of configuring SQL Server backups with batch file-based policies](#)
- [Configure the NetBackup services for SQL Server backups and restores \(batch file-based policies\)](#)
- [About SQL Server security with batch file-based policies](#)
- [Requirements to use batch files with NetBackup for SQL Server](#)
- [Add a batch file-based policy](#)
- [Schedule properties for SQL Server batch file-based policies](#)
- [Add clients to a policy](#)
- [Add batch files to the backup selections list](#)
- [Options for SQL Server backup operations](#)
- [Create a script to backup a remote SQL Server installation](#)
- [About automatic retry of unsuccessful SQL Server backups](#)
- [Configure a batch file-based policy for a user-directed backup of read-only filegroups](#)

- [Configure a batch file-based policy for a user-directed backup of read-write filegroups](#)

About batch file-based policies for SQL Server backups

A batch file-based policy includes a list of SQL Server database clients and a batch file that contains SQL Server backup commands. When a backup is scheduled, NetBackup runs the commands in the batch file for each client in the policy. You create the batch file through the NetBackup MS SQL Client interface, which saves the options you select to a batch file. Or you can create this batch file manually.

The batch file-based policy includes the following criteria:

- Storage unit and media to use
- Policy attributes
- Backup schedules: Automatic schedule (called Full Backup) and application schedule
- Clients to be backed up
- Backup batch files to be run on the clients

Overview of configuring SQL Server backups with batch file-based policies

Table 14-1 Steps to configure SQL Server backups that use batch file-based policies

Step	Action	Description
Step 1	Configure the logon account for the NetBackup services.	The logon account for the NetBackup Client Service and the NetBackup Legacy Network Service must meet certain requirements. See “Configure the NetBackup services for SQL Server backups and restores” on page 24.
Step 2	Configure the batch files for the policy.	See “Requirements to use batch files with NetBackup for SQL Server” on page 190.
Step 3	Configure a batch file-based policy.	See “Add a batch file-based policy” on page 200.

Table 14-1 Steps to configure SQL Server backups that use batch file-based policies *(continued)*

Step	Action	Description
Step 4	If you have a SQL Server availability group or cluster, you must configure the mappings for distributed application restores.	See “Configuring mappings for restores of a distributed applications, clusters, or virtual machines” on page 31.
Step 5	If you have a SQL Server availability group or cluster, you must review the auto-discovered mappings for the hosts in your environment.	See “Reviewing the auto-discovered mappings” on page 26.

Configure the NetBackup services for SQL Server backups and restores (batch file-based policies)

NetBackup uses the NetBackup Client Service and the NetBackup Legacy Network Service to access the SQL Server when it performs backups and restores. With the proper configuration, these services can log on with the Local System account or another account that has the necessary privileges.

The logon account for the services requires the following:

- Both services must use the same logon account.
- The SQL Server “sysadmin” role.
- Apply the sysadmin role manually to the NT AUTHORITY\SYSTEM or the BUILTIN\Administrators group.
- For a SQL Server cluster or SQL Server availability group, configure the NetBackup services on each node in the cluster or availability group.
- For VMware backups, different configuration is required for logon account for the services.

See [“Configure the NetBackup services for SQL Server backups and restores”](#) on page 24.

To configure the NetBackup services for SQL Server backups and restores

- 1 Log on to the Windows host with the account that has the sysadmin role.
- 2 If the SQL Server instance uses standard or mixed security, perform the following steps:
 - Open the NetBackup MS SQL Client.
 - Select **File > Set SQL Server connection properties**.
 - Provide the SQL Server **Userid** and **Password**.

- Click **Apply** > **Close**.
- 3 In the Windows Services application, open the **NetBackup Client Service** entry and click the **Log On** tab.
 - 4 Confirm that **Local System account** is selected.
If you selected a different logon account, stop and restart the service.
 - 5 Open the **NetBackup Legacy Network Service** entry and click the **Log On** tab.
 - 6 Confirm that **Local System account**.
If you selected a different logon account, stop and restart the service.

About SQL Server security with batch file-based policies

NetBackup for SQL Server uses SQL Server backup and restore commands and queries the SQL Server master database. These operations are validated according to the security method you choose when you install SQL Server, either integrated security or standard security. Integrated security refers to the use of Windows authentication in lieu of standard SQL Server-based logons.

Note: Microsoft recommends using integrated security. Unlike SQL Server-based logons, Windows logons can be traced with standard Windows security tools. NetBackup for SQL Server supports both integrated security and standard security for any level of SQL Server.

If you use integrated security, the Windows account you log into is used for authentication. SQL Server ignores any user ID and password that you enter in the NetBackup MS SQL Client or in a batch file.

If you use standard security, then you must supply a SQL Server-based user ID and password. Once you provide these credentials, NetBackup stores this information in the registry (the password is encrypted) under the following registry key:

```
HKEY_CURRENT_USER\SOFTWARE\VERITAS\NETBACKUP\NetBackup for  
Microsoft SQL Server\
```

Requirements to use batch files with NetBackup for SQL Server

NetBackup for SQL Server uses batch files to initiate backup and restore operations. A batch file uses the `.bch` extension and is typically run from the `install_path\DbExt\MsSql\` directory. If you use a SQL Server intelligent policy or a protection plan for backups, the batch files are created automatically. Batch files are automatically created for restores that you perform with the NetBackup web UI.

You must create a batch file if you start operations in any of the following ways:

- NetBackup MS SQL Client
- `dbbackup` command line
- Automatically scheduled backups from policies that use batch files and clients

Rules for using batch files

Review the following information before you create and use batch files:

- Ensure that the batch file resides on the client.
See [“Registering authorized locations used by a NetBackup database script-based policy”](#) on page 290.
- Batch files are in Unicode text.
- A batch file consists of a series of operations that run in sequence. For batch file-based policies, you create batch files for backup operations and restore operations. For SQL Server Intelligent Policy, you create the batch files for restore operations in the same way.
- Each operation consists of a series of `<keyword value>` pairs, which completely define the total operation.
- The keyword is not case-sensitive but the value is. Generally, you can code both the keyword and value in uppercase. The exception is the `NBIMAGE` keyword option. The value must be specified exactly as it appears in the NetBackup server.
- Operations are not nested.
- With the exception of the `BATCHSIZE`, `GROUPSIZE`, `RESTARTTYPE`, `NUMRESTARTS`, and `RESTARTWAITSECONDS` parameters, `<keyword value>` pairs are not global. If you use `BATCHSIZE`, `GROUPSIZE`, `RESTARTTYPE`, `NUMRESTARTS`, or `RESTARTWAITSECONDS` then it must appear only once in your batch file and it must appear in the first operation.

- If `SQLINSTANCE $ALL` is used, then it must appear in the first operation of the batch file. Each operation in the batch file is performed for all SQL Server instances on the client where the batch file is run. Also, it is not necessary to specify an `SQLHOST` or `SQLINSTANCE` on any subsequent operations.
- Within an operation, the *<keyword value>* pairs may appear in any order except that you must terminate each operation with `ENDOPER TRUE`.
- You can include comment lines in your batch file by placing a hash mark (`#`) in the first column.
- `STOPAT`, `RESTORETOMARK`, `RESTORETOMARKAFTERTIME`, `RESTOREBEFOREMARK`, and `RESTOREBEFOREMARKAFTERTIME` are mutually exclusive restore parameters. If either `RESTORETOMARKAFTERTIME` or `RESTOREBEFOREMARKAFTERTIME` are used, then the batch file must also specify a datetime string with the keyword `STOPAFTER`.
- If you remove the `MAXTRANSFERSIZE` keyword from the batch file, the default is 0 or a maximum transfer size of 64 KB. If you remove the `BLOCKSIZE` keyword from the batch file, the default is 0 or a block size of .5 KB. A default value of 0 is also applied if you manually create a batch file without these keywords.

Keywords and values used in batch files

See [“Create a batch file”](#) on page 199.

See [“Requirements to use batch files with NetBackup for SQL Server”](#) on page 190.

[Table 14-2](#) describes the keywords and values that can be used in batch files.

Table 14-2 Keywords and values used in batch files

Keyword and description	Type and values
<code>ALTCLIENT</code> (Same as <code>BROWSECLIENT</code>) - Restores the images from a host other than the local host.	String Default: None Required: No
<code>BACKUPMODEL</code> - Valid only for restore. Indicates whether the backup was originated from a snapshot method.	<code>BACKUPMODEL_</code> <code>CONVENTIONAL</code> , <code>BACKUPMODEL_ SNAPSHOT</code> Default: <code>BACKUPMODEL_</code> <code>CONVENTIONAL</code> Required: No

Table 14-2 Keywords and values used in batch files (*continued*)

Keyword and description	Type and values
BATCHSIZE - Number of backup operations to start simultaneously, per database instance. Applies to all of the operations in the batch file. Must appear before the end of the first operation. Range is 1–32.	Integer Default: 1 Required: No
BLOCKSIZE - Applicable for backup operations only. Block size is calculated as 512 bytes * 2 ^{BLOCKSIZE} . Range is 0–7.	Integer Default: 0 Required: No
BROWSECLIENT (Same as ALTCLIENT) - Restores the images from a host other than the local host.	String Default: None Required: No
CONSISTENCYCHECK - Performs the specified consistency check after the restore has been completed.	FULLINCLUDINGINDICES, FULLEXCLUDINGINDICES, PHYSICALCHECKONLY, CHECKCATALOG Default: None Required: No
CONVERTBACKUP - If no previous full backup exists for the database or filegroup, then NetBackup converts the differential or log 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. See “Converting differential backups to full backups” on page 89. See “Converting log backups to full backups” on page 89.	TRUE, FALSE Default: FALSE Required: No
COPYONLY - If TRUE, SQL Server creates an out-of-band backup so that it does not interfere with the normal backup sequence. The default value is FALSE except for full database Instant Recovery backups. See “Using copy-only snapshot backups to affect how differentials are based” on page 127.	TRUE, FALSE Default: See description Required: No
DATABASE - Name of database. For backup operations, specify value \$ALL to designate all databases (except for tempdb.)	String Default: None Required: Yes

Table 14-2 Keywords and values used in batch files (*continued*)

Keyword and description	Type and values
DBMS - You can specify MSSQL only.	MSSQL Default: MSSQL Required: No
DUMPOPTION - Specifies INCREMENTAL restoring from an incremental backup.	INCREMENTAL Default: None Required: No
ENABLESERVICEBROKER - Enables SQL Server Service Broker after a restore operation. To take effect, RECOVERED STATE must be set to RECOVERED. Include this keyword in each individual RESTORE operation.	TRUE Default: None Required: No
ENDOPER - Terminates each operation that is specified in the batch file.	TRUE Default: None Required: Yes
EXCLUDE - Name of a database to exclude when DATABASE \$ALL is specified in a batch operation. EXCLUDE can be used in a batch file only if DATABASE \$ALL is used.	String Default: None Required: No
GROUPSIZE - The number of databases that are snapped as a single SQL Server backup image. Range is 1-64.	Integer Default: None Required: No
INHIBITALTBUFFER METHOD - Tells NetBackup whether to consider the candidacy of alternate buffer method.	TRUE, FALSE Default: FALSE Required: No
KEEPCDC - Preserves the change data capture settings when a database or log backup is recovered. This option is not valid with the RECOVEREDSTATE NOTRECOVERED option.	TRUE, FALSE Default: FALSE Required: No
MAXTRANSFERSIZE - Maximum transfer size is calculated as 64 KB * 2^MAXTRANSFERSIZE. Range is 0–6.	Integer Default: 0 Required: No

Table 14-2 Keywords and values used in batch files (*continued*)

Keyword and description	Type and values
MOVE - Specifies a filegroup name. Used for the MOVE restore type. For any backups that were made with a batch file-based policy, the PARTIAL restore type also applies.	Filegroup Default: None Required: No
NBIMAGE - Specifies a NetBackup image for the restore operations. See note for NBSERVER. * Required for restore operations.	String Default: None Required: Yes*
NBSCHED - If the NetBackup policy has several Application Backup Policy schedules, use NBSCHED to select amongst them.	String Default: None Required: No
NBSERVER - Specifies which primary server to use for the backup or restore operation. Note: If NBSERVER is not specified in a batch file operation, the primary server defaults to the name that is specified at HKEY_CURRENT_USER\Software\VERITAS\NetBackup\NetBackup for Microsoft SQL Server\DEFAULT_SQL_NB_MASTER_SERVER.	String Default: None Required: No
NUMBUFS - Number of buffers per stripe. Range is 1–32.	Integer Default: 1 Required: No
NUMRESTARTS - The number of times to retry a backup if RESTARTTYPE AUTO is specified. Use this keyword only once in the batch file and in the first operation of the batch file.	1-9 Default: 1 Required: No
OBJECTNAME - Specifies a file or a filegroup name for file or for filegroup backups and restores. * If OBJECTTYPE= FILE or FILEGROUP.	String Default: None Required: Yes*
OBJECTTYPE - Specifies the object you want to back up or restore, a database, transaction log, filegroup, or file.	DATABASE, TRXLOG, FILEGROUP, FILE Default: DATABASE Required: No

Table 14-2 Keywords and values used in batch files (*continued*)

Keyword and description	Type and values
OPERATION - Type of operation, either backup or restore.	BACKUP, RESTORE Default: BACKUP Required: No
PAGE - Ignored for a restore if the backup was performed with SQL Server Intelligent Policy. Specifies a page ID for a page restore operation.	Page ID Default: None Required: No
PARTIAL - Ignored for a restore if the backup was performed with SQL Server Intelligent Policy. Specifies NetBackup perform a partial backup or restore.	TRUE, FALSE Default: FALSE Required: No
PASSWORD - Password for logging into SQL Server. This keyword is ignored if you use integrated security.	String Default: null
PREFERREDDREPLICA - For each operation in the batch file, include this keyword. (All NetBackup versions) TRUE honors your SQL Server backup preferences. FALSE indicates there is no preference for the replica that is used for backup. NONE: The backup is performed on the specified instance. SKIP: Ignores any availability databases on the instance. PRIMARY and PREFERRED apply to availability replicas and to instances that have both standard databases and availability databases. PRIMARY: The primary replica is used for backup. PREFERRED: Honors your SQL Server backup preferences.	NONE, PRIMARY, PREFERRED, SKIP, TRUE, FALSE Default: PRIMARY
RECOVERED STATE - RECOVERED = The database is restored to the recovered state. NOTRECOVERED = The database remains in the loading state following the restore. STANDBY = The database is restored to the standby state. The STANDBYPATH keyword is also required. TRUE and FALSE are synonyms for RECOVERED and NOTRECOVERED.	RECOVERED, STANDBY, NOTRECOVERED, TRUE, FALSE Default: RECOVERED Required: No
RESTARTTYPE Available only for backups. Use AUTO to automatically retry backup of failed objects. Use MANUAL to create a batch file for backing up any of the objects that were not successfully backed up. Use this keyword only once in the batch file and in the first operation of the batch file.	AUTO, MANUAL Default: None Required: No

Table 14-2 Keywords and values used in batch files (*continued*)

Keyword and description	Type and values
RESTARTWAITSECONDS - The time to make a second attempt following a backup failure. Use this keyword only once in the batch file and in the first operation of the batch file.	Integer number Default: 60 Required: No
RESTOREBEFOREMARK - Recovers the transaction log to a point before the occurrence of a transaction log mark.	String Default: None Required: No
RESTOREBEFOREMARK AFTERTIME - Recovers the transaction log to a point before the occurrence of a transaction log mark, but after a point in time (STOPAFTER).	String Default: None Required: No
RESTORECOPYNUM - Allows the agent to recover from non-primary copies. This number represents the copy number to use for restore. Range is 0-10. Copy 0 is the primary copy and a value of 1-10 represents a specific copy. Copy selection is only available with the NetBackup web UI when the user selects the copy along with a storage server and storage location.	Integer Default: 0 Required: No
RESTOREOPTION - Tells NetBackup to use the SQL Server replace option on a restore.	REPLACE Default: None Required: No
RESTOREPAGES - Ignored for a restore if the backup was performed with SQL Server Intelligent Policy. Specifies that NetBackup perform a page restore operation.	TRUE, FALSE Default: FALSE Required: No
RESTORETOMARK - Recovers the transaction log to a transaction log mark.	String Default: None Required: No
RESTORETOMARK AFTERTIME - Recovers the transaction log to a transaction log mark, but after a point in time (STOPAFTER).	String Default: None Required: No

Table 14-2 Keywords and values used in batch files (*continued*)

Keyword and description	Type and values
RESTORETYPE - Applicable only to RESTORE database operations. Full = Full database restore. Move = Database move. The batch file must contain a series of one or more <MOVE><filegroup> and <TO><file 9path> sequences. (batch file-based policies only) Partial = Partial database restore. The sequence for PARTIAL must specify all of the filegroups in the database whose backup image is referenced by the NBIMAGE keyword.	FULL, PARTIAL, MOVE Default: FULL Required: No
ROLLBACKVOLUME - Tells NetBackup to do the recovery of an Instant Recovery backup using the volume rollback method.	TRUE, FALSE Default: FALSE Required: No
SQLCOMPRESSION - Uses SQL Server compression on the backup image. If you enable SQL Server compression, do not enable NetBackup compression.	TRUE, FALSE Default: FALSE Required: No
SQLHOST - Name of SQL Server host. If SQLHOST is not specified in a batch file operation, then the SQL Server host is obtained from HKEY_CURRENT_USER\Software\VERITAS\NetBackup\NetBackup for Microsoft SQL Server\DEFAULT_SQL_HOST. If the SQLINSTANCE keyword is not included, then the default SQL Server instance is assumed for the SQL Host.	String Required: No
SQLINSTANCE - Name of the SQL Server instance. Or for backup operations specify \$ALL to designate all SQL Server instances including the default instance. If SQLINSTANCE \$ALL is used, then it must appear in the first operation of the batch file. Each operation in the batch file is performed for all SQL Server instances on the client where the batch file is executed. Also, it is not necessary to specify an SQLHOST or SQLINSTANCE on any subsequent operations.	String Required: No
STANDBYPATH - Specify a fully- qualified file path to use for the standby redo log.	String Default: None Required: No
STOPAFTER - Specifies datetime for RESTORETOMARK options. The datetime string is formatted as YYYY/MMDDHH:MM:SS.	Datetime string Default: None Required: No

Table 14-2 Keywords and values used in batch files (*continued*)

Keyword and description	Type and values
STOPAT - Specifies the point-in-time recovery of a transaction log. The datetime string is formatted as YYYY/MMDDHH:MM:SS .	Datetime string Default: None Required: No
STORAGEIMAGE - Used for restoring a database that was backed up using a grouped Snapshot Client snapshot. STORAGEIMAGE identifies the image with which the physical files are associated.	String Default: None Required: No
STRIPES - Number of stripes. Range is 1–32.	Integer Default: 1 Required: No
TO - Specifies a filegroup destination path. Required for each MOVE keyword. Also must sequentially follow each MOVE entry. The value may be delimited with single quotes.	File path Default: None Required: No
TRACELEVEL - Trace level.	MIN, MID, MAX Default: MIN Required: No
TRXOPTION - SQL Server transaction log backup options. If NOTRUNC is not selected, then the transaction log can be backed up and truncated. If TAILLOG is selected, the tail log is backed up and restored.	NOTRUNC, TAILLOG Default: None Required: No
USERID - User ID for logging into SQL Server. This keyword is ignored if you use integrated security.	String Default: sa Required: No
VDITIMEOUTSECONDS - Time-out interval for SQL Server Virtual Device Interface.	Integer Default: 300 Required: No
VERIFYONLY - Tells SQL Server to verify a backup image but not to restore it.	TRUE, FALSE Default: FALSE Required: No

Table 14-2 Keywords and values used in batch files (*continued*)

Keyword and description	Type and values
VERIFYOPTION - Valid for the databases that have an active page. STOPONERROR performs verification and stops if a verification error occurs. CONTINUEAFTERERROR performs verification but continues if a verification error occurs.	NONE, STOPONERROR CONTINUEAFTERERROR Default: NONE Required: No

Create a batch file

You can use any of the backup or restore dialog boxes to create a batch file that contains a NetBackup for SQL Server script.

Or you can launch the script from the `dbbackupx` command line program or through the NetBackup scheduler. See the example batch files.

[NetBackup for SQL Server sample batch files](#)

To create a batch file

- 1 Open the NetBackup MS SQL Client.
- 2 Select **File > Backup SQL Server objects** or **File > Restore SQL Server objects**.
- 3 Select the object you want to back up or restore.
- 4 Select the backup or restore options.

See [“Options for SQL Server backup operations”](#) on page 207.

See [“Options for NetBackup for SQL Server restores”](#) on page 222.

- 5 In the **Backup script** or **Restore script** group, click **Save**.
- 6 Click **Backup** or **Restore**.
- 7 Specify the following folder for the batch file:

`install_path\NetBackup\DbExt\MsSql\` folder.

Batch files must reside on the host from which they executed. If you perform actions on a remote host, the batch file must reside on that remote host.

See [“Registering authorized locations used by a NetBackup database script-based policy”](#) on page 290.

- 8 Give the file a unique name with the extension `.bch`.

- 9 Click **Save**.

Alternatively, you can select the name of an existing file and NetBackup appends the new script to it.

- 10 Click **Yes** to open and edit the batch file.

Run batch files

Once you have created a batch file, you manually run it from the NetBackup for SQL Server interface.

To run a batch file

- 1 Open the NetBackup MS SQL Client.
- 2 Log on to the host and instance you want to access.

See [“Select the SQL Server host and instance \(NetBackup MS SQL Client\)”](#) on page 216.
- 3 Select **File > Manage script files**.
- 4 Double-click the batch file.
- 5 Click **Start**.
- 6 To monitor the operation, select **File > View status**.

Add a batch file-based policy

This topic describes how to create a batch file-based policy that uses clients and batch files to perform backups.

Note: To perform multistreamed backups and restores, or if you have multiple network interfaces, you need to perform other configuration.

See [“Configure multistriped backups of SQL Server”](#) on page 88.

See [“Configuration and requirements for SQL Server backups with multiple NICs”](#) on page 245.

To add a batch file-based policy

- 1 Sign in to the primary server as a user that has the RBAC Administrator role or a role that can manage policies. If your site has more than one primary server, choose the one on which you want to add the policy.
- 2 On the left, select **Protection > Policies**.
- 3 Select **Add**.

- 4 Type a unique name for the new policy.
- 5 In the **Policy type** list, select **MS-SQL-Server**.
 The database agent policy type does not appear in the list unless your primary server has a license for the database agent.
- 6 Complete the entries on the **Attributes** tab.
 See [“About policy attributes”](#) on page 76.
- 7 On the **Instances and databases** tab, select **Clients for use with batch files**.
 The tab name changes to the name **Clients**. The **Backup selections** tab now lets you specify and browse for scripts.
- 8 Add other policy information as follows:
 - Add schedules.
 See [“Schedule properties for SQL Server batch file-based policies”](#) on page 201.
 - Add clients.
 See [“Add clients to a policy”](#) on page 205.
 - Add batch files to the backup selections list.
 See [“Add batch files to the backup selections list”](#) on page 206.
- 9 When you have added all the schedules, clients, and backup selections you need, select **Create**.

Schedule properties for SQL Server batch file-based policies

Each policy has its own set of schedules. These schedules initiate automatic backups and specify when a user can initiate operations. Some schedule properties that have a different meaning for database backups than for file system backups. Other schedule properties vary according to your specific backup strategy and system configuration. See the [NetBackup Administrator's Guide, Volume I](#).

Table 14-3 Description of schedule properties

Property	Description
Type of backup	Specifies the type of backup that this schedule can control. The selection list shows only the backup types that apply to the policy you want to configure. See “Schedule backup types for batch file-based policies” on page 202.

Table 14-3 Description of schedule properties (*continued*)

Property	Description
Schedule type	You can schedule an automatic backup in one of the following ways: ■ Calendar The Calendar option lets you schedule the backup operations that are based on specific dates, recurring week days, or recurring days of the month. ■ Frequency The Frequency specifies the period of time that can elapse until the next backup operation begins on this schedule. For example, assume that the frequency is 7 days and a successful backup occurs on Wednesday. The next full backup does not occur until the following Wednesday.
Effect of type of schedule on retention period	The type of schedule you select affects the retention period as follows: ■ Frequency-based scheduling Set a retention period that is longer than the frequency setting for the schedule. For example, if the frequency setting is set to one week, set the retention period to be more than one week. The NetBackup scheduler compares the latest record of the automatic backup schedule to the frequency of that automatic backup schedule. This comparison is done to determine whether a backup is due. So if you set the retention period to expire the record too early, the scheduled backup frequency is unpredictable. However, if you set the retention period to be longer than necessary, the NetBackup catalog accumulates unnecessary records. ■ Calendar-based scheduling The retention period setting is not significant for calendar-based scheduling.

Schedule backup types for batch file-based policies

[Table 14-4](#) shows that the backup types you can specify for a SQL Server batch file-based policy that uses clients and batch files. Intelligent Policies have a different set of backup types.

Table 14-4 Backup types for batch file-based policies

Backup type	Description
Application Backup	The application backup schedule enables user-controlled NetBackup operations from the client. These operations include those initiated from the client and those initiated by a full schedule on the primary server. NetBackup uses the application backup schedule when the user starts a backup manually. Configure at least one application backup schedule for each database policy. The Default-Application-Backup schedule is configured automatically as an application backup schedule.

Table 14-4 Backup types for batch file-based policies (*continued*)

Backup type	Description
Full Backup	This schedule specifies the dates and times for NetBackup to automatically start backups as indicated in the batch file (full, differential, or transaction log). NetBackup runs the batch files in the order that they appear in the file list. If there is more than one client in the policy, the batch files are run on each client. See “Keywords and values used in batch files” on page 191. See “Converting differential backups to full backups” on page 89.

Configure an application backup schedule

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

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

For any policies that include read-only filegroups, consider creating a schedule with a retention level set to infinity. This level can enable you to avoid redundant backups.

To configure an application backup schedule

- 1 Open the policy and select the **Schedules** tab.
- 2 Select the schedule that is named **Default-Application-Backup** and select **Edit**.
- 3 Specify the other properties for the schedule.
See [“Schedule properties for SQL Server batch file-based policies”](#) on page 201.
- 4 Select **Add**.

Example application backup schedule

Assume the following:

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

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

Table 14-5 Example settings for a NetBackup for SQL Server application backup schedule

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

Configure automatic backup schedules

If you put multiple batch files in the same policy, they run during each automatic backup session for that policy. You may have a variety of SQL Server backup operations that you want to run on different schedules. In this case, you may want to create multiple policies each with an automatic backup schedule that is different. Then assign each batch file to the policy that uses the appropriate automatic backup schedule.

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

To configure an automatic backup schedule

- 1 Open the policy and select the **Schedules** tab.
- 2 Click **Add**.
- 3 Specify a unique name for the schedule.
- 4 Select the **Full backup** schedule.
See [“Schedule backup types for batch file-based policies”](#) on page 202.
- 5 Specify the other properties for the schedule.
See [“Schedule properties for SQL Server batch file-based policies”](#) on page 201.
- 6 Select **Add**.

Example automatic backup schedule

Table 14-6 shows example settings for an automatic backup schedule.

Table 14-6 Example settings for a NetBackup for SQL Server automatic backup schedule

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

Add clients to a policy

The client list is the list of hosts on which your batch files are run during an automatic backup. A NetBackup client must be in at least one policy but can be in more than one.

To add clients to a policy

- 1 Open the policy and select the **Clients** tab.
- 2 Before you can add clients, you must select **Clients for use with batch files** on the **Instances and databases** tab.
- 3 Select **Add**.
- 4 Type the name of the client and select the hardware and operating system of the client.

If SQL Server is installed in a cluster, specify the virtual name of the SQL Server as the client name.

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

See [“Reviewing the auto-discovered mappings”](#) on page 26.

See [“Configuring mappings for restores of a distributed applications, clusters, or virtual machines ”](#) on page 31.

- 5 Select **Add**.

Add batch files to the backup selections list

The backup selections list in a database policy has a different meaning than for non-database policies. For example, in a Standard or Microsoft Windows policy, the list contains files and directories to be backed up. In a database policy, you can specify batch files to run. (For NetBackup for SQL Server, the scripts are called batch files and have the `.bch` extension.) Batch files describe the backup operations you want to start. You can start them by initiating manual or scheduled operations from the NetBackup server. These files reside on the client and direct the operation of NetBackup for SQL Server and SQL Server.

Add batch files if you want a policy that runs scheduled backups. All batch files that are listed in the backup selections list are run for manual backups and for automatic backup schedules. NetBackup runs the batch files in the order that the batch files appear in the backup selections list.

To add batch files to the backup selections list

- 1 Ensure that the batch file resides on the client.

See [“Registering authorized locations used by a NetBackup database script-based policy”](#) on page 290.
- 2 Open the policy you want to edit or create a new policy.
- 3 Before you can add batch files, you must do the following:
 - On the **Instances and databases** tab, select **Clients for use with batch files**.
 - On the **Clients** tab, add one or more clients.
- 4 Select the **Backup selections** tab.
- 5 Select **Add**.
- 6 In the **Script** box, type the full path name of a batch file on the client, then click **Add**.

For example:

`install_path\NetBackup\DbExt\Mssql\bkup.bch`

You must indicate the full pathname of the batch file.
- 7 Add any other batch files.
- 8 Select **Add** to add the batch files to the backup selections list.

Options for SQL Server backup operations

[Table 14-7](#) describes the options that are available when you perform backups from the MS SQL Client. These options appear in the **Backup Microsoft SQL Server Objects** dialog box after you select **File > Backup SQL Server objects**.

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.

Table 14-7 Options for SQL Server backup operations

Option	Description
Expand database	This pane lets you traverse live databases. You can expand the SQL Server instance to view its databases. Expand each database to view its filegroups or expand a filegroup to view its files. You can select any object in this pane to view its constituent objects in the right-hand pane.
Select database(s) for backup from <i>instance</i> <i>host\instance</i>	Select the objects that you want to back up from this pane. This pane displays the list of constituent database objects of the selected host and instance in the left-hand pane. You can select one or more objects (databases) in this pane.
Type of Backup	The following backup types are available: ■ Full Create a full database backup. ■ Full differential Create a differential backup. ■ transaction log Create a transaction log backup. This type of backup is only available for databases. When you select this type of backup, you then need to select a backup option from the Transaction log backup options list. ■ Read/write filegroups Create a backup of read-write filegroups in a database. ■ Differential on read/write filegroups Create a differential backup of read-write filegroups in a database. ■ Create a template for partial backup Create a backup of only the selected filegroups in a database. ■ Create a template for partial differential backup Create a differential backup of only the selected filegroups in a database.

Table 14-7 Options for SQL Server backup operations (*continued*)

Option	Description
Transaction log backup options	The following options are available when you have chosen a transaction log backup type: ■ Back up and truncate transaction log Back up the transaction log and remove the inactive part of the transaction log. ■ Back up transaction log, but do not truncate it Back up a transaction log without truncating it. ■ Back up and restore tail log Back up and recover the tail log from disk.
Use SQL compression	Select this option if you want to use SQL Server to compress the backup image. If you enable SQL Server compression, do not enable NetBackup compression.
Backup script	■ Launch Immediately Start the backup operation immediately. Launch immediately is disabled if you are logged into a SQL Server instance that is not on the local host. If you generate a script for a non-local host, then it must be executed on that host. ■ Save Generate a script that can be started at a later time.
Back up	In the right-hand pane, choose one of the following backup options: ■ Selected Back up only the objects selected. ■ All but selected Back up all of the objects, except those selected. ■ All Back up all of the objects.
Stripes	Set the number of backup stripes that you want SQL Server to create for your backup. Type a number from 1 to 32. 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. See “Configure multistriped backups of SQL Server” on page 88.
Resume options for this selection	■ Do not resume unsuccessful backups ■ Retry from the beginning Restart failed backups after waiting 60 seconds.

Table 14-7 Options for SQL Server backup operations (*continued*)

Option	Description
NetBackup policy	If this host is the NetBackup primary server, then this list includes all active policies of type MS-SQL-Server. You can select one of these policies or type the name of a policy. The default is <any>. If you select the default, then NetBackup selects which MS-SQL-Server policy to use.
Page verification	This option is enabled for objects have a page verification type that is either torn page detection or checksum. All of the objects in the right-hand pane must have the proper verification type. This indicates a performance penalty when you use page verification. ■ Do not perform verification Do not perform page verification before you run the backup. ■ Perform verification Perform page verification when you run the backup and stop the backup if a verification error is encountered.
Backup	Start a database backup or generate a database backup script. This option is enabled only when you select an object to back up.

Create a script to backup a remote SQL Server installation

You can use the NetBackup MS SQL Client to create a script to back up databases on a remote host. Generated batch files must be saved on the remote host. You can launch the operation from the local installation of NetBackup for SQL Server, from an automatic backup policy, or from a manual backup.

To create a script to backup of a remote SQL Server installation

- 1 Select the host and instance you want to access.
See [“Select the SQL Server host and instance \(NetBackup MS SQL Client\)”](#) on page 216.
- 2 Select **File > Backup SQL Server objects**.
- 3 Select the options for the operation.
See [“Options for SQL Server backup operations”](#) on page 207.
Save is enabled in the backup dialog box. **Launch immediately** is disabled because the generated script must be run on the remote host that you are logged on to.
- 4 Click **Backup**.

- 5 In the **Save Script As** dialog box, navigate to the `install_path\NetBackup\DbExt\MsSql\` folder on the remote host, and save the batch file there.
- 6 Launch the backup operation.
 Do one of the following:
 - Run the operation from the local installation of NetBackup for SQL Server.
 - Create a new policy that includes the remote SQL Server client. Add the batch file to the **Backup selections** list in the policy.

About automatic retry of unsuccessful SQL Server backups

NetBackup for SQL Server provides the following options to retry unsuccessful backup attempts.

Automatic retry	NetBackup for SQL Server keeps track of the unsuccessful backups that may have resulted from the execution of a batch file. When the initial backup attempt is complete, the agent rewrites the batch file, including only those operations that failed. The rewritten batch file is launched automatically.
Manual retry	A manual retry is similar to an automatic retry except that NetBackup does not launch the rewritten batch file. Instead it is written to the <code>install_path\dbext\mssql\temp</code> directory. The user can then choose when to run the new batch file.

To use automatic retry, add the following line to your batch file.

```
RESTARTTYPE AUTO
```

By default, the unsuccessful backups are retried one time automatically after 60 seconds. To change the delay following the unsuccessful attempt, then add the following to your batch file.

```
RESTARTWAITSECONDS <integer>
```

You can also specify the number of retries. Add the following to your batch file.

```
NUMRESTARTS <1 to 9>
```

To use manual retry, add the following line to your batch file.

```
RESTARTTYPE MANUAL
```

Note: All of the keyword-value pairs that are described in this topic are only permitted in the first operation of the batch file.

Configure a batch file-based policy for a user-directed backup of read-only filegroups

When you separate read-only and read-write filegroups in your backup strategy, you can reduce total media usage and the total time you spend on backup operations. To back up read-only filegroups you must first create a separate policy for this type of backup. You can also verify that all read-only filegroups are backed up.

See [“View SQL Server read-only backup sets \(NetBackup MS SQL Client\)”](#) on page 211.

To back up read-only filegroups

- 1 Open the NetBackup MS SQL Client interface.
- 2 Create a batch file that includes the read-only filegroups.
 All read-only filegroups must be included in some combination of full, or individual filegroup and file backups. You only need to perform this backup one time.
- 3 In the NetBackup web UI, create a backup policy for read-only filegroups.
 - In the Application Backup schedule, set the **Retention** level of **Infinite**.
 - Add the batch file that you created to the backup selections list.
- 4 Back up the read-only filegroups.
- 5 If necessary, confirm that all read-only groups are backed up by viewing the read-only backup set.

See [“View SQL Server read-only backup sets \(NetBackup MS SQL Client\)”](#) on page 211.

View SQL Server read-only backup sets (NetBackup MS SQL Client)

If you perform periodic backups only on read-write filegroups, you can verify if you have retained backups of the read-only filegroups.

To view read-only backup sets

- 1 Open the NetBackup MS SQL Client interface.
- 2 Browse for the backup images that contain the read-only backup sets.
 See [“Browsing for SQL Server backup images \(NetBackup MS SQL Client\)”](#) on page 225.
- 3 In the **Restore Microsoft SQL Server Objects** dialog box, expand the instance name.
- 4 Right-click the database and select **Properties**.
- 5 Click the "Read-only backup set" tab.
 If the database does not contain read-only filegroups, then the message "This database does not contain any read-only filegroups." is shown. If backups do not exist for all of the read-only filegroups, then a list of the filegroups that were not backed up is shown. Finally, if a backup is found of all of the read-only filegroups, then the name appears of the latest image that contains this backup.
- 6 If there are any read-only filegroups that are not backed up, back them up as soon as possible. These backups ensure you can perform a full recovery.
- 7 Click **OK**.

Configure a batch file-based policy for a user-directed backup of read-write filegroups

When you separate read-only and read-write filegroups in your backup strategy, you can reduce total media usage and the total time you spend on backup operations. More information is available on backing up read-only filegroups.

See [“Configure a batch file-based policy for a user-directed backup of read-only filegroups”](#) on page 211.

Note: Immediately back up any filegroup when you change it from read-write to read-only.

To back up read-write filegroups

- 1 Open the NetBackup MS SQL Client interface.
- 2 Select **File > Backup SQL Server objects**.
- 3 In the left pane, select the database instance.
- 4 In the right pane, select one or more databases that you want to back up.

- 5 Select the **Type of Backup**, as follows:
 - To perform a full backup of the read-write filegroups, select **Read-write filegroups**.
 - To perform a differential backup of the read-write filegroups, select **Differential on read-write filegroups**.
- 6 Select the backup options.
See [“Options for SQL Server backup operations”](#) on page 207.
- 7 From the **Backup script** group, select **Save**.
- 8 Click **Backup**.

Note the location where the batch file is saved. This batch file is added to the policy that backs up the read-write filegroups.
- 9 Open the NetBackup web UI.
- 10 Create a backup policy for read-write filegroups.
 - Create a **Full backup** schedule with the wanted retention period.
 - Add the batch file that you created to the backup selections list.
- 11 (Optional) Manually back up the read-write filegroups.

If you do not perform a manual backup at this time, the backup runs automatically through the schedule you created in step 10.

Performing backups and restores with the NetBackup MS SQL Client

This chapter includes the following topics:

- [About the NetBackup MS SQL Client](#)
- [Start the NetBackup MS SQL Client for the first time](#)
- [Select the SQL Server host and instance \(NetBackup MS SQL Client\)](#)
- [About viewing the properties of the objects selected for backup](#)
- [Perform a user-directed backup of SQL Server databases \(NetBackup MS SQL Client\)](#)
- [Perform a user-directed backup of SQL Server transaction logs \(NetBackup MS SQL Client\)](#)
- [Perform a user-directed backup of SQL Server database filegroups \(NetBackup MS SQL Client\)](#)
- [Perform a user-directed backup of SQL Server database files \(NetBackup MS SQL Client\)](#)
- [Perform a partial database backup \(NetBackup MS SQL Client\)](#)
- [Options for NetBackup for SQL Server restores](#)
- [Browsing for SQL Server backup images \(NetBackup MS SQL Client\)](#)
- [Restore a SQL Server database backup \(NetBackup MS SQL Client\)](#)

- [Stage a full SQL Server database recovery \(NetBackup MS SQL Client\)](#)
- [Restore SQL Server filegroup backups \(NetBackup MS SQL Client\)](#)
- [Recover a SQL Server database from read-write filegroup backups \(NetBackup MS SQL Client\)](#)
- [Restore SQL Server read-only filegroups \(NetBackup MS SQL Client\)](#)
- [Restore SQL Server database files \(NetBackup MS SQL Client\)](#)
- [Restore a SQL Server transaction log image without staging a full recovery \(NetBackup MS SQL Client\)](#)
- [Perform a SQL Server database move \(NetBackup MS SQL Client\)](#)
- [About performing a SQL Server page-level restore \(NetBackup MS SQL Client\)](#)
- [Redirect a SQL Server database to a different host \(NetBackup MS SQL Client\)](#)
- [Perform a restore of a remote SQL Server installation \(NetBackup MS SQL Client\)](#)
- [Restoring multistreamed SQL Server backups](#)
- [About using bplist to retrieve SQL Server backups](#)
- [About NetBackup for SQL Server backup names](#)

About the NetBackup MS SQL Client

The NetBackup MS SQL Client is the original interface that was used to perform user-directed backups of SQL Server and to create scripts for backup policies and for restore operations. Much of the functionality in this interface is now available in the NetBackup web UI.

Start the NetBackup MS SQL Client for the first time

This topic describes how to start the NetBackup MS SQL Client for the first time. For subsequent sessions, the agent remembers the information that you provided.

To start the NetBackup MS SQL Client for the first time

- 1 If you use SQL Server integrated security, log on to the Windows host with the Windows account that has permissions to perform SQL Server backups and restores.
- 2 Open the NetBackup MS SQL Client.
- 3 When you are prompted to provide the logon parameters, click **OK**.
- 4 Select the SQL Server host and instance that you want to log into.
- 5 If the SQL Server host and instance use standard or mixed security, provide the SQL Server user ID and password.
- 6 Click **Apply > Close**.

Select the SQL Server host and instance (NetBackup MS SQL Client)

Use this procedure to set which SQL Server host and the instance that you want the NetBackup MS SQL Client to access.

The user ID and password are only required if the host uses standard or mixed security. If applicable, you only need to provide these credentials when you first open the NetBackup MS SQL Client.

To select the SQL Server host and instance

- 1 Open the NetBackup MS SQL Client.
- 2 Select **File > Set SQL Server connection properties**.
- 3 From the **Host** list, select the SQL Server host.

You can type a host name if it does not appear in the list. If you select a remote host and click **Apply**, the **Host type** is shown as "remote".
- 4 From the **Instance** list, select the SQL Server instance.

You can type an instance name if it does not appear in the list. You can designate the default instance either by setting the **Instance** value to <default> or to empty (no spaces).
- 5 Click **Apply > Close**.

About viewing the properties of the objects selected for backup

You can view the properties of any object in the **Backup Microsoft SQL Server Objects** dialog box by right-clicking the object. [Table 15-1](#) describes the properties of objects that are selected for backup.

To view the properties of an object that is selected for backup

- 1 Select **File > Backup SQL Server objects**.
- 2 In the **Backup Microsoft SQL Server Objects** dialog box, in the right pane, right-click an object and select **Properties**.
- 3 When you finish, click **OK**.

Table 15-1 Properties of the objects that are selected for backup

Property	Description
Object type	Database, database filegroup, database file, or transaction log.
Object name	Name of the object.
Parent (database, instance, filegroup, etc.)	Name of the object's parent.
SQL Server instance	SQL Server instance the object belongs to.
File size	The size of the component files. This size should closely match the size of a backup snapshot.
Data size	Size of the backup stream. Applies to databases only.
Page verification	The type of SQL Server page verification that is configured for selected databases, filegroups, and logical files. The available values are: none, torn page detection, or checksum.
Read-only/read-write	The attribute that is applied to the filegroup.
On-line/off-line	The status of the filegroup.
Path	(Database files only) The absolute path of the database file.

Perform a user-directed backup of SQL Server databases (NetBackup MS SQL Client)

This procedure describes how to perform a database backup.

To perform a user-directed backup of a SQL Server database

- 1 Open the NetBackup MS SQL Client interface.
- 2 Select the host and instance you want to access.

See [“Select the SQL Server host and instance \(NetBackup MS SQL Client\)”](#) on page 216.
- 3 Select File > **Backup SQL Server objects**.
- 4 In the **Backup Microsoft SQL Server Objects** dialog box, in the left pane, select the database instance.
- 5 In the right pane, select one or more databases that you want to back up.
- 6 Select the **Type of Backup**.

Select one of the following:
 - To perform a full backup, select **Full Backup**.
 - To back up the database with the differential option, select **Perform differential backup**.
- 7 Select the backup options.

See [“Options for SQL Server backup operations”](#) on page 207.
- 8 Click **Backup**.
- 9 When you are prompted to start the backup, click **Yes**.
- 10 To view the progress of the backup, select File > **View status**.

Perform a user-directed backup of SQL Server transaction logs (NetBackup MS SQL Client)

This procedure describes how to perform a transaction log backup.

Caution: Ensure that the entire sequence of transaction logs that are generated following any database backup are maintained on the same NetBackup server. Back up all transaction logs to the same facility and do not allow any logs to expire before the others.

To back up a transaction log

- 1 In SQL Server, set the **Recovery Model** setting to either **Full** or **Bulk-logged**.
- 2 Open the NetBackup MS SQL Client interface.
- 3 Select the host and instance you want to access.
 See [“Select the SQL Server host and instance \(NetBackup MS SQL Client\)”](#) on page 216.
- 4 Select **File > Backup SQL Server Objects**.
- 5 In the **Backup Microsoft SQL Server Objects** dialog box, in the left pane, select the database instance.
- 6 In the right pane, select one or more databases whose transaction logs you want to back up.
- 7 In the **Type of Backup** list, select **transaction log**.
- 8 From the list, select the transaction log option. For more information, see the following table.

Back up and truncate transaction log	Back up the transaction log and remove the inactive part of the transaction log.
Truncate transaction log, but don't back it up	Truncate the log without performing a backup.
Back up and restore tail log	Back up and recover the tail log from the disk.

- 9 Select the backup options.
 See [“Options for SQL Server backup operations”](#) on page 207.
- 10 Click **Backup**.
 To view the progress of the backup, select **File > View status**.

Perform a user-directed backup of SQL Server database filegroups (NetBackup MS SQL Client)

More information is available on how to use read-write and read-only filegroups in your backup strategy.

See [“Configure a batch file-based policy for a user-directed backup of read-write filegroups”](#) on page 212.

See [“Configure a batch file-based policy for a user-directed backup of read-only filegroups”](#) on page 211.

To back up a database filegroup

- 1** Open the NetBackup MS SQL Client interface.
- 2** Select the host and instance you want to access.

See [“Select the SQL Server host and instance \(NetBackup MS SQL Client\)”](#) on page 216.
- 3** Select File > **Backup SQL Server objects**.
- 4** In the **Backup Microsoft SQL Server Objects** dialog box, in the left pane, expand the instance name.
- 5** Select a database whose filegroups you want to back up.
- 6** In the right pane, select one or more filegroups that you want to back up.
- 7** Select the backup options.

See [“Options for SQL Server backup operations”](#) on page 207.
- 8** Click **Backup**.

To view the progress of the backup, select File > **View status**.

Perform a user-directed backup of SQL Server database files (NetBackup MS SQL Client)

This procedure describes how to back up database files.

To back up a database file

- 1** Open the NetBackup MS SQL Client interface.
- 2** Select the host and instance you want to access.

See [“Select the SQL Server host and instance \(NetBackup MS SQL Client\)”](#) on page 216.
- 3** Select File > **Backup SQL Server objects**.
- 4** In the **Backup Microsoft SQL Server Objects** dialog box, in the left pane, expand the instance name and database.
- 5** In the left pane, select the filegroup that contains the files you want to back up.
- 6** In the right pane, select one or more files that you want to back up.

- 7 Select the backup options.
See [“Options for SQL Server backup operations”](#) on page 207.
- 8 Click **Backup**.
To view the progress of the backup, select File > **View status**.

Perform a partial database backup (NetBackup MS SQL Client)

This procedure describes how to create a script to perform a partial database backup. This type of back is only available for batch file-based policies.

To perform a partial database backup

- 1 Open the NetBackup MS SQL Client interface.
- 2 Select the host and instance you want to access.
See [“Select the SQL Server host and instance \(NetBackup MS SQL Client\)”](#) on page 216.
- 3 Select **File > Backup SQL Server objects**.
- 4 In the **Backup Microsoft SQL Server Objects** dialog box, in the left pane, select the database instance.
- 5 In the right pane, select a database that you want to back up.
- 6 For the **Type of Backup**, select one of the following:
 - **Create a template for partial backup.**
 - **Create a template for partial differential backup.**
- 7 Select the backup options.
See [“Options for SQL Server backup operations”](#) on page 207.
- 8 Click **Backup**.
- 9 In the **Save Script As** dialog box, specify a file name and click **OK**.
- 10 When you are prompted to open the template, click **Yes**.

- 11 Edit the template by uncommenting the filegroups that you want to include in the backup. You must uncomment at least one filegroup.

For example, replace:

```
#
# If you wish to include filegroup DBA_FG1 in the partial backup,
# then remove the hash mark that precedes the following line.
#FILEGROUP DBA_FG1
```

with:

```
#
# If you wish to include filegroup DBA_FG1 in the partial backup,
# then remove the hash mark that precedes the following line.
FILEGROUP DBA_FG1
```

- 12 When you are finished modifying the template, save it.
- 13 To run the backup, select File > **Manage script files**, select the script you created, and click **Start**.

Options for NetBackup for SQL Server restores

Table 15-2 describes the options that are available when you perform restores.

Table 15-2 Options for restore operations

Option	Description
Scripting	These scripting options are available for restoring from a database image: ■ Restore selected object Produce a script that performs a database restore. This script is the default option. ■ Create a move template Create a script template for moving the selected database. ■ Restore read-only filegroups Restore the most recent backup of every read-only filegroup. ■ Create a page restore template Create a template for restoring a database, filegroup, or file from the pages that are contained in the selected backup image. The Microsoft SQL Server service must have full access permission to the folder <code>install_path\Netbackup\dbext\mssql\temp</code>. ■ Verify backup image, but don't restore This option is only available if the image was backed up with the page verification option. NetBackup processes the image for errors, but does not perform a restore.

Table 15-2 Options for restore operations (*continued*)

Option	Description
Use replace option	Restore with the SQL Server replace option.
Recovery	Specify one of the SQL Server recovery options. Not recovered Use this option during a restore if additional backup images must be applied to the database following the current restore. When you use this option, the database is left in a loading state. Recovered Restore the last image in a restore sequence. After the recovery operation, the database is ready for use. If you do not select this option, the database is in an intermediate state, and is not usable. If Recovered is selected when an intermediate backup is applied, you cannot continue to restore backups; you must restart the restore operation from the beginning. 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 <code>install_path\NetBackup\logs\SQLStandBy\</code>. The account that runs the Microsoft SQL Server service must have full access permission to the <code>SQLStandBy</code> folder. The database is placed in "standby" state following the restore.
Consistency check	Select the consistency check to perform after the restore. Output from the consistency check is written to the SQL Server client progress log. You cannot select consistency checking unless the database is restored to the recovered state. If you select consistency checking for a staged recovery, then the check occurs following the last restore. None Do not perform consistency checking. 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. Physical check only Perform a low overhead check of the physical consistency of the SQL Server database. This option only checks the integrity of the physical structure of the page and record headers. It also checks the consistency between the pages' object ID and index ID and the allocation structures. Full check, including indexes Include indexes in the consistency check. Any errors are logged. Check catalog Check for consistency in and between system tables in the specified database.

Table 15-2 Options for restore operations (*continued*)

Option	Description
Page verification	Note: A performance penalty can happen when you use page verification. These options are available if the source object was backed up with torn page detection or checksum verification. ■ Do not perform verification Do not include page verification in the restore script. ■ Perform verification Include page verification in the restore script and stop the restore if an error is encountered.
Stage full recovery	Perform a complete database restore using the recovery set that NetBackup found.
Restore selected transaction log	Restore only the selected transaction log.
Transaction log recovery options	This list contains the controls for you to restore a transaction log. You can restore the log to a point in time that precedes the time when the transaction log was dumped. ■ To point in time Recover the transaction log to a point in time. ■ To transaction log mark Recover the transaction log to a transaction log mark. ■ To transaction log mark but after Recover the transaction log to a transaction log mark but after a point in time. ■ Before transaction log mark Recover the transaction log recovered to a point before the occurrence of a transaction log mark. ■ Before transaction log mark but after Recover the transaction log to a point before the occurrence of a transaction log mark but after a point in time. ■ Entire transaction log Restore the entire log.
Transaction log mark	The transaction log mark you want to use for recovery.
YYYY, MM, DD, HH, MM, SS am, pm	Specify the time to which you want the transaction logs restored.
Launch immediately	Start the restore operation immediately. Launch immediately is disabled if you are logged into a SQL Server instance that is not on the local host. If you generate a script for a non-local host, you must run it from on that host.
Save	Generate a script that can be started at a later time.

Table 15-2 Options for restore operations (continued)

Option	Description
Restore	Start the restore or generate a restore script.

Browsing for SQL Server backup images (NetBackup MS SQL Client)

This procedure describes how to browse for a backup image from which you want to restore.

If you have multiple NICs, backups from a UNIX server, or a NetBackup client name with a qualified domain name or an IP address, see the following:

See [the section called “How NetBackup resolves SQL Server host and instance names”](#) on page 226.

To browse for backup images

- 1 Open the NetBackup MS SQL Client.
- 2 Change the host and instance you want to access.

See [“Select the SQL Server host and instance \(NetBackup MS SQL Client\)”](#) on page 216.
- 3 Select **File > Restore SQL Server objects**.
- 4 Select the **SQL Host** whose backup images you want to browse, or type its name.
- 5 Indicate the **Source Client**, if applicable.
 - When the NetBackup client name and the host name are different you also need to also provide the **Source Client** name. For example, if the NetBackup client name is the network interface name.
 - For Intelligent Policies, you also need to indicate the **Source Client** if you add or register the instance with a host name that is different than the NetBackup client name.
- 6 (Optional) In the **Database name filter** box, provide a keyword or query to match databases with that name. Filtering on the database name can significantly reduce the time it takes for NetBackup to return the list of backup images.

- 7 Select the date range to search and click **OK**.
- 8 Continue with the applicable instructions for how to restore the objects.
 - See [“Restore a SQL Server database backup \(NetBackup MS SQL Client\)”](#) on page 227.
 - See [“Stage a full SQL Server database recovery \(NetBackup MS SQL Client\)”](#) on page 227.
 - See [“Restore SQL Server filegroup backups \(NetBackup MS SQL Client\)”](#) on page 228.
 - See [“Recover a SQL Server database from read-write filegroup backups \(NetBackup MS SQL Client\)”](#) on page 229.
 - See [“Restore SQL Server read-only filegroups \(NetBackup MS SQL Client\)”](#) on page 230.
 - See [“Restore SQL Server database files \(NetBackup MS SQL Client\)”](#) on page 230.
 - See [“Restore a SQL Server transaction log image without staging a full recovery \(NetBackup MS SQL Client\)”](#) on page 231.
 - See [“Perform a SQL Server database move \(NetBackup MS SQL Client\)”](#) on page 232.
 - See [“About performing a SQL Server page-level restore \(NetBackup MS SQL Client\)”](#) on page 234.

How NetBackup resolves SQL Server host and instance names

To ensure that NetBackup displays the backup images you want, consider the following special cases:

- If backups are performed on a different network, the images are stored under the network interface name and not the NetBIOS name.
See [“Restore SQL Server when you have multiple NICs \(NetBackup MS SQL Client\)”](#) on page 249.
- Backups from a UNIX server. Since UNIX names are case-sensitive, you must provide the exact client name in the **Source Client** box field.
SQL Host: TIGER
Source Client: Tiger
- The NetBackup client name is a qualified domain name. The SQL Server host name or registered host name (Intelligent Policies) is the NetBIOS name. Specify the **SQL Host** as the NetBIOS name and the **Source Client** as the fully qualified domain name.

SQL Host: Tiger

Source Client: tiger.apexworks.com

- The NetBackup client name is an IP address. The SQL Server host name or registered host name (Intelligent Policies) is the NetBIOS name. Specify the **SQL Host** as the NetBIOS name and the **Source Client** as the IP address:

SQL Host: Tiger

Source Client: 10.80.136.68

Restore a SQL Server database backup (NetBackup MS SQL Client)

This topic describes how to restore a database from a full database or differential database backup.

To restore a database backup

- 1 Open the NetBackup MS SQL Client.
- 2 Browse for the backup images you want to restore.
See [“Browsing for SQL Server backup images \(NetBackup MS SQL Client\)”](#) on page 225.
- 3 Expand the database instance and the database.
- 4 Select the database image that you want to restore, as follows:
 - To restore a full backup, select the image of the database backup.
 - To restore a full backup and a differential database backup, click the "+" and select a differential backup.
- 5 Select the restore options.
See [“Options for NetBackup for SQL Server restores”](#) on page 222.
- 6 Click **Restore**.

Stage a full SQL Server database recovery (NetBackup MS SQL Client)

This topic describes how to stage a full database recovery. Alternatively, you can restore without staging a full recovery.

See [“Restore a SQL Server transaction log image without staging a full recovery \(NetBackup MS SQL Client\)”](#) on page 231.

To stage a full database recovery

- 1 Open the NetBackup MS SQL Client.
- 2 Browse for a backup image that contains the point in time to which you want to recover.

See [“Browsing for SQL Server backup images \(NetBackup MS SQL Client\)”](#) on page 225.
- 3 Expand the database instance.
- 4 Click the "+" next to the database that contains the transaction log backup you want to restore.
- 5 Select the transaction log image that includes the point in time from which you want to recover.
- 6 Select **Stage full recovery**.

When you view the properties of the transaction log, a **Recovery Set** tab displays.

The recovery set can include any combination of backup images that are sufficient for staging the full recovery. These can include full database, filegroup, and differentials.

- 7 Click **Restore**.

Restore SQL Server filegroup backups (NetBackup MS SQL Client)

This topic describes how to restore a backup of a filegroup. If you scheduled backups only include read-write filegroups, see the following topics.

See [“Recover a SQL Server database from read-write filegroup backups \(NetBackup MS SQL Client\)”](#) on page 229.

See [“Restore SQL Server read-only filegroups \(NetBackup MS SQL Client\)”](#) on page 230.

Note: If you attempt to restore a single differential backup without first restoring the preceding database backup file, SQL Server halts the load process. An error such as 4305 or 4306 is displayed. If you plan to restore a single differential, then you are responsible for first restoring the database backup file. You can avoid this problem by backing up the entire sequence of transaction logs. Also back up the differential backup and the backup file to the same NetBackup server. Then you can restore the entire sequence of backup objects.

See [“Stage a full SQL Server database recovery \(NetBackup MS SQL Client\)”](#) on page 227.

To restore a filegroup backup

- 1 Open the NetBackup MS SQL Client.
- 2 Browse for the backup images you want to restore.
 See [“Browsing for SQL Server backup images \(NetBackup MS SQL Client\)”](#) on page 225.
- 3 Expand the database instance and database.
- 4 Expand the filegroup and select a filegroup image to restore, as follows:
 - To restore a full backup, select the image of the filegroup backup.
 - To restore a differential filegroup backup, click the "+" next to the full backup and select the differential backup.
- 5 Click **Restore**.

Recover a SQL Server database from read-write filegroup backups (NetBackup MS SQL Client)

NetBackup for SQL Server automatically generates the most efficient recovery path when you select a transaction log image for restore. The recovery path can be based on read-write filegroups if you use them in your backup strategy. After restoring the read-write filegroups, you can bring the database online without having to restore the read-only filegroups.

To recover a database from read-write filegroups

- 1 Open the NetBackup MS SQL Client.
- 2 Browse for the backup images you want to restore.
 See [“Browsing for SQL Server backup images \(NetBackup MS SQL Client\)”](#) on page 225.

- 3 Expand the database instance.
- 4 Expand the database that contains the read-write filegroups you want to restore.
- 5 Select the transaction log backup.
- 6 Right-click the transaction log backup and select **Properties**.
- 7 On the **Recovery set** tab, verify that a complete backup set is available and click **OK**.
- 8 Click **Restore**.

See [“Restore SQL Server read-only filegroups \(NetBackup MS SQL Client\)”](#) on page 230.

Restore SQL Server read-only filegroups (NetBackup MS SQL Client)

This topic describes how to restore read-only filegroups.

To restore read-only filegroups

- 1 Open the NetBackup MS SQL Client.
- 2 Browse for the backup images you want to restore.

See [“Browsing for SQL Server backup images \(NetBackup MS SQL Client\)”](#) on page 225.

Be sure that the start date for the Time Filter is early enough to include the timestamp of the earliest backup of the read-only filegroups.

- 3 Expand the database instance.
- 4 Select the database that contains the read-only filegroups you want to restore.
In the **Scripting** list, **Restore read-only filegroups** is selected.
The restore option is enabled if a full set of read-only filegroups is available.
- 5 Click **Restore**.

Restore SQL Server database files (NetBackup MS SQL Client)

This topic describes how to restore database files.

To restore a database file

- 1 Open the NetBackup MS SQL Client.
- 2 Browse for the backup images you want to restore.
See [“Browsing for SQL Server backup images \(NetBackup MS SQL Client\)”](#) on page 225.
- 3 Expand the database instance and the database.
- 4 Expand the filegroup and the file.
- 5 Select the database file image that you want to restore.
- 6 Click **Restore**.

Restore a SQL Server transaction log image without staging a full recovery (NetBackup MS SQL Client)

This topic describes how to restore a transaction log image without staging a full recovery. Alternatively, you can stage a full recovery.

See [“Stage a full SQL Server database recovery \(NetBackup MS SQL Client\)”](#) on page 227.

To restore a transaction log without staging a full recovery

- 1 Open the NetBackup MS SQL Client.
- 2 Browse for the backup images you want to restore.
See [“Browsing for SQL Server backup images \(NetBackup MS SQL Client\)”](#) on page 225.
- 3 Expand the database instance.
- 4 Select the transaction log image that you want to restore.
- 5 Select **Restore selected transaction log**.
- 6 Click **Restore**.

Perform a SQL Server database move (NetBackup MS SQL Client)

Note: NetBackup only supports a database move of a backup with FileStream enabled if the backup is stream-based.

A database move lets you use a full set of backup images to copy an existing database to a location under a different name. Database move operations can only be carried out when your selection includes a database image. This move can occur either when you directly select the database backup image, or when NetBackup finds a recovery set that contains a database backup image.

To perform a database move

- 1 Open the NetBackup MS SQL Client.
- 2 Browse for the backup images you want to restore.
See [“Browsing for SQL Server backup images \(NetBackup MS SQL Client\)”](#) on page 225.
- 3 Expand the database instance.
- 4 Select the database backup image that you want to restore.
- 5 From the **Scripting** list, select **Create a move template**.
When you create a move script, the capability to perform an immediate launch is disabled. You must edit the script to specify certain destination parameters.
- 6 Click **Restore**.
- 7 Indicate a file name and click **Save > Yes**.

8 Change the database name in the template to the name of the database to restore to.

For example, replace:

```
# Replace the database name in the following line with the name of the database that you
# want to move to. Also remove the hash mark <#> which precedes the keyword <DATABASE>.
#
# DATABASE "DatabaseA"
```

with:

```
# Replace the database name in the following line with the name of the database that you
# want to move to. Also remove the hash mark <#> which precedes the keyword <DATABASE>.
#
DATABASE "DatabaseB"
```

9 Change the path for the database files that you want to restore.

You must uncomment at least one file. For example, replace:

```
# Replace the file path <C:\Microsoft SQL Server\MSSQL.3\MSSQL\DATA\DBA_FG1_File1.ndf>
# with a new file path. Also remove the hash mark <#> which precedes the keyword <TO>.
# The target of the MOVE keyword must be "DBA_FG1_File1".
MOVE "DBA_FG1_File1"
#TO "C:\Microsoft SQL Server\MSSQL.3\MSSQL\DATA\DBA_FG1_File1.ndf"
```

with:

```
# Replace the file path <C:\Microsoft SQL Server\MSSQL.3\MSSQL\DATA\DBA_FG1_File1.ndf>
# with a new file path. Also remove the hash mark <#> which precedes the keyword <TO>.
# The target of the MOVE keyword must be "DBA_FG1_File1".
MOVE "DBA_FG1_File1"
TO "C:\Microsoft SQL Server\MSSQL.3\MSSQL\DATA\DBB_FG1_File1.ndf"
```

10 Change the database file path.

For example, replace:

```
# Replace the file path <C:\Microsoft SQL Server\MSSQL.3\MSSQL\DATA\DatabaseA.mdf>
# with a new file path. Also remove the hash mark <#> which precedes the keyword <TO>.
# The target of the MOVE keyword must be "DatabaseA".
MOVE "DatabaseA"
#TO "C:\Microsoft SQL Server\MSSQL.3\MSSQL\DATA\DatabaseA.mdf"
```

with:

```
# Replace the file path <C:\Microsoft SQL Server\MSSQL.3\MSSQL\DATA\DatabaseA.mdf>
# with a new file path. Also remove the hash mark <#> which precedes the keyword <TO>.
# The target of the MOVE keyword must be "DatabaseA".
MOVE "DatabaseA"
TO "C:\Microsoft SQL Server\MSSQL.3\MSSQL\DATA\DatabaseB.mdf"
```

- 11 Make similar changes to the template for any differential backups or transaction log backups you want to move.
- 12 When you finish modifying the template, save it.
- 13 To run the restore, select **File > Manage script files**.
- 14 Select the script that you created and click **Start > Yes**.

About performing a SQL Server page-level restore (NetBackup MS SQL Client)

Note: Page-level restores are only applicable for batch file-based policies.

Use page-level restore to recover only the pages that are corrupted. If many pages are corrupt, then a full database recovery may be faster.

When you select the page restore option, NetBackup for SQL Server creates a page restore template.

This template includes the following parts:

- A page restore operation that you can modify by inserting the IDs of the pages that you want to restore.
- A series of transaction log images for recovering the database to the current point in time.

- A tail-log backup and recovery operation, which is required to bring the database online.

About SQL page-level restore requirements and limitations

The following requirements and limitations exist when you perform SQL Server page-level restores:

- Pages can be restored from the following backup types: Database, filegroup, file, read-write filegroups, and partial database.
- Your SQL Server must use either the full or bulk-logged recovery model.
- SQL Server sometimes cannot recover the specific pages that you request if they contain critical information about the definition of the database itself. For example, you cannot use page-level restore for the first page in a database file. When you detect that page-level restore does not work, you need to use full database recovery.
- A maximum of 1000 pages can be recovered from a backup image through a page-level restore.

Perform a SQL Server page-level restore (NetBackup MS SQL Client)

This topic describes how to perform page-level restores. Note that the Microsoft SQL Server service must have full access permission to the folder

`install_path\netbackup\dbext\mssql\temp`.

To perform a SQL Server page-level restore

- 1 Open the NetBackup MS SQL Client.
- 2 Obtain a list of corrupt pages in the database.
- 3 Browse for the backup images you want to restore.
 See [“Browsing for SQL Server backup images \(NetBackup MS SQL Client\)”](#) on page 225.
- 4 Expand the database instance and the database.
- 5 Select the database backup image that contains the pages you want to restore.
- 6 From the **Scripting** list, select **Create a page restore template**.
- 7 Click **Restore**.
- 8 Type a file name for the page restore script and click **Save > Yes**.

- 9 Edit the page first operation the page IDs that you want to replace.

For example, replace:

```
#
# Create one or more page restore requests. These use the following format
#PAGE file-id:page-id
```

with

```
#
# Create one or more page restore requests. These use the following format
PAGE 1:14
PAGE 1:20
```

- 10 When you finish modifying the template, save it.
- 11 Select the script you created and click **Start > Yes**.

Redirect a SQL Server database to a different host (NetBackup MS SQL Client)

You can use a database move operation to redirect a backup to a client that is different from the client that performed the backup. NetBackup creates a template that you edit to indicate the host and location where you want to redirect the restore. The new location can be a different instance on the same host, a different host, or a different file path. The move operation also lets you restore the database under a different name than the original one.

Note: The destination host and instance of a move or restore operation is the one that you log into. For move or restore operations designate the source (or browse) host and the instance when you select **File > Restore SQL Server objects**.

To redirect a database to another location on a different host

- 1 Establish permissions for redirected restores on the primary server.
See [“Configuring permissions for redirected restores”](#) on page 161.
- 2 The server that backed up the database you want to restore must appear in the server list of the destination host. If the server is not in the list, add it.
See [“About selecting a primary server ”](#) on page 237.
- 3 Open the NetBackup MS SQL Client.
- 4 Select **File > Set SQL Server connection properties**.

- 5** From the **Host** list, select the host you want to restore to.
- 6** From the **Instance** list, select the database instance.
To select the default instance, either select **<default>** or leave the field empty.
- 7** Click **Apply** and then **Close**.
- 8** Select **File > Set NetBackup client properties**.
- 9** From the **Current NetBackup Server** list, select the NetBackup primary server.
This server contains the SQL Server backup images that you want to restore on the destination host. The clients must both use the same primary server.
See [“About selecting a primary server”](#) on page 237.
- 10** Click **OK**.
- 11** Browse for the backup images you want to restore.
For the **SQL Host** list, select the host that has the database you want to restore.
See [“Browsing for SQL Server backup images \(NetBackup MS SQL Client\)”](#) on page 225.
- 12** Browse for the database that you want to move.
- 13** From the **Scripting** list, select **Create a move template**.
- 14** Click **Restore**.
- 15** Enter a file name and click **Save > Yes**.
- 16** Edit the template to designate the name that you want to use for the destination database. Also include the file paths that you want to use for each of the database files.

About selecting a primary server

When you perform a move, the backup images must be available on the host machine that acts as the NetBackup primary server for the destination host. If this server is contained in the server list of the destination host, click **File > Set NetBackup client properties** to select the current primary server.

If the server is not in the server list of the destination host you must duplicate the images onto removable media (with a unique ID). Then transport that media to the primary server that the destination host uses, and import the images to that server. After the images are imported, continue with the instructions for performing a move. A server may not appear in the server list because the server is remote or has access limitations.

See [“Perform a SQL Server database move \(NetBackup MS SQL Client\)”](#) on page 232.

Perform a restore of a remote SQL Server installation (NetBackup MS SQL Client)

You can use NetBackup for SQL Server to restore databases on a remote host. Generated batch files must be saved on the remote host. You can launch the operation from the local installation of NetBackup for SQL Server.

To perform a restore of a remote SQL Server installation

- 1 Open the NetBackup MS SQL Client.
- 2 Select the host and instance you want to access.

See [“Select the SQL Server host and instance \(NetBackup MS SQL Client\)”](#) on page 216.
- 3 Browse for the backup images you want to restore.

See [“Browsing for SQL Server backup images \(NetBackup MS SQL Client\)”](#) on page 225.
- 4 Select the options for the operation.

See [“Options for NetBackup for SQL Server restores”](#) on page 222.

Save is enabled in the restore dialog box. **Launch immediately** is disabled because the generated script must run on the remote host that you are logged on to.
- 5 Click **Restore**.
- 6 Navigate to the `install_path\NetBackup\DbExt\MsSql\` folder on the remote host, and save the batch file there.
- 7 Run the operation from the local installation of NetBackup for SQL Server.

Restoring multistreamed SQL Server backups

When you perform a restore from a backup that used multiple stripes, NetBackup automatically performs the restore using the same number of stripes. Select the object you want to restore and NetBackup finds all of the related backups and restore them. Upon restore, all of the streams must also be available at the same time.

About conventional backups using multiple streams

If you specified multiple stripes for a non-snapshot backup, then the number of backup streams that you specified were created. NetBackup names these streams, for example:

```
juneberry.MSSQL7.COLE.db.pubs.~.7.001of003.20140908200234..C  
juneberry.MSSQL7.COLE.db.pubs.~.7.002of003.20140908200234..C  
juneberry.MSSQL7.COLE.db.pubs.~.7.003of003.20140908200234..C
```

To create your own batch file to restore a striped object, specify only the first stripe name with the NBIMAGE keyword. NetBackup for SQL Server finds the remaining ones automatically. More information is available about the backup names that are used for SQL Server objects.

See [“About using bplist to retrieve SQL Server backups”](#) on page 241.

About snapshot backup methods using multiple streams

If you specified multiple stripes for any Snapshot Client backup, then NetBackup divides the number of component files equally among the number of stripes. (A Snapshot Client backup streams the frozen image to tape.) If the number of files is less than the specified number of stripes, then the agent performs the backup using only as many stripes as there are files.

Note: NetBackup ignores the multistream directive for Instant Recovery backups.

With SQL Server backups that are performed with Snapshot Client, NetBackup identifies all of the backup streams by the same name. They are differentiated by NetBackup by their backup IDs.

```
juneberry.MSSQL7.COLE.db.Northwind.~.7.001of003.20141012131132..C
```

Restoring a multistreamed SQL Server backup with fewer devices than it was backed up with

In your recovery environment, you may have fewer drives available for restores than you used for backups. In this situation, SQL Server times out while it waits for the additional backup images to be mounted. To prevent this time out, modify the number of drives that are available for restore.

Consider, for example, if you had performed a backup using 5 drives, and only 2 are available for recovery. In the policy, change the **Stripes** value from **5** to **2**. If you use batch files, change the stripes parameter from `STRIPES 5` to `STRIPES 2`. This change causes SQL Server to request two backup images at a time until all five images are restored.

About conventional backups using multiple streams

If you specified multiple stripes for a non-snapshot backup, then the number of backup streams that you specified was created. NetBackup names these streams, for example:

```
juneberry.MSSQL7.COLE.db.pubs.~.7.001of003.20140908200234..C  
juneberry.MSSQL7.COLE.db.pubs.~.7.002of003.20140908200234..C  
juneberry.MSSQL7.COLE.db.pubs.~.7.003of003.20140908200234..C
```

To create your own batch file to restore a striped object, specify only the first stripe name with the NBIMAGE keyword. NetBackup for SQL Server finds the remaining ones automatically. More information is available about the backup names that are used for SQL Server objects.

See [“About using bplist to retrieve SQL Server backups”](#) on page 241.

About snapshot backup methods using multiple streams

If you specified multiple stripes for any Snapshot Client backup, which streams the frozen image to tape, then NetBackup divides the number of component files equally among the number of stripes. If the number of files is less than the specified number of stripes, then the agent performs the backup using only as many stripes as there are files.

Note: NetBackup ignores the multistream directive for Instant Recovery backups.

With SQL Server backups performed with Snapshot Client, NetBackup identifies all of the backup streams by the same name. They are differentiated by NetBackup by their backup IDs.

```
juneberry.MSSQL7.COLE.db.Northwind.~.7.001of003.20141012131132..C
```

Restoring a multistreamed SQL Server backup with fewer devices than it was backed up with

In your recovery environment, you may have fewer drives available for restores than you used for backups. In this situation, SQL Server times out while it waits for the additional backup images to be mounted. To prevent this time out, modify the recovery batch file to specify the number of drives that are available for restore.

Consider, for example, if you had performed a backup using 5 drives, and only 2 are available for recovery. In the recovery batch file, change the stripes parameter from `STRIPES 5` to `STRIPES 2`. This change causes SQL Server to request two backup images at a time until all five images are restored.

About using bplist to retrieve SQL Server backups

You can use the `bplist` command to obtain restore images. Use this command if you plan to manually create a restore script, rather than through the NetBackup for SQL Server interface. See the [NetBackup Commands Reference Guide](#) for complete information about `bplist`.

To extract all of the NetBackup for SQL Server backups from a specific server for a specific client, run the following command from the Windows command prompt.

```
install_path\NetBackup\bin\bplist -C client -t 15 -S server -R \
```

where *client* is the host machine on which NetBackup for SQL Server resides and *server* is the host machine of NetBackup server.

The following example shows how to obtain the list of SQL Server backups that were backed up from client *juneberry* to server *Cole*:

```
C:\Program Files\NetBackup\bin\bplist -C juneberry -t 15 -S cole -R \
juneberry.MSSQL7.JUNEBERRY.db.pubs.~.7.001of003.20140920101716..C:\
juneberry.MSSQL7.JUNEBERRY.db.pubs.~.7.002of003.20140920101716..C:\
juneberry.MSSQL7.JUNEBERRY.db.pubs.~.7.003of003.20140920101716..C:\
juneberry.MSSQL7.JUNEBERRY.fil.pubs.pubsnew.7.001of001.20140919175149..C:\
juneberry.MSSQL7.JUNEBERRY\NEWINSTANCE.trx.abc.~.7.001of001.20140902170920..C:\
juneberry.MSSQL7.JUNEBERRY\NEWINSTANCE.fg.abc.PRIMARY.7.001of001.20140902170824.C:\
juneberry.MSSQL7.JUNEBERRY\NEWINSTANCE.db.Howard's
Barbeque.~.7.001of001.20140901085255..C:\
juneberry.MSSQL7.JUNEBERRY\NEWINSTANCE.inc.Howard's
Barbeque.~.7.001of001.20140903108552..C:\
juneberry.MSSQL7.COLE.db.pubs.~.7.001of001.20140907100101..C:\
juneberry.MSSQL7.COLE.db.pubs.~.7.001of001.20140908200234..C:\
```

Note: The colon and backslash that terminate each line are not part of the backup name.

See [“About NetBackup for SQL Server backup names”](#) on page 241.

About NetBackup for SQL Server backup names

The backup name is a string that consists of the following components. These components are separated by a delimiter that is specified by the character that precedes the “C” at the end of the backup image name. Backup images for standalone instance databases or read-scale availability groups include the host

and the instance name. Backup images for advanced and basic availability groups include the cluster name, availability group node name, and availability group name.

Figure 15-1 Backup image name for a database filegroup





Figure 15-2 Backup image name for availability database



1 - Named instances are formatted as *<host>\<instance-name>*. The default instance is the name of the host machine.

2 - The object types are as follows:

db	database
inc	database differential
trx	transaction log
fg	filegroup
fgd	filegroup differential
fil	file

- 3 - The name of the file or filegroup if the object type is a file or filegroup; otherwise the symbol ~ is used.
- 4 - The block size.
- 5 - Stripes are specified as *<stripe number>of<total stripes>*. Non-striped backups are always 001of001. For striped backups, *<total stripes>* is the total number of stripes for the backup. *<stripe number>* is the count number of the backup for that backup, starting with 001.
- 6 - The format of the timestamp is YYYYMMDDHHMMSS. The timestamp for availability group backup images reflects Coordinated Universal Time (UTC). For standard database backup images, the timestamp reflects the time zone that is configured for the NetBackup server.
- 7 - The delimiter, which immediately precedes the version indicator. For standard database images, this character is a period (.) by default. For availability database images, the character is a forward slash (/). However, if a period or slash is used in any of the fields, the delimiter may be another character.
- 8 - "C" is applied to all SQL Server backup image names, regardless of the NetBackup version.
- 9 - Backup images for AG databases are formatted as
<WindowsServerFailoverCluster>.<nodename>/MSSQL7/<AGname>.

Using NetBackup for SQL Server with multiple NICs

This chapter includes the following topics:

- [Configuration and requirements for SQL Server backups with multiple NICs](#)
- [Configure the NetBackup client with the private interface name](#)
- [Configure backups of SQL Server when you have multiple NICs \(SQL Server Intelligent Policies\)](#)
- [Configure backups for SQL Server when you have multiple NICs \(batch file-based policies\)](#)
- [Restore SQL Server when you have multiple NICs \(NetBackup MS SQL Client\)](#)
- [Configure backups of a SQL Server cluster when you have multiple NICs \(SQL Server Intelligent Policies\)](#)
- [Configure backups of a SQL Server cluster when you have multiple NICs \(batch file-based policies\)](#)
- [Create a batch file for backups of a SQL Server cluster when you have multiple NICs \(batch file-based policies\)](#)
- [Restore a SQL Server cluster when you have multiple NICs \(NetBackup MS SQL Client\)](#)

Configuration and requirements for SQL Server backups with multiple NICs

Many administrators want to reserve a separate network interface for their SQL Server host machines that are used for routing backup traffic. This type of environment requires additional configuration for backup policies and the NetBackup client that backs up SQL Server. Special configuration is also required to perform restores.

Note: If you have a SQL Server cluster in a private network, you must configure the mappings for distributed application restores. Also review the auto-discovered mappings for the hosts in your environment.

See [“Configuring mappings for restores of a distributed applications, clusters, or virtual machines”](#) on page 31.

See [“Reviewing the auto-discovered mappings”](#) on page 26.

Ask your NetBackup administrator for assistance.

The following distinct network resources exist in a multi-NIC environment:

- The public name of each SQL Server host (for example, `sqlhost1` and `sqlhost2`)
- The private interface name that is used to back up each of the SQL Server hosts (for example, `sqlhost1-NB` and `sqlhost2-NB`)

The following additional resources exist for a SQL Server cluster in a multi-NIC environment:

- The public virtual name of the SQL Server (for example, `virtsql`)
- The private virtual name of the SQL Server (for example, `virtsql-NB`)

The following requirements exist to use NetBackup for SQL Server in a multi-NIC environment:

- Install the NetBackup client on the SQL Server using the private name of the SQL Server host as the NetBackup client name.
Alternatively, you can configure the NetBackup client name after installation.
See [“Configure the NetBackup client with the private interface name”](#) on page 246.
- For intelligent policies, configure a backup policy that includes the private interface name of the host or client.
See [“Configure backups of SQL Server when you have multiple NICs \(SQL Server Intelligent Policies\)”](#) on page 247.
See [“Configure backups of a SQL Server cluster when you have multiple NICs \(SQL Server Intelligent Policies\)”](#) on page 250.

- For batch file-based policies, configure a backup policy that includes the private interface name of the host or client.
See [“Configure backups for SQL Server when you have multiple NICs \(batch file-based policies\)”](#) on page 248.
See [“Configure backups of a SQL Server cluster when you have multiple NICs \(batch file-based policies\)”](#) on page 251.
Note that if you want to protect a SQL Server cluster with a batch file-based policy, you must edit the backup batch file. The `BROWSECLIENT` parameter must indicate the private name of the SQL Server host or virtual SQL Server.
- Configure permissions to allow all nodes in the cluster to browse for backups across the private interface (redirected restores). The administrator can allow all clients or allow single clients to browse and restore a backup that is performed over the multi-NIC connection.
See [“Configuring permissions for redirected restores”](#) on page 161.

Configure the NetBackup client with the private interface name

To perform backups over a private network interface, NetBackup must use the private name of the client. If you installed the NetBackup client using the public interface name, follow this procedure to configure the NetBackup client name as the private interface name.

For cluster environments, additional configuration is required. In that case, NetBackup must use the private virtual name of the SQL Server cluster.

See [“Configure backups of clustered SQL Server instances \(batch file-based policies\)”](#) on page 133.

To configure the NetBackup client with the private interface name

- 1 Open the Backup, Archive, and Restore interface.
- 2 Select **File > NetBackup Client Properties**.
- 3 Select the **General** tab.
- 4 In the **Client name** box, specify the private name of the client.

For example, the private name for the computer `sqlhost1` is `sqlhost1-NB`.

Configure backups of SQL Server when you have multiple NICs (SQL Server Intelligent Policies)

This topic describes how to create a SQL Server Intelligent Policy to protect a SQL Server when you have multiple NICs. The following configuration changes must be made to allow for backups and restores over a private interface:

- Install the NetBackup client on the SQL Server using the private name of the SQL Server host as the NetBackup client name.
 Alternatively, you can configure the NetBackup client name after installation. See [“Configure the NetBackup client with the private interface name”](#) on page 246.
- The backup policy must include the private interface name of the SQL Server host.
 During instance discovery NetBackup automatically adds an instance with the NetBackup client name. If you installed the NetBackup client using the private interface name, NetBackup uses the private name when it performs backups.

To configure a backup policy for a SQL Server in a cluster with a multi-NIC (SQL Server Intelligent Policies)

- 1 If you installed the NetBackup client on the SQL Server host using the public interface name, follow the procedure to configure the NetBackup client name as the private interface name.
 See [“Configure the NetBackup client with the private interface name”](#) on page 246.
- 2 Sign in to the primary server as a user that has the RBAC Administrator role or a role that can manage policies.
- 3 Find and register the instance that has the private interface name of the SQL Server host (`sqlhost1-NB`).
- 4 On the left, select **Protection > Policies**.
 Create a new policy or open an existing policy.
- 5 On the **Instances and databases** tab, select **Protect availability groups** or **Protect intelligent groups**.
- 6 Select **Add**.

- 7 To add the instances or databases that you want to protect, select or expand the instance that has the private interface name of the SQL Server (`sqlhost1-NB`).
 See [“Add instances to a policy”](#) on page 80.
 See [“Add databases to a policy”](#) on page 81.
- 8 Add other policy information as follows:
 - Add schedules.
 See [“Schedule properties for SQL Server Intelligent Policies”](#) on page 77.
 - Add database objects to the backup selections list.
 See [“Add filegroups or files to the backup selections list”](#) on page 83.
 - (Optional) Make changes to any tuning parameters.
 See [“Performance tuning and configuration options ”](#) on page 84.

Configure backups for SQL Server when you have multiple NICs (batch file-based policies)

This topic describes how to configure a batch file-based policy using batch files to protect SQL Server with a multi-NIC. The following configuration changes must be made to allow for backups and restores over a private interface:

- Install the NetBackup client on the SQL Server using the private name of the SQL Server host as the NetBackup client name.
 Alternatively, you can configure the NetBackup client name after installation.
 See [“Configure the NetBackup client with the private interface name”](#) on page 246.
- The backup policy must include the private interface name of the SQL Server host.

To configure backups for SQL Server when you have multiple NICs (batch file-based policies)

- 1 If you installed the NetBackup client on the SQL Server host using the public interface name, follow the procedure to configure the NetBackup client name as the private interface name.
 See [“Configure the NetBackup client with the private interface name”](#) on page 246.
- 2 Sign in to the primary server as a user that has the RBAC Administrator role or a role that can manage policies.

- 3 On the left, select **Protection > Policies**.
Create a new policy or open an existing policy.
- 4 On the **Clients** tab, add a new client.
For the Client name, provide the private interface name. For example, the public name is `sqlhost1`. The private interface that is used to back up `sqlhost1` is `sqlhost1-NB`.
- 5 Create a batch file that includes the private interface name of the virtual SQL Server. Then add this batch file to the backup selections list.

See [“Create a batch file for backups of a SQL Server cluster when you have multiple NICs \(batch file-based policies\)”](#) on page 252.

See [“Add batch files to the backup selections list ”](#) on page 206.
- 6 Add other policy information as follows:
 - Add schedules.
See [“Schedule properties for SQL Server batch file-based policies ”](#) on page 201.

Restore SQL Server when you have multiple NICs (NetBackup MS SQL Client)

To restore a SQL Server in a multi-NIC environment with the NetBackup MS SQL Client, you need to do the following:

- Connect to SQL Server host using the public name of the host.
- To browse for backup images, specify the public name of the SQL Server for the **SQL Host** name. Specify the private name of the SQL Server for the **Source Client**.

If you use SQL Server policies in a cluster environment, you must follow a different procedure:

See [“Restore a SQL Server cluster when you have multiple NICs \(NetBackup MS SQL Client\)”](#) on page 253.

To restore SQL Server when you have multiple NICs

- 1 Open the NetBackup MS SQL Client.
- 2 Select **File > Set SQL Server connection properties**.
- 3 In the **Host** box, specify the public name of the SQL Server host.
- 4 Select **OK**.

Configure backups of a SQL Server cluster when you have multiple NICs (SQL Server Intelligent Policies)

- 5 Select **File > Restore SQL Server objects**.
- 6 In the **SQL Host** box, specify the public name of the SQL Server host (`sqlhost1`).
- 7 In the **Source Client** box, specify the private interface name of the SQL Server host (`sqlhost1-NB`).
- 8 Select **OK**.
A dialog box opens that shows the SQL Server backups that the **SQL Host** made on the private network interface.
- 9 Continue with the restore as normal.
See [“Restore a SQL Server database backup \(NetBackup MS SQL Client\)”](#) on page 227.

Configure backups of a SQL Server cluster when you have multiple NICs (SQL Server Intelligent Policies)

This topic describes how to create a SQL Server Intelligent Policy to protect a SQL Server when you have multiple NICs. During instance discovery NetBackup automatically adds an instance with the NetBackup client name. For a virtual SQL Server in a multi-NIC environment, you must add and register the instance with the private interface name of the virtual SQL Server. This name is the instance name that you add to the backup policy.

To configure a backup policy for a SQL Server cluster with a multi-NIC (SQL Server Intelligent Policies)

- 1 Sign in to the primary server as a user that has the RBAC Administrator role or a role that can manage policies.
- 2 Manually add a new instance and register it. For the **Host**, provide the private interface name of the virtual SQL Server (`virtsql-NB`).
See [“Manually add a SQL Server instance”](#) on page 71.
- 3 On the left, select **Protection > Policies**.
Create a new policy or open an existing policy.
- 4 On the **Instances and databases** tab, select **Protect availability groups** or **Protect intelligent groups**.
- 5 Select **Add**.

- 6 To add the instances or databases that you want to protect, select or expand the instance that has the private interface name of the virtual SQL Server (`virtsql-NB`). Or, select the intelligent groups that you want to protect.
 See [“Add instances to a policy”](#) on page 80.
 See [“Add databases to a policy”](#) on page 81.
 See [“Add intelligent groups to a policy”](#) on page 82.
- 7 Add other policy information as follows:
 - Add schedules.
 See [“Schedule properties for SQL Server Intelligent Policies”](#) on page 77.
 - Add database objects to the backup selections list. Filegroups and files cannot be selected for policies that use intelligent groups.
 See [“Add filegroups or files to the backup selections list”](#) on page 83.
 - (Optional) Make changes to any tuning parameters.
 See [“Performance tuning and configuration options”](#) on page 84.

Configure backups of a SQL Server cluster when you have multiple NICs (batch file-based policies)

This topic describes how to create a batch file-based policy to protect a SQL Server cluster with a multi-NIC. When you create the backup policy, it must include a client that has the private interface name of the virtual SQL Server. The public name of the host should not be used.

To configure backups of a SQL Server when you have multiple NICs (batch file-based policies)

- 1 Sign in to the primary server as a user that has the RBAC Administrator role or a role that can manage policies.
- 2 On the left, select **Protection > Policies**.
 Create a new policy or open an existing policy.
- 3 On the **Clients** tab, add a new client.
 For the client name, use the private interface name of the virtual SQL Server. For example, `virtsql-NB`.

Create a batch file for backups of a SQL Server cluster when you have multiple NICs (batch file-based policies)

- 4 Create a batch file that includes the private interface name of the virtual SQL Server. Then add this batch file to the backup selections list.

See [“Create a batch file for backups of a SQL Server cluster when you have multiple NICs \(batch file-based policies\)”](#) on page 252.

See [“Add batch files to the backup selections list ”](#) on page 206.

- 5 Add other policy information as follows:

- Add schedules.

See [“Schedule properties for SQL Server batch file-based policies ”](#) on page 201.

Create a batch file for backups of a SQL Server cluster when you have multiple NICs (batch file-based policies)

This topic describes how to create a batch file for a batch file-based policy to protect a SQL Server cluster with a multi-NIC connection. To create the batch file you need to connect to the SQL Server host using the public name of the virtual SQL Server. The batch file must include the private name of the virtual SQL Server.

To create a batch file for SQL Server cluster backups with a multi-NIC connection

- 1 On any node in the SQL Server cluster, open the NetBackup for SQL Server interface.
- 2 Select **File > Set SQL Server connection properties**.
- 3 In the **Host** box, specify the public name of the virtual SQL Server host (virtsql).
- 4 Select **Apply** and **Close**.
- 5 Select **File > Backup SQL Server objects**.
- 6 Select the databases to back up.
- 7 Select the backup options.

See [“Options for SQL Server backup operations”](#) on page 207.

Note: Do not attempt to perform an immediate backup from the backup dialog box. The generated batch files must be modified before they can be run successfully.

8 From the **Backup script** options, select **Save**.

9 Select **Backup**.

A batch file similar to the following is created:

```
OPERATION BACKUP
DATABASE "ACCOUNTING"
SQLHOST "VIRTSQL"
NBSERVER "THOR"
BROWSECLIENT "VIRTSQL"
MAXTRANSFERSIZE 0
BLOCKSIZE 7
ENDOPER TRUE
```

10 Change the line value that is associated with the `BROWSECLIENT` from the public name of the virtual SQL Server to the private name.

```
OPERATION BACK
UPDATABASE "ACCOUNTING"
SQLHOST "VIRTSQL"
NBSERVER "THOR"
BROWSECLIENT "VIRTSQL-NB"
MAXTRANSFERSIZE 0
BLOCKSIZE 7
ENDOPER TRUE
```

11 Place the modified batch file on all nodes in the cluster or in a shared location. This way it is available for scheduled backups.

Backups are done regardless of which node is active when a backup is initiated.

Restore a SQL Server cluster when you have multiple NICs (NetBackup MS SQL Client)

To restore a SQL Server cluster in a multi-NIC environment with the NetBackup MS SQL Client, you need to do the following:

- Connect to virtual SQL Server host using the public name of the host.
- To browse for backup images, specify the public name of the virtual SQL Server for the **SQL Host** name. Specify the private name of the virtual SQL Server for the **Source Client**.
- Create a batch file for the restore and change the `BROWSECLIENT` parameter to indicate the private name of the virtual SQL Server.

For a non-cluster environment, you must follow a different procedure:

See [“Restore SQL Server when you have multiple NICs \(NetBackup MS SQL Client\)”](#) on page 249.

To restore a SQL Server cluster when you have multiple NICs

- 1 On a specific node in the cluster, open the NetBackup MS SQL Client.
- 2 Select **File > Set SQL Server connection properties**.
- 3 In the **Host** box, specify the public name of the virtual SQL Server host (virtsql).
- 4 Select **Apply** and **Close**.
- 5 Select **File > Restore SQL Server objects**.
- 6 In the **Backup History Options** dialog box, specify the following.

SQL Host	Public name of the virtual SQL Server (virtsql).
Source Client	Private name of the virtual SQL Server (virtsql-NB).

- 7 Select **OK**.
- 8 Select the databases to restore.
See [“Options for NetBackup for SQL Server restores”](#) on page 222.

Note: Do not try to perform an immediate restore from the restore dialog box. The generated batch files must be modified before they can be run successfully.

- 9 Select the restore options.
- 10 From the **Restore script options**, select **Save**.

11 Select **Restore**.

The NetBackup MS SQL Client generates a batch file that is similar to the following.

```
OPERATION RESTORE
OBJECTTYPE DATABASE
DATABASE "ACCOUNTING"
NBIMAGE "SQLHOST1.MSSQL7.VIRTSQL.db.ACCOUNTING.~.7.001of001.20040306111309..C"
SQLHOST "VIRTSQL"
NBSERVER "THOR"
BROWSECLIENT "VIRTSQL"
MAXTRANSFERSIZE 0
BLOCKSIZE 7
RESTOREOPTION REPLACE
RECOVEREDSTATE RECOVERED
ENDOPER TRUE
```

12 Change the line value that is associated with `BROWSECLIENT` from the public name of the virtual SQL Server to the private name.

```
OPERATION RESTORE
OBJECTTYPE DATABASE
DATABASE "ACCOUNTING"
NBIMAGE "SQLHOST1.MSSQL7.VIRTSQL.db.ACCOUNTING.~.7.001of001.20040306111309..C"
SQLHOST "VIRTSQL"
NBSERVER "THOR"
BROWSECLIENT "VIRTSQL-NB"
MAXTRANSFERSIZE 0
BLOCKSIZE 7
RESTOREOPTION REPLACE
RECOVEREDSTATE RECOVERED
ENDOPER TRUE
```

13 Select **File > Manage script files**.

14 Select the modified batch file and select **Start**.

Performance and troubleshooting

This chapter includes the following topics:

- [NetBackup for SQL Server performance factors](#)
- [About monitoring NetBackup for SQL Server operations](#)
- [About debug logging for SQL Server troubleshooting](#)
- [Troubleshooting credential validation](#)
- [Troubleshooting VMware backups](#)
- [SQL Server log truncation failure during VMware backups of SQL Server](#)
- [Intelligent group backup completes with a status code 1 when database not found in the SQL Server environment](#)
- [Set the maximum trace level for NetBackup for SQL Server](#)
- [Reporting of unsuccessful filegroup or file backups](#)
- [About minimizing timeout failures on large SQL Server database restores](#)
- [SQL Server restore fails when you restore a SQL Server compressed backup image as a single stripe or with multiple stripes](#)
- [Incorrect backup images are displayed for availability group clusters](#)
- [A restore of a SQL Server database fails with Status Code 5, or Error \(-1\), when the host name of the SQL Server or the SQL Server database name has trailing spaces](#)

- A move operation fails with Status Code 5, or Error (-1), when the SQL Server host name, the database name, or the database logical name has trailing spaces
- Unable to discover or browse availability group replicas
- SQL Server policy backups that use intelligent groups may fail with a status code 200 when the associated credentials in a credential rule are invalid
- About disaster recovery 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 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 36.

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 batch file-based 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.

Using read-write and read-only filegroups

You can significantly reduce backup time and the storage media that is needed if you periodically back up only the read-write filegroups. Then keep a single backup of read-only filegroups, which is retained infinitely. You can set the retention level in the schedule.

About monitoring NetBackup for SQL Server operations

Use the Activity monitor to monitor NetBackup for SQL Server operations.

The agent also creates its own progress reports that you can view in the NetBackup MS SQL Client interface. Select **File > View status** to view the reports. The reports are saved in `install_path\NetBackup\logs\user_ops\MsSql\logs`.

Job details and progress reports include the following types of information:

- Summary information about the operation
- Information about the operation as it progresses
- Any error conditions or warnings that cause the operation to fail
- The final outcome of the operation, whether it succeeded or failed, and how long it took

The progress reports also provide additional details for operations, including the following:

- The SQL Server commands that NetBackup included in the batch file for operation.

```
OPERATION BACKUP
DATABASE "TestDB1"
OBJECTTYPE DATABASE
COPYONLY FALSE
BLOCKSIZE 7
MAXTRANSFERSIZE 6
NUMBUFS 2
STRIPES 1
SQLCOMPRESSION FALSE
VERIFYOPTION NONE
```

- The NetBackup server that performed the backup, the SQL Server instance and host you selected for the backup, and other policy information.

```
NBSERVER "servera"
SQLINSTANCE "SQL2K14"
SQLHOST "SERVERA"
POLICY "sql-server"
NBSCHED "full"
INF - Setting backup catalog name to: servera
```

- Progress of the backup or restore operation and any errors or failures that SQL Server encountered.

```
USER - Operation inhibited by NetBackup for Microsoft SQL
Server: Only a full or incremental database backup can be performed
on database <Archive> because it uses the simple recovery model or
has 'truncate log on checkpoint' set.
INF - OPERATION #1 of batch
C:\NBU\Veritas\NetBackup\dbext\mssql\temp\__01_35_42_508_00.bch
FAILED with STATUS 1 (0 is normal). Elapsed time = 6(6) seconds.
INF - Results of executing
<C:\Program Files\Cohesity NetBackup\NetBackup\dbext\mssql\temp\__01_35_42_
<0> operations succeeded. <1> operations failed.
INF - The following object(s) were not backed up successfully.
INF - Archive
```

About debug logging for SQL Server troubleshooting

NetBackup offers a comprehensive set of debug logs for troubleshooting issues that can occur during NetBackup operations. You can create individual logs or use a script to create all NetBackup debug logs. For details on the contents of these debug logs, see the [NetBackup Troubleshooting Guide](#).

Backup operation debug logs

After you perform a backup, debug logging information is placed in the *install_path\NetBackup\logs* directory. A subdirectory is created for each process. The debug log file is named `ALL_ADMINS.mmdyy_0000x.log`. For Veritas Unified Logging (VxUL), the log file is in a format that is standardized across Cohesity products.

Client	Refer to the following logs: ■ bphdb (scheduled backups only) ■ dbclient ■ ncfncbs (VxUL) ■ nbdisco (VxUL) ■ user_ops\mssql\logs
Primary server	nbars (VxUL)
Snapshot backups	Refer to the following logs: ■ bpbkar (Snapshot Client) ■ nbfsd (Snapshot Client) ■ bppfi Instant Recovery
VMware backups	For ASC issues and failures, the following logs are created on the VM that is backed up: ■ bpbkar ■ dbclient ■ ncfncbs (VxUL)

Restore operation debug logs

The following logs apply to restore operations.

Client	Refer to the following logs: ■ bpbkar (Snapshot Client) ■ bpfis (Snapshot Client) ■ bppfi (Instant Recovery) ■ dbclient ■ user_ops\mssql\logs
VMware restores from snapshots using Replication Director	See the Cohesity VSS provider logs. See “VSS provider logs” on page 263.

Create all debug logs

To create all debug logs

- ◆ Run the following batch file:

```
install_path\NetBackup\logs\mklogdir.bat
```

Set the debug level on a SQL Server client

Information is also available about the **Client Trace Level**.

See [“Set the maximum trace level for NetBackup for SQL Server”](#) on page 268.

To set the debug level

- 1 Open the **Backup, Archive, and Restore** interface.
- 2 Select **File > NetBackup Client Properties**.
- 3 Click the **Troubleshooting** tab.
- 4 Set the **General** debug level.
- 5 Set the **Verbose** debug level.
- 6 Set the **Database** debug level.
- 7 Click **OK** to save your changes.

VSS provider logs

The VSS provider records its activities in Windows Event Logs. Debug logs are also available at the following location:

```
install_path\Veritas VSS provider\logs
```

Enabling VSS provider logging in the registry

Enable the VSS provider logging on the NetBackup computer where SQL Server is installed.

To enable VSS provider logging in the registry

- 1 Log on as administrator on the computer where NetBackup is installed.
- 2 Open Regedit.
- 3 Open the following key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Veritas\NetBackup\CurrentVersion\Config
```

- 4 Create a new DWORD value named **CreateDebugLog**.

- 5 Right-click on the new value and click **Modify**.
- 6 In the **Value data** box, enter **1**.
- 7 Click **OK**.

Increasing the VSS provider log debug level

To increase the log debug level modify both the pre-freeze-script.bat and post-thaw-script.bat files in the C:\Windows folder. Add the `-log` parameter to the script, at the line where `BeVssRequestor.exe` is called. VMware determines which script is invoked.

To increase the VSS provider log debug level

- 1 Change the following line in the pre-freeze-script.bat:

```
BeVssRequestor.exe -pre2 -logscreen !SkipExReplica! !SkipSQL!  
!VMBackupType! !ExcludeList!
```

to:

```
BeVssRequestor.exe -pre2 -logscreen !SkipExReplica! !SkipSQL!  
!VMBackupType! !ExcludeList! -log
```

- 2 Also change the following line in the post-thaw-script.bat:

```
BeVssRequestor.exe -post2 -logscreen !SkipExReplica! !SkipSQL!  
!VMBackupType! !ExcludeList!
```

to:

```
BeVssRequestor.exe -post2 -logscreen !SkipExReplica! !SkipSQL!  
!VMBackupType! !ExcludeList! -log
```

Troubleshooting credential validation

[Table 17-1](#) describes the reasons that validation can fail for an instance, replica, or instance group. You can also select the bell icon to go to the **Events** page. This page shows any audit events or errors that occurred with credentials rules or credentials.

Table 17-1 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.

Table 17-1 Reasons for credential validation failure (*continued*)

Status code or error	Description	Explanation
46	The validation operation timed out waiting for a response from the client	You cannot connect to the host because the host is down.
41	Validation of operating system user/password failed for client: <client>.	- The host name is correct, but the username 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 25.
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 “Configure the NetBackup services for SQL Server backups and restores” on page 24.

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 the 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 17-2 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 111.
	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 Enable 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.
	If Enable T-SQL snapshots is enabled, ensure that the SQL Server VSS Writer service is disabled on the guest virtual machine (SQL Server client) to prevent the ASC job failure.
ASC job fails with status code 142.	The NetBackup version on the primary server, media server, and the client must be at version 10.4 to support T-SQL snapshot backups. Legacy VMWare-ASC backups are supported for back-level versions. The ASC job may fail with status code 142 if you attempt T-SQL snapshot backups on back-level NetBackup versions.

Table 17-2 Issues with using a VMware policy to protect databases
(continued)

Issue	Explanation
You can recover the entire virtual machine from the backup, but you cannot recover the databases individually.	You did not select Application protection option Microsoft SQL Server on the VMware tab in the policy, which allows recovery of the databases from the virtual machine backups.
Transaction log backups fail.	You must first perform a full VMware backup without log truncation (Truncate logs option).
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.
Unable to recover from SQL Server agent differential backup.	If you enabled the Enable T-SQL backups option for VMware backups and the backup failed, NetBackup is not able to inform SQL Server that the backup failed. The next differential backup becomes invalid because there is no full backup on which to base the incremental backup. This issue is resolved after the next full backup is successful.

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.

Intelligent group backup completes with a status code 1 when database not found in the SQL Server environment

NetBackup maintains a list of SQL Server databases in its asset database. Consider a situation where a SQL Server database exists in the NetBackup asset database but not in the live SQL Server environment. In that case, when NetBackup performs an intelligent group backup that contains that database, logs a message, and reports a status code 1:

```
08:31:05.561 [5736.15584] <2> debuglog: <16> bphdb do_script:
ERR - exit status: 5465 Could not find database objects during backup.
```

```
08:31:05.561 [5736.15584] <2> debuglog: <16> bphdb do_script:
INF - Partial success status so switch status from <5465> to <1>
```

On the **Databases** tab, NetBackup indicates the **Database state** as **Deleted** for that database. To resolve this issue you can do one of the following:

- Review the intelligent group in the policy and adjust the query.
- Run the backup job again. NetBackup then marks the database as inactive.
- If the database is now available, perform discovery.

Set the maximum trace level for NetBackup for SQL Server

Note: For SQL Server backups, this feature is only available with batch file-based policies.

You can set the maximum trace level in the NetBackup MS SQL Client or in the batch file. The maximum level produces large amounts of output, usually appropriate only for internal debugging.

To set the maximum trace level in the NetBackup MS SQL Client

- 1 Open the NetBackup MS SQL Client.
- 2 Select **File > Set NetBackup client properties**.
- 3 In the **Client Trace Level** group, select **Maximum**.

To set the maximum trace level in the backup or restore batch file

- 1 Open the NetBackup MS SQL Client.
- 2 Select **File > Manage script files**.
- 3 Select the batch file you want to change and click **Open File**.
- 4 Add the following line:
`TRACELEVEL MAX`
- 5 Save the file.

See [“About minimizing timeout failures on large SQL Server database restores”](#) on page 269.

Reporting of unsuccessful filegroup or file backups

If you select specific databases and specific filegroups or files in a backup policy, NetBackup reports any unsuccessful filegroup or file backups differently than if you select an entire instance (`DATABASE $ALL`). Consider the following scenarios:

- **Scenario 1** - For `SQLINSTANCE1` (`DATABASE $ALL` or all the databases), back up the filegroups `FG1`, `FG2`, and `FG3`. If NetBackup cannot back up `FG1`, `FG2`, or `FG3`, NetBackup skips the backup of the filegroup for that database. The parent job completes with a status 0.
- **Scenario 2** - For `DATABASEA` and `DATABASEC` in `SQLINSTANCE1`, back up the filegroups `FG1`, `FG2`, and `FG3`. If NetBackup cannot back up any of these filegroups for `DATABASEA` or `DATABASEC`, the parent job completes with a status 2. The job details indicate that one or more of the filegroups that you selected were not backed up.

About minimizing timeout failures on large SQL Server database restores

A large SQL Server restore may fail with a Client Read Timeout error before any data has been read from the NetBackup media. This error occurs because the SQL Server may need to pre-write the database files before the restore operation begins. The time that is required for this process is a function of certain factors: the size of the database files and the speed at which your host machine can write to disk. For example, consider that your system can perform disk writes at the rate of 60 megabytes per second and you have a 2.4 terabyte database. Then it takes at least

SQL Server restore fails when you restore a SQL Server compressed backup image as a single stripe or with multiple stripes

12 hours for SQL Server to prep the disk before the actual restore can begin. In reality, the delay may be even longer than what you calculate by as much as 20% to 40%.

The timeout problem can be resolved by increasing the NetBackup **Client read timeout** setting. In the client host properties, change the properties of each client that contains a database that you may need to restore. The default for the **Client read timeout** setting is 300 seconds (5 minutes). If you have any clients which contain large SQL Server databases, you may need to set this value much higher.

You can eliminate file initialization during SQL Server restores. See the following topic:

See “[NetBackup for SQL Server performance factors](#)” on page 257.

SQL Server restore fails when you restore a SQL Server compressed backup image as a single stripe or with multiple stripes

This issue occurs when SQL Server is busy with the buffer of compressed data and cannot process all the data that is sent within a certain length of time. By default in Windows Server, TCP connections must close after the TCP connection state has been set to FIN_WAIT_2 for two minutes. Refer to the following Microsoft article for more information:

<https://support.microsoft.com/en-us/kb/923200/>

Note: If the **TCPFinWait2Delay** value does not exist, you must create it as a REG_DWORD registry value. Otherwise, Windows uses the default value of **240**.

To increase the time that TCP connections may remain in the FIN_WAIT_2 state

- 1 On the NetBackup media server, open `regedit.exe`.
- 2 Locate and select the following registry subkey:
`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters`
- 3 Double-click on **TCPFinWait2Delay**.
- 4 Enter a value of **300**.
- 5 Restart the media server.

- 6 After the restore completes successfully, remove the registry setting or change the setting to its original value.

When you increase the value of this setting it has an adverse effect for all TCP/IP connections. This higher value could cause port exhaustion for other applications that run on the media server.

- 7 Restart the media server.

Incorrect backup images are displayed for availability group clusters

You can perform backups of multiple availability group clusters that have the same short cluster name but that exist in different domains. However, it is important to use the fully qualified domain name (FQDN) of the Windows Server Failover Clustering (WSFC) cluster when you browse for backups. In the NetBackup MS SQL Client, for the **Source Client** enter the FQDN of the WSFC cluster. If you use the short cluster name, NetBackup may not display the correct list of backup images.

A restore of a SQL Server database fails with Status Code 5, or Error (-1), when the host name of the SQL Server or the SQL Server database name has trailing spaces

When the host name of a SQL Server or a SQL Server database name has one or more trailing spaces, NetBackup does not generate the restore script correctly. The trailing spaces in the SQL Server host name or the database name are truncated in the script. To successfully perform a restore, you must create and edit a restore script in the NetBackup MS SQL Client.

In the script, edit the `DATABASE` and the `NBIMAGE` lines to include the correct SQL Server host name or SQL Server database name. For example, assume that the server host name is "ACCT", you use the default instance, and that the database name is "DatabaseA". Notice the trailing spaces after the server host name and the database name.

Change the following lines:

```
DATABASE "DatabaseA"  
NBIMAGE "ACCT.MSSQL7.ACCT.db.DatabaseA.~.7.001of001.20151118121736..C"
```

To:

A move operation fails with Status Code 5, or Error (-1), when the SQL Server host name, the database name, or the database logical name has trailing spaces

```
DATABASE "DatabaseA"
NBIMAGE "ACCT.MSSQL7.ACCT .db.DatabaseA .~.7.001of001.20151118121736..C"
```

A move operation fails with Status Code 5, or Error (-1), when the SQL Server host name, the database name, or the database logical name has trailing spaces

If the SQL Server host name, database name, or database logical name has one or more trailing spaces, a move operation fails with Status Code 5 or Error (-1). To successfully perform a move operation, you must create and edit a move script in the NetBackup MS SQL Client.

For information on a workaround for this issue, please see the following tech note on the Cohesity Support website:

<https://support.cohesity.com/s/article/article-100031572>

Unable to discover or browse availability group replicas

You must have the Microsoft SQL Server Native Client version 11.0.7462 ODBC driver or later installed on the availability group replicas to be able to discover and to browse databases on a read-scale availability group.

In the web UI, a read-scale availability group is not discovered, but no error message is given.

SQL Server policy backups that use intelligent groups may fail with a status code 200 when the associated credentials in a credential rule are invalid

A SQL Server policy backup can fail with a status code 200 during scheduled backups when the following details apply:

- Instances that are associated with the intelligent groups in the policy were discovered and registered using credential rules.

- Intelligent groups are used in the policy to query the NetBackup environment for the databases.

A warning message: "No databases were found for instance {0}." can be found in the job details. This message indicates an issue where NetBackup attempted to access the SQL Server instances or databases and the NetBackup asset database appears to be empty.

Error from job details:

```

Jan 15, 2025 11:45:00 AM - Info nbjm (pid=3606803) starting backup job (jobid=5) for
client <mssql_client>, policy <mssql_auto_protect_policy>, schedule Full
Jan 15, 2025 11:45:00 AM - Info nbjm (pid=3606803) requesting MEDIA_SERVER_ONLY resources
from RB for backup job (jobid=5, request id:{71277C90-D368-11EF-A88B-C711544AD613})
Jan 15, 2025 11:45:00 AM - requesting resource <media_bdstu>
Jan 15, 2025 11:45:00 AM - requesting resource <media_server>.NBU_CLIENT.MAXJOBS.
<media_server>
Jan 15, 2025 11:45:00 AM - granted resource <media_server>.NBU_CLIENT.MAXJOBS.
<media_server>
Jan 15, 2025 11:45:00 AM - estimated 0 kbytes needed
Jan 15, 2025 11:45:00 AM - begin Parent Job
Jan 15, 2025 11:45:00 AM - begin Application Resolver: Step By Condition
Operation Status: 0
Jan 15, 2025 11:45:00 AM - end Application Resolver: Step By Condition; elapsed time
0:00:00
Jan 15, 2025 11:45:00 AM - begin Application Resolver: Resolver Discovery
Jan 15, 2025 11:45:01 AM - Warning nbpem (pid=3606866) No databases were found
for instance {0}.
Jan 15, 2025 11:45:01 AM - Error nbpem (pid=3606866) backup of client <mssql_client>
exited with status 200 (The scheduler found that no backups are due. Or, the target
hosts do not need to be upgraded.)
Operation Status: 0
Jan 15, 2025 11:45:01 AM - end Application Resolver: Resolver Discovery; elapsed time
0:00:01
Jan 15, 2025 11:45:01 AM - begin Application Resolver: Policy Execution Manager
Preprocessed
Jan 15, 2025 11:45:01 AM - Warning nbpem (pid=3606866) No databases were found for
instance {0}.
Operation Status: 200
Jan 15, 2025 11:45:01 AM - end Application Resolver: Policy Execution Manager
Preprocessed; elapsed time 0:00:00
Jan 15, 2025 11:45:01 AM - begin Application Resolver: Stop On Error
Operation Status: 0
Jan 15, 2025 11:45:01 AM - end Application Resolver: Stop On Error; elapsed time 0:00:00

```

Operation Status: 200

Jan 15, 2025 11:45:01 AM - end Parent Job; elapsed time 0:00:01

The scheduler found that no backups are due. Or, the target hosts do not need to be upgraded. (200)

Environment

- Any supported Windows version
- Any supported SQL Server version
- NetBackup 11.0 or later

Cause

An issue exists with the credentials that are configured in the credential rule. The rule may contain an invalid username or password. Or, the user does not have the necessary permissions to access the instances or databases. This situation can occur because NetBackup does not validate a credential when the credential rule is created. Consider a case where a NetBackup administrator receives a request from the SQL Server administrator to create a credential rule. Later when the SQL Server administrator activates the credential rule, that administrator finds that the credential that was selected for the rule has an outdated username or password.

Solution

Validate the configured credentials for the SQL Server instances or databases that are included in the policy.

To validate the credentials

- 1 On the **Instances** tab, select the check box for instance.
- 2 Select **Manage credentials**.
- 3 Select the option **Select from existing credentials**.
- 4 Select the credential that is associated with the credential rule.
- 5 Select **Next**.
- 6 Review the messages that the validation displays. Then modify the credential as needed.

About disaster recovery of SQL Server

SQL Server corrects itself automatically from temporary or minor problems. However, most disasters are beyond the scope of the automatic recovery feature. For example, if a database becomes severely corrupted, or there is a catastrophic failure, recovery is initiated by the system administrator.

User-initiated recovery can entail either restoring the entire server, including the SQL Server databases, from full system backups. Or recovery can include restoring only the SQL Server databases to a newly-installed or other available SQL Server.

Restoring the entire server has the added benefit of recovering other applications and data which may have resided on the server at the time of failure. Restoring be accomplished using one of the following methods:

- Manual recovery of the server. This method involves manually restoring the server from full system backups.
See [“Preparing for disaster recovery of SQL Server”](#) on page 275.
- NetBackup Bare Metal Restore. BMR automates system recovery by restoring the operating system, system configuration, and all system files and data files. See the [NetBackup Bare Metal Restore Administrator's Guide](#) for more information.

After recovery of the server is complete, or after the new server installation is available, recovery of the SQL Server databases can begin.

Preparing for disaster recovery of SQL Server

When you develop your SQL Server disaster recovery plan you need to plan how to recover from corruption of the master database. You also need to plan for loss of your host machine. If the master database has been corrupted, then SQL Server does not start. When disaster happens you may need to rebuild the system databases. This process, however, does not recreate the schema information of your application databases. To recover your database schema use NetBackup to restore your latest backup of the master database.

Disaster recovery of SQL Server assumes that you have already put in place a strategy to recovery from other sorts of data loss. Data loss can include disk, software, and human error. To prepare for disaster recovery you need to make frequent backups of the master database. Do frequent backups after you have added or dropped databases or carried out other operations that may result in schema definitions.

Recovering SQL Server databases after disaster recovery

For the purposes of disaster recovery, you should only restore to a new installation of SQL Server. However, you can restore an existing installation of SQL Server with other active databases. The server should be running the same version of Windows on the same hardware platform. It also should be running the same version of SQL Server with the same service pack as the original server.

To recover SQL Server databases

- 1 If you want to restore to an existing SQL Server, choose from one of the following:
 - For a new SQL Server installation or when the master database is intact, continue with step 4.
 - If the master database is corrupt, you must first rebuild the master database. Continue with step 2.

- 2 Refer to the following article for instructions on how to rebuild the master database. Click the “Other Versions” drop-down list to select the correct SQL Server version.

<http://msdn.microsoft.com/en-us/library/ms144259.aspx>

Look for the information that describes how to rebuild system databases for a default instance from the command prompt.

- 3 When the rebuild is complete, restart the SQL Server services if necessary.
- 4 To begin the restore of the master database, start SQL Server in single-user mode.

The procedure to start SQL Server in single-user mode is described in the following article:

<http://msdn.microsoft.com/en-AU/library/ms188236.aspx>

Click the “Other Versions” drop-down list to select the correct SQL Server version.

- 5 Open the NetBackup MS SQL Client interface.
- 6 Locate all the media that is required to perform the restore operations.
- 7 Select **File > Restore SQL Server objects**.
- 8 Select the backup image that contains the copy of the master database you want to restore.

Select only the master database at this time.
- 9 Click **Restore**.

- 10** Restart the SQL Server service after the restore completes.
- 11** Continue with the restore of the remaining SQL Server databases.

Follow the instructions for restoring SQL databases, differentials, transaction logs, files, and filegroups.

When all of the restore operations have completed successfully, then the recovery of the SQL Server databases is complete.

After the recovery is complete, Cohesity recommends that you perform a full database backup as soon as possible.

Other configurations

This appendix includes the following topics:

- [Configuring multiplexed backups of SQL Server](#)
- [Restoring a multiplexed SQL Server backup](#)
- [About SQL Server backups and restores in an SAP environment](#)
- [Configure NetBackup to support database log-shipping](#)
- [Backing up SQL Server in an environment with log shipping](#)
- [About NetBackup for SQL Server with database mirroring](#)

Configuring multiplexed backups of SQL Server

Multiplexing lets you interleave multiple backups to the same tape. This feature is useful if you have many simultaneous backups that use the same tape drive.

However, multiplexing can interfere with SQL Server recovery due to how SQL Server requests streams during a restore. If you enabled multiplexing for multistreamed backups, see the information on how to perform restores. To restore a multiplexed backup, you must configure the restore for one stripe.

See [“Restoring multistreamed SQL Server backups”](#) on page 238.

Configure the following to create a multiplexed backup:

- In the backup policy, select the number of **Stripes** you want to use.
For SQL Server Intelligent policy, configure this setting on the **Microsoft SQL Server** tab. For batch file-based policies, configure the **Stripes** setting when you create the backup batch file.
- In the schedules for your policy, set **Media multiplexing** to the number of backup stripes that you want to use.

For batch file-based policies, enable multiplexing in the “Application Backup” schedule.

- In the storage units that are associated with this schedule, select **Enable Multiplexing** and set **Maximum streams per drive** to the number of stripes that you want to use.

Restoring a multiplexed SQL Server backup

In most cases, Cohesity does not recommend multiplexing multiple SQL Server streams from the same backup to a single tape. However, you may want to do use this method if you vault or export backup images. During the restore of this type of multiplexed backup, NetBackup may time out while it tries to synchronize access to data blocks from the backup tape. To prevent this time out, change the **Stripes** value to **1**. For a batch file-based policy, in the recovery batch file change the value from `STRIPES N` to `STRIPES 1`.

When you change this value it causes the restore to be performed in a single-stream. NetBackup presents the *N* backup images to SQL Server one at a time. The tape is rewound between the restore of each image.

About SQL Server backups and restores in an SAP environment

Note: SQL Server in an SAP environment is not supported for SQL Server Intelligent Policy.

With NetBackup you can perform scheduled SAP backups, in accordance with a predefined backup strategy, or manual backups. These backups may not be planned and may be necessary in exceptional situations. The practices that are described here are based on the practices SAP recommends in SAP/MS SQL Server DBA in CCMS.

The NetBackup backup and restore procedures for the SAP R/3 database are identical to the NetBackup procedures with any other SQL Server database.

You can create scripts to perform full or differential backups of databases and backups of transaction logs. In addition to the database backups and restores, NetBackup also provides the capabilities to back up the SAP file systems.

Creating batch files for automatic backups in for SQL Server in an SAP environment

NetBackup for SQL Server uses batch files to initiate database backup and restore operations. A batch file must be created for database backups and for transaction log backups. These batch files must then be added to the backup selections list in the backup policies that you created.

See [“Create a batch file for database backups”](#) on page 280.

See [“Create a batch file for transaction log backups”](#) on page 280.

Create a batch file for database backups

This topic describes how to create a batch file for database backups.

To create a script for database backups

- 1 Open the NetBackup MS SQL Client.
- 2 Select **File > Backup SQL Server objects**.
- 3 In the **Backup Microsoft SQL Server Objects** dialog box, in the left pane, expand the database instance.
- 4 In the right pane, select the R/3 database.
- 5 From the **Type of Backup** list, select the type of backup you want to perform, **Full** or **Full differential**.
- 6 Under **Backup Script**, select **Save**.
- 7 Click **Backup**.
- 8 Specify a file name and click **Save**.

Alternatively, you can select the name of an existing file, and NetBackup appends the new script to it.

- 9 Click **Yes** to open and edit the batch file.

Create a batch file for transaction log backups

This topic describes how to create a batch file for transaction log backups.

To create a batch file for transaction log backups

- 1 Before starting a transaction log backup, the database administrator should set the **Transaction log backup options** database option to off. This option on the SQL Server interface applies to the databases.

The entire sequence of transaction logs generated following any database dump must be maintained on the same NetBackup server. NetBackup for SQL Server requires that you follow these guidelines in devising your backup strategy to ensure success in restoring your database.

- 2 Select File > **Backup SQL Server objects**.
- 3 In the **Backup Microsoft SQL Server Objects** dialog box, in the left pane, expand the database instance.
- 4 In the right pane, select the R/3 database.
- 5 For the **Type of Backup**, select **transaction log**.
- 6 Under **Backup Script**, select **Save**.
- 7 Click **Backup**.
- 8 Specify a file name and click **Save**.

Alternatively, you can select the name of an existing file, and NetBackup appends the new script to it.

- 9 Click **Yes** to open and edit the batch file.

Monitoring backups on SQL Server

Check scheduled backups regularly to ensure that they completed successfully.

Always check the following:

- That the most recent backup has run successfully.
See [“About monitoring NetBackup for SQL Server operations”](#) on page 260.
- All the backups in the backup cycle are executed according to the schedule.
Gaps in a backup sequence can have serious consequences in a subsequent attempt to restore the database.

Restoring the R/3 database

This topic describes how to restore the R/3 database.

Determine how to perform the restore based on the following scenarios:

- If you have scheduled differential backups, review the information for that type of restore.

See [“About including differential backups in a restore operation”](#) on page 282.

- If the R/3 database disk system is damaged or the transaction log disk system is damaged, follow the instructions for that scenario.
See [“Restore the R/3 database after a disk crash”](#) on page 282.
- To perform a regular restore of the R/3 database, follow the instructions for that type of restore.
See [“Restore the database backups and transaction log backups”](#) on page 283.

About including differential backups in a restore operation

If you incorporated differential backups in the backup strategy, the restore process differs depending on the type of backups available.

Determine how to perform the restore based on which of the following differential backups you have:

- If differential backups were made after the last full database backup, restore the last database backup that is followed by the most recent differential backup. Then apply all subsequent transaction logs.
- If no differential backups were made since the last full database backup, restore the last full database backup and then apply all subsequent transaction logs.
- If several differential backups are available but the latest one cannot be read, restore the most recent full database backup. And restore the latest readable differential backup and apply all subsequently created transaction logs.

Restore the R/3 database after a disk crash

This topic describes how to restore the database when the R/3 database disk system is damaged or the transaction log disk system is damaged. This process is only applicable to a configuration with three disk systems: one system for the R/3 database, one for the R/3 transaction logs and one for all others.

Note: The R3 database must not be in use when you are performing a restore operation. Make sure that all SAP services are stopped before you attempt a restore with NetBackup.

Warning: If the disk system on which the R/3 database resides is damaged, it is vital to immediately back up the currently active transaction log. This log backup is done to prevent loss of data. Without a backup of the current log, the database can only be restored to the status at the time of the last transaction log backup. If work has been carried out on the R/3 system since then, this work is lost.

To restore the R/3 database after a disk crash

- 1 Back up the current transaction log.
- 2 Replace damaged disks.

Replacing damaged disks in a RAID disk system is normally a straightforward procedure. If you are uncertain how to proceed, see the documentation of your hardware vendor to learn how to handle the disks. The new disks must be formatted and assigned the same drive letter as the old disks.

- 3 Restore the database logs and transaction logs.

The central phase of a restore operation is the reloading of the database backup and the application of the available transaction logs. When the database backup is reloaded, the database files are automatically recreated. The data is copied from the backup device to the newly created files. Once this copy has been done, the transaction logs are applied in the same sequence as they were originally made. In a final step, open transactions that were not completed at the time of the database failure are rolled back.

Restore the database backups and transaction log backups

The NetBackup MS SQL Client provides for automatic staging. By selecting the latest transaction log backup, NetBackup automatically restores the previous full database backup. It also restores any optional differential backups and subsequent transaction log backups. You can also use the option to specify a point in time to which to restore to.

Note: The R3 database must not be in use when performing a restore operation. Make sure that all SAP services are stopped before you attempt a restore with NetBackup.

Warning: To restore the R/3 database you first restore the most recent database backup and then the subsequent transaction logs. During the entire procedure, do not execute any transactions and do not shut down the database server. A server shutdown would write a checkpoint to the log and as a result you would not be able to restore further transaction logs.

To restore the database backups and transaction log backups

- 1 Restore the most recent database backup.
- 2 Restore the latest differential database backup (if available).

- 3 Restore all succeeding transaction log backups.
- 4 Restore the latest transaction log backup.

About policy configuration for SQL Server in an SAP environment

To automatically perform backups of an SAP environment, you need to create backup policies. A backup policy with the "MS-SQL-Server" policy type that is selected must be created for R/3 database backups. Batch files, which initiate the backup of the database and transaction logs, must be added to the backup selections list in the policy.

Information is available for how to create the batch files that are needed and how to configure backup policies.

See ["Creating batch files for automatic backups in for SQL Server in an SAP environment"](#) on page 280.

For backups of the executables disk (a file-system backup), a backup policy must be created with the Windows policy type selected.

For information on Windows policies, see the [NetBackup Administrator's Guide, Volume I](#).

About manual backups of SQL Server in an SAP environment

The administrator on the primary server can use the NetBackup web UI to manually run an automatic backup schedule. This schedule can be for an "MS-SQL-Server" policy, where the R/3 database is specified in the backup script.

For more information, see the section on manual backups in the [NetBackup Administrator's Guide, Volume I](#).

Configure NetBackup to support database log-shipping

Log shipping is a SQL Server feature that may be employed to enhance the overall availability of your installation. It uses a primary server, which contains the active database, a monitor, and one or more secondary servers. Under log shipping, copies of the transaction log are supplied to the secondary servers on a per-transaction basis to the secondary servers. This configuration allows each secondary server to be in a standby state in case the primary goes offline.

To use log-shipping with NetBackup, both the primary and the secondary should be set up as clients of the same primary server. During transaction log backups truncation is not performed for these databases even if the policy or the protection

plan settings enable truncation. SQL Server requires the transaction logs to keep the log-chain intact for the log-shipping feature to function correctly. A message displays in the progress log if this situation occurs: `Truncate option ignored on transaction log backup for log shipping primary database: LogShippingDB.`

To configure NetBackup to support database log-shipping

- 1 The hosts that contain both databases should specify the same primary server in their server lists.
- 2 Any policy or any protection plan that backs up the primary should also specify the host that contains the secondary database.

See [“Backing up SQL Server in an environment with log shipping”](#) on page 285.

- 3 On the primary server, configure permissions for redirected restores for both the primary and the secondary server.

See [“Configuring permissions for redirected restores”](#) on page 161.

Backing up SQL Server in an environment with log shipping

Many sites also use the secondary server to off-load certain activities from the primary to minimize its load. However, a backup should *not* be performed on a secondary (or standby) server. Databases must always be backed up on the primary server and restored on the primary server. This requirement is based on the Microsoft SQL Server restriction that is outlined in Microsoft knowledge base article 311115.

If you try to perform a backup on the secondary server, you see a message in the `dbclient` log similar to the following:

```
16:33:26 [1208,2348] <16> CODEBAccess::LogODBCerr: DBMS MSG - ODBC message. ODBC return
code <-1>, SQL State <37000>, Message Text <[Microsoft][ODBC SQL Server Driver][SQL
Server]Database 'Mumbo' is in warm-standby state (set by executing RESTORE WITH
STANDBY) and cannot be backed up until the entire load sequence is completed.>
```

NetBackup can automatically skip these databases. In the settings for the SQL Server policy (intelligent policy or batch file-based policy) or the SQL Server protection plan, select **Skip unavailable (offline, restoring, etc.) databases**.

See [“NetBackup for SQL Server performance factors”](#) on page 257.

About NetBackup for SQL Server with database mirroring

Note: Database mirroring is not supported for SQL Server Intelligent Policy.

Database mirroring is a software solution that increases the availability of a SQL Server database. It uses two database instances (normally on different hosts), which contain copies of the same SQL Server database. These databases are identical in both name and content. The copies are the principal and the mirror. The mirror serves as a hot standby to the principal, where transactions take place. The mirror is very closely synchronized with the principal through transaction log porting. It is immediately available in case the principal fails.

The primary consideration when you establish your backup and restore procedures for database mirroring is that these operations are only available on the principal database.

For a complete description of database mirroring refer to the *SQL Server Books Online*.

Configure NetBackup to support database mirroring

To use database mirroring with NetBackup, both the principal and the mirror should be set up as clients of the same primary server.

To configure NetBackup to support database mirroring

- 1 The hosts that contain both databases should specify the same primary server in their server lists.
- 2 Any policy that is used to back up the principal should also specify the host that contains the mirror database.

See [“Perform simultaneous backups for mirrored partners”](#) on page 287.
- 3 On the primary server, configure permissions for a redirected restore for both mirroring partners.

See [“Configuring permissions for redirected restores”](#) on page 161.
- 4 (Conditional) If you specify the fully-qualified domain name (FQDN) for the client in the backup policy, you need to create an alias for the short client name. This alias lets you successfully browse for a backup image and restore it in a mirrored environment. NetBackup attempts to find a mirrored partner backup image using the short name of the client host (for example, `client1`). However,

the backup image in this case is stored using the FQDN (for example, `client1.domain.com`).

You can create an alias in one of the following ways:

- On the NetBackup client, create the following touch file:
`install_path\dbext\mssql\ClientNameMapping.txt`
 Add an entry <short name of client host> <FQDN of client host>.
 For example:
`client1 client1.domain.com`
- On the NetBackup primary server, use the `bpclient` command to create the alias:

```
bpclient -client client_name -M master_server -add_alias alias_name
```

For example:

```
bpclient -client client1.domain.com -M primary.domain.com -add_alias hpe013-vm02
```

You must use the FQDN for the `-client` argument.

Perform simultaneous backups for mirrored partners

Since backups can occur only on the principal, you must take steps to ensure that you don't miss any scheduled backups due to failover. Establish a procedure to simultaneously initiate backups for both partners, but suppress the operation on the mirror.

When you restore a mirrored database, you must restore it to the node currently in the principal role. See *SQL Server Books Online*.

To simultaneously initiate backups for both partners

- 1 Create a policy with a backup schedule for the principal.
- 2 Add the host that contains the mirroring partner to the client list.
- 3 Create a batch file and add it to the backup selections list.
- 4 Create a batch file on the mirroring partner that has the same name as the batch file specified in the backup selections policy.

The batch file on the mirroring partner should be identical to the one used on the principal, with one exception. The value for `SQLHOSTS` and `SQLINSTANCE` are different.

Restore a mirrored database backup image

Note: Before you restore a mirrored database, you must remove the mirroring attribute.

For mirrored databases, NetBackup can create backup images on either or on both the principal and the mirror server. The **Restore Database** dialog box displays any backups images from both servers. To determine which partner the backup was taken from, look at the property page for the image. To view backup images you can select the **Host name** that contains either of the mirroring partners, provided that NetBackup performed backups for that partner.

For example, assume that mirroring partners are as follows. All of the backups were done on `HostB`, though the principal is currently on `HostA`:

- **Principal**
 Host name: `HostA`
 SQL Server instance: `Solaria`
 Database: `Accounting`
- **Mirror**
 Host name: `HostB`
 SQL Server instance: `Moonbeam`
 Database: `Accounting`

If backup images were created exclusively on `HostA` or on both `HostA` and `HostB`, you can view the images from both partners. Select `HostA` in the **SQL Host** list.

To restore a mirrored backup image

- 1** Disable mirroring on the principal mirror.
 You can use the appropriate commands in SQL Server Management Studio or use `ALTER DATABASE` directly.
- 2** On the principal server, open the NetBackup MS SQL Client.
 When you restore a mirror database, you must run the NetBackup MS SQL Client from the principal server. See *SQL Server Books Online* for information on how to determine which partner is the principal.
 In the previous example, the principal is `HostA`.
- 3** On the **File** menu, select **Restore SQL Server Objects**.
- 4** In the **Backup History Options** dialog box, from the **SQL host** list select the mirror server.
 In the previous example, the mirror is `HostB`.

5 Click **OK**.

6 Proceed with the restore as normal.

NetBackup creates a recovery script for the database that includes images from both partners, as appropriate.

Register authorized locations

This appendix includes the following topics:

- [Registering authorized locations used by a NetBackup database script-based policy](#)

Registering authorized locations used by a NetBackup database script-based policy

During a backup, NetBackup checks for scripts in the default script location and any authorized locations. The default, authorized script location for UNIX is `usr/opensv/netbackup/ext/db_ext` and for Windows is `install_path\netbackup\dbext`. If the script is not in the default script location or an authorized location, the policy job fails. You can move any script into the default script location or any additional authorized location and NetBackup recognizes the scripts. You need to update the policy with the script location if it has changed. An authorized location can be a directory and NetBackup recognizes any script within that directory. An authorized location can also be a full path to a script if an entire directory does need to be authorized.

If the default script location does not work for your environment, use the following procedure to enter one or more authorized locations for your scripts. Use `nbsetconfig` to enter an authorized location where the scripts reside. You can also use `bpsetconfig`, however this command is only available on the primary or the media server.

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

For more information about registering authorized locations and scripts, review the knowledge base article:

<https://support.cohesity.com/s/article/article-100039639>

To add an authorized location

- 1 Open a command prompt on the client.
- 2 Use `nbsetconfig` to enter values for an authorized location. The client privileged user must run these commands.

The following examples are for paths you may configure for the Oracle agent. Use the path that is appropriate for your agent.

- On UNIX:

```
[root@client26 bin]# ./nbsetconfig
nbsetconfig>DB_SCRIPT_PATH = /Oracle/scripts
nbsetconfig>DB_SCRIPT_PATH = /db/Oracle/scripts/full_backup.sh
nbsetconfig>
<ctrl-D>
```

- On Windows:

```
C:\Program Files\Cohesity NetBackup\NetBackup\bin>nbsetconfig
nbsetconfig> DB_SCRIPT_PATH=c:\db_scripts
nbsetconfig> DB_SCRIPT_PATH=e:\oracle\fullbackup\full_rman.sh
nbsetconfig>
<ctrl-Z>
```

Note: Review the [NetBackup Command Reference Guide](#) for options, such as reading from a text file and remotely setting clients from a NetBackup server using `bpsetconfig`. If you have a text file with the script location or authorized locations listed, `nbsetconfig` or `bpsetconfig` can read from that text file. An entry of `DB_SCRIPT_PATH=none` does not allow any script to run on a client. The `none` entry is useful if an administrator wants to completely lock down a server from running scripts.

Registering authorized locations used by a NetBackup database script-based policy

- 3** (Conditional) Perform these steps on any clustered database or agent node that can perform the backup.
- 4** (Conditional) Update any policy if the script location was changed to the default or authorized location.