

NetBackup™ for Microsoft SQL Server Administrator's Guide

for Windows

Release 10.1

VERITAS™

NetBackup™ for Microsoft SQL Server Administrator's Guide

Last updated: 2022-09-08

Legal Notice

Copyright © 2022 Veritas Technologies LLC. All rights reserved.

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

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

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

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

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

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

Veritas Technologies, LLC
2625 Augustine Drive
Santa Clara, CA 95054

<http://www.veritas.com>

Technical Support

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

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

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

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

Documentation feedback

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

NB.docs@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1	Introducing NetBackup for SQL Server	11
	Overview of NetBackup for SQL Server	11
Chapter 2	Installation	15
	Planning the installation of NetBackup for SQL Server	15
Chapter 3	Host configuration and job settings	17
	Configuring SQL Server hosts	17
	Installing the Veritas VSS provider for vSphere	18
	Configuring the NetBackup services for SQL Server backups and restores	19
	Configure local security privileges for SQL Server	21
	Reviewing the auto-discovered mappings in Host Management	22
	Configuring mappings for restores of a distributed applications, clusters, or virtual machines	27
	Configuring the primary server host name for the SQL Server agent	28
	Configure the number of jobs allowed for backup operations	29
	Configure the Maximum jobs per client setting	30
Chapter 4	Managing SQL Server objects for use with SQL Server Intelligent Policies	32
	About the Applications utility	32
	About discovery of SQL Server objects	33
	Discovering instances on demand	34
	Discover advanced or basic availability groups on demand	34
	Discover read-scale availability groups	34
	About registering SQL Server instances and availability replicas	35
	About credentials used with SQL Server Intelligent Policy	35
	Registering a SQL Server instance or availability replica	38
	Registering instances or availability replicas with an instance group	39
	Registering instances or availability replicas automatically	42

Authorizing a DBA to register instances or availability replicas with the <code>nbsqladm</code> command	43
Deleting SQL Server objects from the Applications utility	43
Manually add a SQL Server instance	44
Deactivating or activating an instance	45
Cleaning up instances	46

Chapter 5	Configuring backups with SQL Server Intelligent Policy	47
	About SQL Server Intelligent Policies	48
	Creating a SQL Server Intelligent Policy	48
	About policy attributes	49
	About schedule properties	50
	Schedule backup types for SQL Server Intelligent Policies	51
	Adding instances to a policy	53
	Adding databases to a policy	55
	Adding filegroups or files to the backup selections list	57
	Manually adding files or filegroups to the backup selections list	59
	Adding instance groups to a backup policy	59
	About tuning parameters for SQL Server backups	60
	Configuring multistriped backups of SQL Server	64
	Backing up read-only filegroups	64
	Backing up read-write filegroups	65

Chapter 6	Performing restores of SQL Server	67
	Starting the NetBackup MS SQL Client for the first time	68
	Selecting the SQL Server host and instance	68
	Browsing for SQL Server backup images	69
	Options for NetBackup for SQL Server restores	71
	Restoring a SQL Server database backup	73
	Staging a full SQL Server database recovery	74
	Restoring SQL Server filegroup backups	74
	Recovering a SQL Server database from read-write filegroup backups	75
	Restoring SQL Server read-only filegroups	76
	Restoring SQL Server database files	76
	Restoring a SQL Server transaction log image without staging a full recovery	77
	Performing a SQL Server database move	77
	About performing a SQL Server page-level restore	79
	Configuring permissions for redirected restores	81
	Redirecting a SQL Server database to a different host	83

	About selecting a primary server	84
	Performing a restore of a remote SQL Server installation	84
	Restoring multistreamed SQL Server backups	85
	About using bplist to retrieve SQL Server backups	86
	About NetBackup for SQL Server backup names	87
Chapter 7	Protecting SQL Server data with VMware backups	90
	About protecting an application database with VMware backups	90
	Limitations of VMware application backups	91
	About configuring NetBackup for VMware backups that protect SQL Server	92
	Configuring the NetBackup services for a VMware backup that protects SQL Server	93
	Configuring a VMware backup policy to protect SQL Server	94
	Configuring a VMware policy to protect SQL Server using Replication Director to manage snapshot replication	96
	Restore SQL Server databases from a VMware backup	98
Chapter 8	Configuring backups with Snapshot Client	99
	About NetBackup Snapshot Client for SQL Server	99
	How SQL Server operations use Snapshot Client	100
	Configuration requirements for SQL Server snapshot and Instant Recovery backups	103
	Configuring a snapshot policy for SQL Server	104
	Configuring a policy for Instant Recovery backups of SQL Server	106
	Using copy-only snapshot backups to affect how differentials are based	108
	Creating a copy-only backup (legacy SQL Server policies)	109
	Creating an Instant Recovery backup that is not copy-only (legacy SQL Server policies)	109
	About SQL Server agent grouped backups (legacy SQL Server policies)	109
	Viewing the progress of a grouped backup	110
	Restoring a database backed up in a group	111
Chapter 9	Protecting SQL Server availability groups	113
	About protecting SQL Server availability groups	113
	Protecting SQL Server availability groups with intelligent policies	115
	Prerequisites for protecting SQL Server availability groups	115

Configuring a backup policy to protect a SQL Server availability group	116
Protecting SQL Server availability groups with legacy policies	120
About protecting the preferred replica in a SQL Server availability group (legacy backup policies)	121
About protecting a specific node in a SQL Server availability group (legacy backup policies)	126
Protect a SQL Server availability group that crosses NetBackup domains	129
Browsing for SQL Server availability group backup images	130
Restoring a SQL Server availability database to a secondary replica	131
Restoring a SQL Server availability database to the primary and the secondary replicas	132
Restoring an availability database when an availability group crosses NetBackup domains	133

Chapter 10 Protecting SQL Server in a cluster environment

Configuring backups of clustered SQL Server instances (SQL Server Intelligent Policy)	135
Configuring backups of clustered SQL Server instances (legacy SQL Server policies)	137
Performing a restore of a virtual SQL Server instance	138

Chapter 11 Configuring backups with legacy SQL Server policies using clients and batch files

About legacy SQL Server policies	140
About configuring backups with legacy SQL Server policies	140
Configuring the NetBackup services for SQL Server backups and restores (legacy SQL Server policies)	141
About SQL Server security with NetBackup legacy backup policies	142
About using batch files with NetBackup for SQL Server	143
Keywords and values used in batch files	144
Creating a batch file	152
Running batch files	153
Adding a new SQL Server legacy policy	153
About schedule properties	154
Legacy policy backup types	155
Converting differential backups to full backups	156

Configuring an application backup schedule	157
Example application backup schedule	157
Configuring automatic backup schedules	158
Example automatic backup schedule	158
Adding clients to a policy	159
Adding batch files to the backup selections list	160
Selecting the SQL Server host and instance	161
Options for SQL Server backup operations	162
About viewing the properties of the objects selected for backup	165
Performing user-directed backups of SQL Server databases	166
Performing user-directed backups of SQL Server transaction logs	166
Performing user-directed backups of SQL Server database filegroups	167
Performing user-directed backups of read-only filegroups	168
Performing user-directed backups of read-write filegroups	169
Performing user-directed backups of SQL Server database files	170
Performing partial database backups	171
Performing a backup of a remote SQL Server installation	172
About file checkpointing with NetBackup for SQL Server	173
About automatic retry of unsuccessful SQL Server backups	174

Chapter 12

Using NetBackup for SQL Server with multiple NICs	176
About configuration of SQL Server backups with multiple NICs	176
Configuring the NetBackup client with the private interface name	178
Configuring backups of SQL Server when you have multiple NICs (SQL Server Intelligent Policies)	179
Configuring backups for SQL Server when you have multiple NICs (legacy SQL Server policies)	180
Performing restores of SQL Server when you have multiple NICs	181
Configuring backups of a SQL Server cluster when you have multiple NICs (SQL Server Intelligent Policies)	182
Configuring backups of a SQL Server cluster when you have multiple NICs (legacy SQL Server policies)	183
Creating a batch file for backups of a SQL Server cluster when you have multiple NICs (legacy SQL Server policies)	183
Performing restores of a SQL Server cluster when you have multiple NICs	185

Chapter 13	Performance and troubleshooting	188
	What are the components of NetBackup for SQL Server?	189
	How does NetBackup for SQL Server back up a database?	191
	How does NetBackup for SQL Server recover a database?	191
	Performing a manual backup	192
	About debug logging for SQL Server troubleshooting	192
	Setting the debug level	194
	Veritas VSS provider logs	194
	NetBackup for SQL Server performance factors	195
	About monitoring NetBackup for SQL Server operations	198
	Setting the maximum trace level for NetBackup for SQL Server	200
	Troubleshooting credential validation	201
	Reporting of unsuccessful filegroup or file backups	202
	About minimizing timeout failures on large SQL Server database restores	202
	Troubleshooting VMware backups	203
	SQL Server log truncation failure during VMware backups of SQL Server	204
	SQL Server restore fails when you restore a SQL Server compressed backup image as a single stripe or with multiple stripes	205
	Incorrect backup images are displayed for availability group clusters	206
	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	206
	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	207
	Unable to discover or browse availability group replicas	207
	About disaster recovery of SQL Server	207
	Preparing for disaster recovery of SQL Server	208
	Recovering SQL Server databases after disaster recovery	208
Appendix A	Other configurations	210
	Configuring multiplexed backups of SQL Server	210
	Restoring a multiplexed SQL Server backup	211
	About SQL Server backups and restores in an SAP environment	211
	Creating batch files for automatic backups in for SQL Server in an SAP environment	212
	Monitoring backups on SQL Server	213
	Restoring the R/3 database	213

About policy configuration for SQL Server in an SAP environment	216
Configuring NetBackup to support database log-shipping	216
Backing up SQL Server in an environment with log shipping	217
About NetBackup for SQL Server with database mirroring	217
Configuring NetBackup to support database mirroring	218
Performing simultaneous backups for mirrored partners	219
Restoring a mirrored database backup image	219
Appendix B Register authorized locations	221
Registering authorized locations used by a NetBackup database script-based policy	221

Introducing NetBackup for SQL Server

This chapter includes the following topics:

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

Overview of NetBackup for SQL Server

NetBackup for SQL Server provides the capability for backups and restores of SQL Server databases. NetBackup offers the following types of SQL Server backup policies:

- **SQL Server Intelligent Policies.** A single policy protects multiple SQL Server instances that are spread over multiple clients. You select instances for a policy from a list of instances that are automatically discovered in the NetBackup environment.
- **Legacy policies,** using clients and batch files. These policies include a list of SQL database clients and a batch file that contains SQL backup commands to run when the backup is scheduled.

The NetBackup MS SQL Client lets you perform the following operations:

- View discovered instances, databases, or availability groups.
- Restore databases and database components.
- Configure restore options.
- Monitor restore operations.
- (SQL Server legacy policies) Backups of databases and database components and configuration of backup options.

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

Table 1-1 NetBackup for SQL Server features

Feature	Description
NetBackup integration	Full integration with the NetBackup primary server and media manager. Job monitoring from the server and the NetBackup MS SQL Client interface.
SQL Server Intelligent Policy	The following benefits are included: ■ Create a single policy to protect multiple SQL Server instances or instance databases or a policy to protect availability groups or availability databases. Instances can be spread over multiple clients. ■ Include a full, differential, and transaction log backup in the same policy. ■ Schedule frequent backups of transaction logs. ■ You are not required to know SQL Server commands or to write and use batch files. Instead, this feature automatically generates the batch files at run-time.
Management of SQL Server objects for use with Intelligent Policies	NetBackup automatically discovers SQL Server instances and availability groups in the environment. You can also perform manual discovery. After instances or replicas are registered, a SQL Server Intelligent Policy can be created that protects instances or availability groups. The user can also use instance groups to organize instances or replicas and, optionally, automatically register those objects.
Authentication and credentials	SQL Server Intelligent Policy supports the following: ■ Windows authentication and 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 SQL Server DBA can manage the SQL Server credentials and instance registration independently from the NetBackup administrator, with the <code>nbsqladm</code> command.

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

Feature	Description
Backup and restore features	The following features are available for backups and restores: ■ Backups and are managed entirely by the NetBackup server from a central location. Administrators can schedule automatic, unattended backups for instances on local or remote hosts across the network. ■ 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 and restores of only read-write filegroups. ■ Support for high availability (HA) environments, including SQL Server clusters and availability groups. ■ An administrator that uses the NetBackup MS SQL Client can browse backups and select the ones to be restored. ■ Restore SQL Server objects to different locations (redirected restores). ■ When the Encryption attribute is enabled, NetBackup encrypts the backup for the instances or clients that are listed in the policy. ■ Ability to use multiple stripes during a backup. ■ Tuning options that can improve the performance of backups.
Stream-based backups and restores	Stream-based backup and restore of SQL Server objects with SQL Server's high-speed virtual device interface.
Snapshot backups and instant access databases	NetBackup can perform backups of SQL Server with snapshot methodology. Also available are off-host backups, Instant Recovery, and backups with a hardware provider.
Support for VMware backups that protect SQL Server	Support for application-consistent backups of VMware computers using the VMware intelligent policy. The VMware intelligent policy includes three features that NetBackup for SQL Server supports: VMware snapshots, Replication Director (RD) snapshots, and Accelerator. Only full backups are supported on these three variations of the VMware intelligent policy. See the following documents for more information. NetBackup for VMware Administrator's Guide NetBackup Replication Director Solutions Guide NetBackup Administrator's Guide, Volume I
NetBackup encryption	When the Encryption attribute is enabled, NetBackup encrypts the backup for the instances or clients that are listed in the policy.
Multistreaming	Ability to use multiple stripes during a backup.

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

Feature	Description
Legacy SQL Server policies	Support for the legacy backup policies that use batch files and a list of clients.

Installation

This chapter includes the following topics:

- [Planning the installation of NetBackup for SQL Server](#)

Planning the installation of NetBackup for SQL Server

[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.	

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. The NetBackup for SQL Server agent is installed with the NetBackup client software. To use the new features that are included in NetBackup for SQL Server in NetBackup 10.1, upgrade your NetBackup for SQL Server clients to NetBackup 10.1. The NetBackup media server must use the same version as or a higher version than the NetBackup for SQL Server client.	Note the following: ■ 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.
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.	Review the following requirements: ■ The NetBackup server software is installed and configured to work in a NetBackup cluster. ■ The NetBackup client software is installed and operational on each node to which NetBackup can failover. ■ A valid license must exist for NetBackup for SQL Server on each node where NetBackup server resides. See the Software Compatibility List (SCL). See the NetBackup Installation Guide. See the NetBackup Clustered Master Server Administrator's Guide.

Host configuration and job settings

This chapter includes the following topics:

- [Configuring SQL Server hosts](#)
- [Installing the Veritas VSS provider for vSphere](#)
- [Configuring the NetBackup services for SQL Server backups and restores](#)
- [Configure local security privileges for SQL Server](#)
- [Reviewing the auto-discovered mappings in Host Management](#)
- [Configuring mappings for restores of a distributed applications, clusters, or virtual machines](#)
- [Configuring the primary server host name for the SQL Server agent](#)
- [Configure the number of jobs allowed for backup operations](#)
- [Configure the Maximum jobs per client setting](#)

Configuring SQL Server hosts

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 Veritas VSS provider.	See “Installing the Veritas VSS provider for vSphere” on page 18.
Step 3	(Conditional) For SQL Server Intelligent Policies, register the SQL Server credentials.	Add the SQL Server credentials that you need to register the instances or availability replicas. See “About credentials used with SQL Server Intelligent Policy” on page 35. Note: To use gMSA credentials, you must use the credential option Use credentials that are defined locally on the client.
Step 4	Configure the NetBackup Client Service and the NetBackup Legacy Network Service.	This configuration allows access to the SQL Server when NetBackup performs backups and restores. See “Configuring the NetBackup services for SQL Server backups and restores” on page 19.
Step 5	(Conditional) For SQL Server Intelligent Policies, 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 21.
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 NetBackup Management > Hosts properties on the master server. See “Reviewing the auto-discovered mappings in Host Management” on page 22.

Installing the Veritas VSS provider for vSphere

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

To install the Veritas VSS provider

- 1 Browse to the following location:
`install_path\Veritas\NetBackup\bin\goodies\`
- 2 Double-click on the **Veritas VSS provider for vSphere** shortcut.
- 3 Follow the prompts.
- 4 When the utility has completed, restart the computer if prompted.
- 5 Following the restart, the utility resumes. Follow the prompts to complete the installation.

To uninstall the Veritas VSS provider

- 1 In the Control Panel, open **Add or Remove Programs** or **Programs and Features**.
- 2 Double-click on **Veritas VSS provider**.
The uninstall program does not automatically reinstall the VMware VSS provider.

Configuring the NetBackup services for SQL Server backups and restores

For SQL Server intelligent policies, NetBackup uses the NetBackup Client Service and the NetBackup Legacy Network Service to access the SQL Server when it performs backups and restores.

Note the following requirements for the NetBackup services logon account:

- The account has the fixed server role “sysadmin”. You can use a domain account, a member of BUILTIN\Administrators, or another account that has this role.
- (non-VMware backups) If you want to use Local System for the logon account, apply the SQL Server sysadmin role manually to the NT AUTHORITY\SYSTEM or the BUILTIN\Administrators group.
- (VMware backups) You must use an account other than the Local System account as the logon account. Both services must use the same logon account.
- (VMware backups) If you choose to truncate logs, ensure that the account that runs the Microsoft SQL Server Service has full permissions for the NetBackup Legacy Network Service `temp` directory.

This directory is `C:\Users\user\AppData\Local\Temp`. *User* is the account that runs the NetBackup Legacy Network Service.

- To use a gMSA account for backups and restores, you must create a credential with the option **Use credentials that are defined locally on the client**.

- 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 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 21.

- 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:
 - **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 Host Management

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 23.

See [the section called “Auto-discovered mappings for a SQL Server cluster in a multiple NIC environment”](#) on page 25.

See [the section called “Manually map host names”](#) on page 26.

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 appear in the Host Management properties on the primary server. You can also use the `nbhostmgmt` command to manage the mappings. See the [NetBackup Administrator's Guide, Volume I](#) for more details on Host Management properties.

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
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 Administration Console, expand **Security Management > Host Management**.
- 2
- At the bottom of the **Hosts** pane, click the **Mappings for Approval** 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
- If a mapping is valid, right-click on a host entry and click **Approve**.

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

- 4 When you finish approving the valid mappings for the hosts, click on the **Hosts** tab at the bottom of the **Hosts** pane.

For hosts `client01.lab04.com` and `client02.lab04.com`, you see **Mapped Host Names/IP Addresses** 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

- 5 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>	Virtual name of the SQL Server cluster
	Physical name of <i>Node 2</i>	Virtual name of the SQL Server cluster
Basic or advanced availability group (primary and secondary)	Primary name	WSFC name
	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 Administration Console, expand **Security Management > Host Management**.
- 2 At the bottom of the **Hosts** pane, click the **Mappings for Approval** 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
client01-bk.lab04.com	clustername-bk.lab04.com
client02-bk.lab04.com	clustername-bk.lab04.com

- 3 If a mapping is valid, right-click on a host entry and click **Approve**.

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 at the bottom of the **Hosts** pane.

For hosts `client01-bk.lab04.com` and `client02-bk.lab04.com`, you may see the following **Mapped Host Names/IP Addresses**.

Host	Mapped Host Names/IP Addresses
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

- In the NetBackup Administration Console, expand **Security Management > Host Management**.
- Click on the **Hosts** tab.
- Right-click in the **Hosts** pane and click **Add Shared or Cluster Mappings**.
For example, provide the name of the virtual name of the cluster. Then click **Select Hosts** 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 primary server, open the NetBackup Administration Console.
- 2
- Select **NetBackup Management > Host Properties > Master Servers**.
- 3
- In the right pane, double-click on the 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 SQL Server

Table 3-4 Example entries for SQL Server

Environment	Application host	Component host
FCI (cluster with two nodes)	Virtual name of the SQL Server cluster	Physical name of <i>Node 1</i>
	Virtual name of the SQL Server cluster	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
	Virtual name of the SQL Server cluster	Physical name of <i>Node 1</i>
	Virtual name of the SQL Server cluster	Physical name of <i>Node 2</i>

Table 3-4 Example entries for SQL Server (continued)

Environment	Application host	Component host
VMware	VM display name, VM BIOS UUID, or VM DNS name (Primary VM identifier other than VM hostname)	Host name of the VM

Configuring the primary server host name for the SQL Server agent

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

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

- **NBSERVER** value.
For 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 legacy 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. See the NetBackup Administrator's Guide, Volume I.
Maximum jobs per client	In a policy, the maximum number of jobs per client that you want to allow. This setting applies to all clients in all policies. It is configured in the primary server host properties in the Global Attributes. See "Configure the Maximum jobs per client setting" on page 30.

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 (Intelligent Policies). Each instance or database that is specified in the policy creates a new backup job. For legacy policies, this setting indicates the maximum that is allowed per client.

To configure the maximum jobs per client

- 1 In the left pane of the NetBackup Administration Console, expand **NetBackup Management > Host Properties**.
- 2 Select **Master Server**.
- 3 In the right pane, double-click the server icon.
- 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 legacy policies, use the following formula to calculate a smaller value for the **Maximum jobs per client** setting:

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

Refer to the following definitions:

- number of database_objects* (Intelligent Policies) The number of databases, filegroups, or files that you want to back up in parallel.
- number_of_streams* The number of backup streams between the database server and NetBackup. If striping is not used, each separate stream starts a new backup job on the client. If striping is used, each new job uses one stream per stripe.
- number_of_policies* The number of policies of any type that can back up this client at the same time. This number can be greater than one. For example, a client can be in two policies to back up two different databases. These backup windows can overlap.

Managing SQL Server objects for use with SQL Server Intelligent Policies

This chapter includes the following topics:

- [About the Applications utility](#)
- [About discovery of SQL Server objects](#)
- [About registering SQL Server instances and availability replicas](#)
- [Deleting SQL Server objects from the Applications utility](#)
- [Manually add a SQL Server instance](#)
- [Deactivating or activating an instance](#)
- [Cleaning up instances](#)

About the Applications utility

NetBackup displays the instances and availability groups that it discovers in the **Applications > Microsoft SQL Server** node of the NetBackup Administration Console, along with any instances you add manually. The properties for an instance, replica, or instance group indicate the name of any Intelligent Policies that protect those objects. Legacy policies (that use clients and batch files) are not reflected in the Applications utility. The **Microsoft SQL Server** node contains the following subnodes:

- **All Instances**

Contains all SQL Server instances that NetBackup discovers or that you manually added. Instances that belong to an availability group are also included in this list.

- **Instance Groups**

Contains any instance groups that you created. You can use instance groups to organize instances or replicas and to register all objects in the group with a single set of credentials.

- **Availability Groups**

Displays all SQL Server availability groups that NetBackup discovers.

About discovery of SQL Server objects

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

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

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

By default, this service reports to the primary server when it finds SQL Server instances. However, the user can turn off discovery for a specific client, with the `bpsetconfig` utility. See the `REPORT_CLIENT_DISCOVERIES` option in the [NetBackup Administrator's Guide, Volume I](#).

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

To discover the instances that you created since the last discovery, select **Actions > Discover Instances**. To update an availability group with any new instances, select the availability group and choose **Actions > Rescan Availability Group**.

Discovering 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 Expand **NetBackup Management > Applications > Microsoft SQL Server > Instances**.
- 2 Select **Actions > Discover Instances**.

Discover advanced or basic availability groups on demand

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

To discover advanced or basic availability groups

- 1 Expand **NetBackup Management > Applications** and select **Microsoft SQL Server**.
- 2 Select **Actions > Discover Availability Groups**.
- 3 From the **Instance** list, select a replica that is part of the availability group.
Note that only registered replicas are shown in this list.
- 4 Click **OK**.

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 Expand **NetBackup Management > Applications > Microsoft SQL Server** and select **All Instances**.
- 2 In the right pane, right-click one of the replicas that is part of the availability group and click **Register**.
- 3 Provide the credentials for the replica.

See [“About registering SQL Server instances and availability replicas ”](#) on page 35.

- 4 Expand **NetBackup Management > Applications** and select **Microsoft SQL Server**.
- 5 Select **Actions > Discover Availability Groups**.
- 6 From the **Instance** list, select a replica that is part of the availability group.
 Note that only registered replicas are shown in this list.
- 7 Click **OK**.

About registering SQL Server instances and availability replicas

All instances and availability group replicas that you want protected as part of a SQL Server Intelligent Policy must be registered with credentials. These credentials must have certain privileges.

See [“About credentials used with SQL Server Intelligent Policy”](#) on page 35.

Instances or replicas can be registered in one of the following ways:

- Manually, for individual instances or replicas.
 See [“Registering a SQL Server instance or availability replica”](#) on page 38.
- Manually, by adding instances or replicas to an instance group.
 See [“Registering instances or availability replicas with an instance group”](#) on page 39.
- Automatically, by configuring an instance group to automatically register newly discovered instances or replicas.
 See [“Registering instances or availability replicas automatically”](#) on page 42.
- Manually, with the `nbsqladm` command.

The NetBackup administrator can also authorize a DBA to manually register instances or replicas.

See [“Authorizing a DBA to register instances or availability replicas with the `nbsqladm` command”](#) on page 43.

About credentials used with SQL Server Intelligent Policy

SQL Server instances or replicas must be registered with Windows credentials that have the proper permissions to perform backup and restore operations. Intelligent Policy supports Windows authentication and Windows Active Directory authentication. It does not support Mixed Mode or SQL Server authentication. Credentials are not supported at the database or the availability group level.

Table 4-1 Options to register credentials

Option to register credentials	Environment and configuration
Use these specific credentials (recommended)	■ The SQL Server DBA provides the NetBackup administrator with the SQL Server user credentials. ■ The SQL Server DBA does not want the NetBackup services running as a privileged SQL Server user on the client. Configuration requirements The user account that is used to register credentials must have the SQL Server “sysadmin” role and be a member of the Windows Administrators group. The NetBackup services can use the Local System logon account. If you want to use a different logon account, that account must also have certain local security privileges. See “Configuring the NetBackup services for SQL Server backups and restores” on page 19. See “Configure local security privileges for SQL Server” on page 21.
Use credentials that are defined locally on the client	■ The user account that installed NetBackup is already running as a SQL Server privileged account. ■ The SQL Server DBA does not want to provide credentials to register instances or replicas. ■ The NetBackup administrator does not have access to the SQL Server credentials. ■ You want to use gMSA credentials. Configuration requirements The user account that is used to register credentials must have the SQL Server “sysadmin” role and be a member of the Windows Administrators group. You must also configure the logon account for the NetBackup Client Service and the NetBackup Legacy Network service. See “Configuring the NetBackup services for SQL Server backups and restores” on page 19.
Add to group and register using group credentials	You want to be able to do one or more of the following: ■ Logically group your instances or replicas in some way. ■ Use a particular tuning parameter to improve the performance for each of the instances or replicas in the group. ■ (Optional) Automatically register new instances or replicas and add them to a group. See “Registering instances or availability replicas with an instance group” on page 39.

Table 4-1 Options to register credentials (*continued*)

Option to register credentials	Environment and configuration
Command line	■ The DBA does not have access to the NetBackup Administration Console. ■ The NetBackup administrator does not have the credentials for SQL Server. ■ The DBA wants to maintain the SQL Server credentials independently of the backup administrator. See the section called “Configuring credentials from the command line” on page 37.

Configuring credentials from the command line

To register an instance or replica from the command line, the following configuration is required:

- The NetBackup administrator must authorize the `nbsqladm` command for a specific DBA or user on a specific host.
On the NetBackup primary server, use `nbsqladm` to authorize the user:

```
nbsqladm [-S master_server] -add_dba host_name user_name
```

If you have multiple NICs, authorize the DBA using the private interface name of the SQL Server host.
For a SQL Server cluster, authorize the DBA for each node in the cluster. (Do not authorize a DBA using the virtual name of the SQL Server cluster.) For the `-host name` provide one of the node names in the SQL Server cluster.
For a SQL Server cluster with multiple NICs, authorize the DBA using the private interface name for each of the nodes in the SQL Server cluster.
- Once a DBA is authorized to use the `nbsqladm` command, the DBA can register instances with the local credentials (`-local_credentials`) or other specific credentials (`-user name -domain name`).

For complete details on the `nbsqladm` command, see the *NetBackup Commands Reference Guide*.

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

When NetBackup discovers a SQL Server cluster, it adds a single entry in the Applications utility. This instance represents all nodes in the cluster. The host name is the virtual name of the SQL Server cluster. When you register 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 in the Applications utility. If you installed the NetBackup client using the public interface name, you must configure the NetBackup client name as the private interface name. Then register the instance with its private interface name. For a SQL Server cluster that uses multiple NICs, add and register the instance with the private virtual name of the SQL Server cluster.

See [“Configuring the NetBackup client with the private interface name”](#) on page 178.

Registering Microsoft SQL Server failover cluster instances (FCIs)

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

Validation of credentials

After you add credentials, NetBackup validates the credentials, marks the instances as registered, and adds the instances to the NetBackup database. NetBackup requests detailed information about the instances or replicas from the NetBackup client and displays it in the **Microsoft SQL Server > Instances** or **Microsoft SQL Server > Availability Group** nodes.

For a SQL Server cluster or if an availability group instance is part of SQL Server cluster, NetBackup validates the credentials on the active node. The credentials must be valid for all nodes in the cluster. For a SQL Server availability group, replicas are registered and validated individually. Note that the registered date reflects the date and time the credential was added or updated. It does not indicate if the credentials are valid.

See [“Troubleshooting credential validation ”](#) on page 201.


Registering a SQL Server instance or availability replica

To protect SQL Server with an Intelligent Policy, you must register the SQL Server instances or availability replicas with credentials. Replicas must be registered individually; however, you can use an instance group to all the replicas with the same credentials.

Refer to [Table 4-1](#) to determine the best option for your environment.

To register a SQL Server instance or availability replica

- 1 In the NetBackup Administration Console, in the left pane, expand **NetBackup Management > Applications > Microsoft SQL Server**.
- 2 Select the instances or replicas, as follows:
 - To select instances, click **All Instances**.
The right pane displays a list of instances.
 - To select replicas, click **Availability Groups > Availability group name**.
The right pane displays a list of replicas for the availability group.
- 3 Select the instances or replicas that you want to register. Any instances or replicas that have previously been registered show a date and time in the **Registered** column.



- 4 Choose **Actions > Register**.
- 5 Select the credentials you want to use.

When you register a cluster instance, NetBackup validates the credentials on the active node. The credentials must be valid for all nodes in the cluster.

See [“About credentials used with SQL Server Intelligent Policy”](#) on page 35.
- 6 Click **OK**.

NetBackup requests detailed information about the instances or replicas from the NetBackup client and displays it in **Applications** node.

Registering instances or availability replicas with an instance group

Instance groups provide the following benefits when you create SQL Server policies:

- When you add an instance group to a policy, that single policy can back up many instances or availability replicas.

- You can configure an instance group to automatically add newly discovered instances or replicas to the group, which are then registered on the fly.
 See [“Registering instances or availability replicas automatically”](#) on page 42.
- All the instances or replicas in the group use the same credentials setting. If you select the setting **Use these specific credentials**, you only need to enter those credentials once.
- In the Applications utility, you can easily see which policies protect which instance groups.

To create an instance group

- 1 In the NetBackup Administration Console, in the left pane, expand **NetBackup Management > Applications > Microsoft SQL Server**.
- 2 Right-click **Instance Groups** and select **New Instance Group**.
- 3 Provide an **Instance Group Name**.
- 4 Select the credentials you want to use.
 This user account must have certain privileges. More information is available to help determine which option best applies for your environment.
 See [“About credentials used with SQL Server Intelligent Policy ”](#) on page 35.
- 5 Click **OK**.
- 6 To add instances or replicas to the group you created, see the following topic.
 See [“Adding an instance or availability replica to an instance group”](#) on page 40.


Adding an instance or availability replica to an instance group

Use an instance group to apply the same credentials to all instances or availability replicas in the group.

To add an instance or availability replica to an instance group

- 1 In the NetBackup Administration Console, in the left pane, expand **NetBackup Management > Applications > Microsoft SQL Server**.
- 2 Add instances or replicas, as follows:
 - To add instances, click **All Instances**. Then select one or more instances that you want to add to an instance group.
 - To add replicas, click **Availability Groups > Availability group name**. Then select one or more replicas that you want to add to an instance group.

For a SQL Server cluster, NetBackup adds a single entry or one instance to the Applications utility. The host name for that instance is the virtual name of the SQL Server cluster.



- 3 From the **Actions** menu, select **Register**.
- 4 Click **Add to group and register using group credentials**.
- 5 From the **Instance Group** list, select the instance group to which you want to add the instances or replicas.
- 6 Click **OK**.

NetBackup requests detailed information about the instances or replicas from the NetBackup client and displays it in **Applications** node.


If you previously registered an instance or replica, its credentials are automatically changed to the group credentials setting.

Validating instance group credentials

Credentials for an instance group are not validated when an instance or a replica is registered automatically. You should periodically validate the credentials for the objects in the group. For a cluster instance, note that NetBackup validates the credentials on the active node. The credentials must be valid for all nodes in the cluster.

To validate group credentials

- 1 Select the instance group.
- 2 Choose **Actions > Properties**.
- 3 Click **Validate**.



NetBackup requests detailed information about the instances or replicas from the NetBackup client and displays it in **Applications** node.

Registering instances or availability replicas automatically

With automatic registration, NetBackup adds newly discovered instances or availability replicas to the instance group that you choose and automatically registers them with group credentials.. Only one instance group can be configured for automatic registration.

Note: Any instances or replicas that were discovered before this instance group was created are not automatically added to the group.

To register instances or availability replicas automatically

- 1 In the NetBackup Administration Console, in the left pane, expand **NetBackup Management > Applications**.
- 2 If necessary, create an instance group.
- 3 Click **Microsoft SQL Server** and choose **Actions > Auto Registration**.
- 4 Select **Automatically register newly discovered instances**.

- 5 From the **Instance Group** list, select the instance group to which you want to add newly discovered instances or replicas.

- 6 Click **OK**.

To validate the credentials for the instances or replicas in the group, see the following topic.

See [“Validating instance group credentials”](#) on page 41.

Authorizing a DBA to register instances or availability replicas with the `nbsqladm` command

The NetBackup administrator can authorize a DBA to use the `nbsqladm` to register instances or replicas if the DBA wants to manage SQL Server credentials independently. From the primary server the NetBackup administrator can control the list of users and hosts that can run `nbsqladm` on the NetBackup client.

For example, the NetBackup administrator can authorize the user `john_smith` on host `winserver.domain.com` with the following command:

```
nbsqladm -add_dba winserver.domain.com john_smith
```

From the NetBackup client, `winserver.domain.com`, `john_smith` can register and manage instances or replicas. For example, the DBA can register an instance with local credentials as follows:

```
nbsqladm -S NBUMaster1 -register_instance hr_city1  
- host winserver.domain.com -local_credentials
```

More information on the `nbsqladm` command is available. See the [NetBackup Commands Guide](#).

Deleting SQL Server objects from the Applications utility

You cannot delete a SQL Server object (for example, an instance) that is part of a policy. First, delete the object from the **Instances and Databases** tab in the policy.

To delete SQL Server objects from the Applications utility

- 1 Expand **NetBackup Management > Applications > Microsoft SQL Server**.
- 2 Choose from the following:
 - Select the **Availability Groups** node. In the right pane, select the availability group that you want to delete.


- Expand **Availability Groups** and select the availability group. In the right pane, select the replica that you want to delete.
 - Select the **Instance Groups** node. In the right pane, select the availability group that you want to delete.
 - Select the **Instances** node. In the right pane, select the availability group that you want to delete.
- 3 Select **Actions > Delete**.

Manually add a SQL Server instance

Newly discovered SQL Server instances on clients are automatically added to the NetBackup database. 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 In the NetBackup Administration Console, in the left pane, expand **NetBackup Management > Applications > Microsoft SQL Server**.
- 2 Right-click **All Instances** and select **New instance**.
- 3 Provide the **Host** name where the instance resides and the **Instance name**.
 - For a SQL Server cluster or multi-NIC environment, add one entry to the Applications utility.
 - For a cluster, the host name is the virtual name of the SQL Server cluster. You do not need to add each node in the cluster.
 - For a multi-NIC environment, the host name is the private interface name of the SQL Server host or of the virtual SQL Server.
 - For a failover cluster instance, enter the virtual name of the SQL Server cluster.
NetBackup enumerates the FCI under the physical node names and the cluster name.



- 4 Click **Edit** to provide credentials and register the instance.

See [“Registering a SQL Server instance or availability replica”](#) on page 38.

You may omit credentials when you add a new instance to the NetBackup database. The instance is marked as unregistered and the **Registered** column is empty.

Deactivating or activating 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 Expand **NetBackup Management > Applications > Microsoft SQL Server**.
- 2 Select the **Instances** node.
- 3 In the right pane, select the instance that you want to deactivate.
- 4 From the **Actions** menu, select **Deactivate**.

To activate an instance

- 1 Expand **NetBackup Management > Applications > Microsoft SQL Server**.
- 2 Select the **Instances** node.
- 3 In the right pane, select the instance that you want to activate.
- 4 From the **Actions** menu, select **Activate**.

Cleaning up instances

This option lets you configure NetBackup to automatically clear orphaned instances from the Applications utility. Orphaned instances are the instances that were discovered at one time but were never registered.

To enable instance cleanup

- 1 Expand **NetBackup Management > Applications > Microsoft SQL Server**.
- 2 Select the **Instances** node.
- 3 Select **Actions > Instance cleanup**.
- 4 Select **Clean up after**.
- 5 Select how often (**days**) that you want NetBackup to perform instance cleanup.
- 6 Click **OK**.

Configuring backups with SQL Server Intelligent Policy

This chapter includes the following topics:

- [About SQL Server Intelligent Policies](#)
- [Creating a SQL Server Intelligent Policy](#)
- [About policy attributes](#)
- [About schedule properties](#)
- [Schedule backup types for SQL Server Intelligent Policies](#)
- [Adding instances to a policy](#)
- [Adding databases to a policy](#)
- [Adding filegroups or files to the backup selections list](#)
- [Manually adding files or filegroups to the backup selections list](#)
- [Adding instance groups to a backup policy](#)
- [About tuning parameters for SQL Server backups](#)
- [Backing up read-only filegroups](#)
- [Backing up read-write filegroups](#)

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

Creating a SQL Server Intelligent Policy

Before you configure an intelligent policy ensure that you have:

- Registered the SQL Server instances that you want to protect.
See [“Registering a SQL Server instance or availability replica”](#) on page 38.
- For SQL Server clusters or availability groups, you configured the mappings for distributed application restores. Also, you reviewed 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 27.
See [“Reviewing the auto-discovered mappings in Host Management”](#) on page 22.

To create a NetBackup for SQL Server Intelligent Policy

- 1 Log on to the primary server as administrator (Windows) or root (UNIX).
- 2 Start the NetBackup Administration Console.
- 3 Expand **NetBackup Management** and select **Policies**.
- 4 Select **Actions > New > Policy**.
- 5 Type a unique name for the new policy and click **OK**.
- 6 In the **Policy type** list, select **MS-SQL-Server**.

- 7 Complete the entries on the **Attributes** tab.
See [“About policy attributes”](#) on page 49.
- 8 Add other policy information as follows:
 - Choose to protect instances or instance groups.
If you choose the instances option, you can select either individual instances or databases.
See [“Adding instances to a policy”](#) on page 53.
See [“Adding databases to a policy”](#) on page 55.
See [“Adding instance groups to a backup policy”](#) on page 59.
 - Add schedules.
See [“About schedule properties”](#) on page 50.
 - (Optional) Select the specific filegroups or files that you want to back up.
By default, NetBackup backs up an entire database.
See [“Adding filegroups or files to the backup selections list”](#) on page 57.
 - (Optional) Make changes to any tuning parameters.
See [“About tuning parameters for SQL Server backups”](#) on page 60.
- 9 When you have completed the policy configuration, click **OK**.
In the **Applications** utility, the properties for an instance, replica, or instance group indicate the name of any Intelligent Policies that protect those objects.

About policy attributes

With a few exceptions, NetBackup manages the policy attributes set for a database backup like a file system backup. Other policy attributes vary according to your specific backup strategy and system configuration.

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

Table 5-1 Policy attribute 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.

Table 5-1 Policy attribute for NetBackup for SQL Server policies (*continued*)

Attribute	Description
Compress	Enables the compression of backups by NetBackup. If you enable NetBackup compression, do not enable SQL Server compression. For more information on advantages and disadvantages of compression, see the NetBackup Administrator's Guide, Volume I .
Keyword phrase	Although you can create a keyword phrase for MS-SQL-Server policies, NetBackup for SQL Server does not record this information with the backup image.
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 99. See “ Configuring a VMware policy to protect SQL Server using Replication Director to manage snapshot replication ” on page 96.

About schedule properties

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 5-2](#) describes how the schedule properties affect a SQL Server Intelligent Policy.

Table 5-2 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 51.

Table 5-2 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, Veritas 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 210.

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 53.
- 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 [“About tuning parameters for SQL Server backups”](#) on page 60.

[Table 5-3](#) shows the backup types that you can specify.

Table 5-3 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 52. See the section called “Configuring high-frequency transaction log backups” on page 52. 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 29.

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 along with full or differential backups, the transaction log backups run at the scheduled time and frequency 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 29.

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 are protected by the policy. The backup completes with a status 0 and the job details indicate that the database was skipped.

Example backup schedules for a policy

[Table 5-4](#) shows an example of the schedules you can create for a single SQL Server Intelligent Policy.

Table 5-4 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


Adding 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 [“Adding databases to a policy”](#) on page 55.

To add instances a policy

- 1 On the **Instances and Databases** tab, click **Protect instances**.
- 2 Click **New**.
 All instances that you registered are displayed.
- 3 In the left pane, select the **Instances** node.
- 4 In the right pane, check 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.



- 5 Click **OK**.

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.



Adding databases to a policy

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

See [“Adding instances to a policy”](#) on page 53.

To add databases to a policy


- 1 On the **Instances and Databases** tab, click **Protect instances**.
- 2 Click **New**.

All instances that you registered are displayed.


- 3 In the left pane, expand the **Databases** node and select the instance that contains the databases that you want to protect.

- 4 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 databases that are hosted on a SQL Server cluster, the **Host Name** represents the virtual name of the SQL Server. (See the following figure.)



- 5 Click **OK**.

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.




Adding 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 Open the policy you want to edit or create a policy.
- 2 Select **Filegroups** or **Files**.
- 3 Click **Browse**.
- 4 In the left pane, select an instance to view the filegroups or files that it contains.

- 5 In the right pane select the filegroups or files.



- 6 Click **OK** 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.



Manually adding files or filegroups to the backup selections list

This topic describes how to manually add SQL Server database filegroups or files to the backup selections list.

To manually add files or filegroups to the backup selections list

- 1 Open the policy you want to edit or create a new policy.
- 2 Click the **Backup Selections** tab.
- 3 Select one of the SQL Server object types to back up:
 - **Filegroups**
 - **Files**
- 4 Click **New**.
- 5 Type the name of a filegroup or file and then click **Add**.
Repeat this step to add any other filegroups or files.
- 6 Click **OK > OK**.

Adding instance groups to a backup policy

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

To add instance groups to a SQL Server Intelligent policy

- 1 Open the policy you want to edit or create a new policy.
- 2 On the **Instances and Databases** tab, click **Protect instance groups**.
- 3 Click **New**.

All instance groups that you registered are displayed.

- 4 Select the instance groups you want to add and click **OK**.

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 click **Preview Instances**.

About tuning parameters for SQL Server backups

The **Microsoft SQL Server** tab 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.

See [“NetBackup for SQL Server performance factors”](#) on page 195.

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 5-5 Tuning parameters for SQL Server backups

Field	Description
Number of backup stripes	This option divides the backup operation into multiple concurrent streams. A stream corresponds to a job in the activity monitor. For example, if the value is 3, each database is backed up using three jobs. This configuration applies in any situation in which SQL Server dumps data faster than your tape drive is capable of writing. The default value for this option is 1. Range is 1–32. See “Configuring multistriped backups of SQL Server” on page 64.
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 master server. The default value for this option is 2, which allows double buffering. You may get slightly better performance by increasing this value to a higher value. Range is 1–32.
Maximum transfer size	(Stream-based backups only) This option is the buffer size used by SQL Server for reading and writing backup images. Generally, you can get better SQL Server performance by using a larger value. This option can be set for each backup operation. Calculated as 64 KB * 2^{MAX_TRANSFER_SIZE}. It ranges in size from 64 KB to 4 MB. The default is 4 MB.
Backup block size	This option applies to stream-based backups only. Sets the incremental size that SQL Server uses for reading and writing backup images and can be set for each backup operation. Calculated as 512 bytes * 2^{BLOCK_SIZE}. The value for this option ranges from 0.5 KB to 64 KB. The default is 64 KB.

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

Field	Description
Parallel backup operations	This option is the number of backup operations to start simultaneously, per database instance. Range is 1–32. The default is 1. You may need to configure other options when you configure two or more parallel backup operations. See “Configure the number of jobs allowed for backup operations” on page 29.
Microsoft SQL Server checksum	Choose one of the following options for SQL Server backup checksums: ■ None. Disables backup checksums. ■ To verify the checksums before the backup, choose one of the following options. Note that these options impose a performance penalty on a backup or restore operation. ■ Continue on error. If the backup encounters a verification error, the backup continues. ■ Fail on error. If the backup encounters a verification error, the backup stops.
Use Microsoft SQL Server compression	Enable this option to use SQL Server to compress the backup image. If you enable SQL Server compression, do not enable NetBackup compression. SQL Server compression is not supported for snapshot backups.
Skip unavailable (offline, restoring, etc.) databases	NetBackup skips any database with a status that prevents NetBackup from successfully backing up the database. These statuses include offline, restoring, recovering, and emergency mode, etc. NetBackup skips the backup of the unavailable database, but continues with the backup of the other databases that the policy includes. The backup completes with a status 0 and the job details indicate that the database was skipped. See “Schedule backup types for SQL Server Intelligent Policies” on page 51.
Copy-only backup	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. See “Using copy-only snapshot backups to affect how differentials are based” on page 108.
Skip read-only file groups	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. See “Backing up read-only filegroups” on page 64. See “Backing up read-write filegroups” on page 65.

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

Field	Description
Convert differential backups to full (when no full exists)	If no previous full backup exists for the database or filegroup, then NetBackup converts a differential backup to a full backup. 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 as follows: ■ If you select a database for a differential backup, the backup is converted to a full database backup. If the Skip read-only file groups option is selected the backup is converted to a full read/write filegroup backup. ■ 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 MS SQL Client. Or, if you expired the backup in NetBackup, you can import the full backups into the NetBackup catalog. Then you can restore both the full and the differential backups with the NetBackup MS SQL Client.
Truncate logs after backup	This option backs up the transaction log and removes the inactive part of the transaction log. This option is enabled by default.

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

Field	Description
Convert log backups to full (when no full exists)	If no previous full backup exists for the database, then NetBackup converts a transaction backup to a full backup. This option also detects if a full recovery database was switched to the simple recovery model and back to the full recovery model. In this scenario, the log chain is broken and SQL Server requires a differential backup before a subsequent log backup can be created. If NetBackup detects this situation, the backup is converted to a differential database backup. Note: NetBackup only converts a transaction log backup if a full backup was never performed on the database. If a full backup does not exist in the NetBackup catalog but SQL Server detects an existing full LSN, NetBackup performs a transaction log backup and not a full. In this situation, you can restore the full backup with native tools and any differentials and log backups with the NetBackup MS SQL Client. Or, if the backup 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 MS SQL Client.
Availability Database Backup Preference	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 and instance groups. ■ None 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.
VDI Timeout (seconds)	Determines the timeout 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 restores is 600. Range is 300–2147483647.

Configuring 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 85.

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 legacy 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 legacy SQL Server 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.

Backing 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 [“Viewing SQL Server read-only backup sets”](#) on page 169.

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 49.
- 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 [“About schedule properties”](#) on page 50.
- 4 Choose to protect instances or instance groups.
See [“Adding instances to a policy”](#) on page 53.
See [“Adding instance groups to a backup policy”](#) on page 59.
- 5 On the **Backup Selections** tab, select **Filegroups**.
See [“Adding filegroups or files to the backup selections list”](#) on page 57.
- 6 Select the filegroups you want to back up.
- 7 When you complete the policy configuration, click **OK**.
- 8 Back up the read-only filegroups.
- 9 If necessary, confirm all read-only groups are backed up by viewing the read-only backup set.
See [“Viewing SQL Server read-only backup sets”](#) on page 169.

Backing 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 [“Backing up read-only filegroups”](#) on page 64.

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 49.
- 3** Create a **Full Backup**, **Differential Incremental Backup**, and **Transaction Log backup** schedule.
See [“About schedule properties”](#) on page 50.
- 4** On the **Instances and Databases** tab, choose to **Protect instances**.
- 5** Add the instances or the databases that contain the read-write filegroups.
See [“Adding instances to a policy”](#) on page 53.
- 6** On the **Backup Selections** tab, select **Whole database**.
- 7** Click the **Microsoft SQL Server** tab.
- 8** Check **Skip read-only file groups**.
See [“About tuning parameters for SQL Server backups”](#) on page 60.
- 9** When you have completed the policy configuration, click **OK**.

Performing restores of SQL Server

This chapter includes the following topics:

- [Starting the NetBackup MS SQL Client for the first time](#)
- [Selecting the SQL Server host and instance](#)
- [Browsing for SQL Server backup images](#)
- [Options for NetBackup for SQL Server restores](#)
- [Restoring a SQL Server database backup](#)
- [Staging a full SQL Server database recovery](#)
- [Restoring SQL Server filegroup backups](#)
- [Recovering a SQL Server database from read-write filegroup backups](#)
- [Restoring SQL Server read-only filegroups](#)
- [Restoring SQL Server database files](#)
- [Restoring a SQL Server transaction log image without staging a full recovery](#)
- [Performing a SQL Server database move](#)
- [About performing a SQL Server page-level restore](#)
- [Configuring permissions for redirected restores](#)
- [Redirecting a SQL Server database to a different host](#)
- [Performing a restore of a remote SQL Server installation](#)

- [Restoring multistreamed SQL Server backups](#)
- [About using bplist to retrieve SQL Server backups](#)
- [About NetBackup for SQL Server backup names](#)

Starting 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 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**.

Selecting the SQL Server host and instance

Use this procedure to set which SQL Server host and the instance that you want the NetBackup MS SQL Client to access.

(Legacy SQL Server policies) 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** drop-down list, select the SQL Server host.

You can type a host name if it does not appear in the drop-down list. If you select a remote host and click **Apply**, the **Host type** is shown as "remote".

- 4 From the **Instance** drop-down list, select the SQL Server instance.

You can type an instance name if it does not appear in the drop-down list. You can designate the default instance either by setting the Instance box to <default> or to empty (no spaces).

- 5 Click **Apply** > Close.

Browsing for SQL Server backup images

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, 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 70.

To browse for backup images

- 1 Change the host and instance you want to access.

See [“Selecting the SQL Server host and instance”](#) on page 68.

- 2 Select **File > Restore SQL Server objects**.

- 3 Select the **SQL Host** whose backup images you want to browse, or type its name.

- 4 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.

- 5 (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.

- 6 Select the date range to search and click **OK**.
- 7 Continue with the applicable instructions for how to restore the objects.
See [“Restoring a SQL Server database backup”](#) on page 73.
See [“Staging a full SQL Server database recovery”](#) on page 74.
See [“Restoring SQL Server filegroup backups”](#) on page 74.
See [“Recovering a SQL Server database from read-write filegroup backups”](#) on page 75.
See [“Restoring SQL Server read-only filegroups”](#) on page 76.
See [“Restoring SQL Server database files”](#) on page 76.
See [“Restoring a SQL Server transaction log image without staging a full recovery”](#) on page 77.
See [“Performing a SQL Server database move”](#) on page 77.
See [“About performing a SQL Server page-level restore”](#) on page 79.

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 [“Performing restores of SQL Server when you have multiple NICs”](#) on page 181.
- 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

Options for NetBackup for SQL Server restores

Table 6-1 describes the options that are available when you perform restores.

Table 6-1 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.
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.

Table 6-1 Options for restore operations (*continued*)

Option	Description
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.
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.

Table 6-1 Options for restore operations (*continued*)

Option	Description
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.
Restore	Start the restore or generate a restore script.

Restoring a SQL Server database backup

This topic describes how to restore a database from a full database or differential database backup.

To restore a database backup

- 1 Browse for the backup images you want to restore.
See [“Browsing for SQL Server backup images”](#) on page 69.
- 2 Expand the database instance and the database.

- 3 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.
- 4 Select the restore options.
See [“Options for NetBackup for SQL Server restores”](#) on page 71.
- 5 Click **Restore**.

Staging a full SQL Server database recovery

This topic describes how to stage a full database recovery. Alternatively, you can restore without staging a full recovery.

See [“Restoring a SQL Server transaction log image without staging a full recovery”](#) on page 77.

To stage a full database recovery

- 1 Browse for a backup image that contains the point in time to which you want to recover.
See [“Browsing for SQL Server backup images”](#) on page 69.
- 2 Expand the database instance.
- 3 Click the "+" next to the database that contains the transaction log backup you want to restore.
- 4 Select the transaction log image that includes the point in time from which you want to recover.
- 5 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.

- 6 Click **Restore**.

Restoring SQL Server filegroup backups

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 [“Recovering a SQL Server database from read-write filegroup backups”](#) on page 75.

See [“Restoring SQL Server read-only filegroups”](#) on page 76.

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 [“Staging a full SQL Server database recovery”](#) on page 74.

To restore a filegroup backup

- 1 Browse for the backup images you want to restore.
 See [“Browsing for SQL Server backup images”](#) on page 69.
- 2 Expand the database instance and database.
- 3 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.
- 4 Click **Restore**.

Recovering a SQL Server database from read-write filegroup backups

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 Browse for the backup images you want to restore.
 See [“Browsing for SQL Server backup images”](#) on page 69.
- 2 Expand the database instance.

- 3 Expand the database that contains the read-write filegroups you want to restore.
- 4 Select the transaction log backup.
- 5 Right-click the transaction log backup and select **Properties**.
- 6 On the **Recovery set** tab, verify that a complete backup set is available and click **OK**.
- 7 Click **Restore**.

See [“Restoring SQL Server read-only filegroups”](#) on page 76.

Restoring SQL Server read-only filegroups

This topic describes how to restore read-only filegroups.

To restore read-only filegroups

- 1 Browse for the backup images you want to restore.
See [“Browsing for SQL Server backup images”](#) on page 69.
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.
- 2 Expand the database instance.
- 3 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.
- 4 Click **Restore**.

Restoring SQL Server database files

This topic describes how to restore database files.

To restore a database file

- 1 Browse for the backup images you want to restore.
See [“Browsing for SQL Server backup images”](#) on page 69.
- 2 Expand the database instance and the database.
- 3 Expand the filegroup and the file.
- 4 Select the database file image that you want to restore.
- 5 Click **Restore**.

Restoring a SQL Server transaction log image without staging a full recovery

This topic describes how to restore a transaction log image without staging a full recovery. Alternatively, you can stage a full recovery.

See [“Staging a full SQL Server database recovery”](#) on page 74.

To restore a transaction log without staging a full recovery

- 1 Browse for the backup images you want to restore.
See [“Browsing for SQL Server backup images”](#) on page 69.
- 2 Expand the database instance.
- 3 Select the transaction log image that you want to restore.
- 4 Select **Restore selected transaction log**.
- 5 Click **Restore**.

Performing a SQL Server database move

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 Browse for the backup images you want to restore.
See [“Browsing for SQL Server backup images”](#) on page 69.
- 2 Expand the database instance.
- 3 Select the database backup image that you want to restore.
- 4 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.

- 5 Click **Restore**.

- 6 Indicate a file name and click **Save > Yes**.
- 7 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"
```

- 8 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"
```

9 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"
```

- 10 Make similar changes to the template for any differential backups or transaction log backups you want to move.
- 11 When you finish modifying the template, save it.
- 12 To run the restore, select **File > Manage script files**.
- 13 Select the script that you created and click **Start > Yes**.

About performing a SQL Server page-level restore

Note: Page-level restores are only applicable for SQL Server legacy backup 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.

Performing SQL Server page-level restores

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 page-level restore

- 1 Obtain a list of corrupt pages in the database.
- 2 Browse for the backup images you want to restore.
 See [“Browsing for SQL Server backup images”](#) on page 69.
- 3 Expand the database instance and the database.
- 4 Select the database backup image that contains pages you want to restore.
- 5 From the **Scripting** list, select **Create a page restore template**.
- 6 Click **Restore**.
- 7 Type a file name for the page restore script and click **Save > Yes**.

- 8 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
```

- 9 When you finish modifying the template, save it.
- 10 Select the script you created and click **Start > Yes**.

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 Administrator's Guide, Volume I](#) 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 `svgrp`, 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 configure the mappings for distributed application restores. You also need to review the auto-discovered mappings for the hosts in your environment.

See [“Reviewing the auto-discovered mappings in Host Management”](#) on page 22.

See [“Configuring mappings for restores of a distributed applications, clusters, or virtual machines”](#) on page 27.

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 `SQLHOST-NB`.

On Windows:

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

On UNIX:

```
/usr/opensv/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.

Redirecting a SQL Server database to a different host

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 81.
- 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 84.
- 3 Select **File > Set SQL Server connection properties**.
- 4 From the **Host** list, select the host you want to restore to.
- 5 From the **Instance** list, select the database instance.
To select the default instance, either select **<default>** or leave the field empty.
- 6 Click **Apply** and then **Close**.
- 7 Select **File > Set NetBackup client properties**.
- 8 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 84.
- 9 Click **OK**.
- 10 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”](#) on page 69.

- 11 Browse for the database that you want to move.
- 12 From the **Scripting** list, select **Create a move template**.
- 13 Click **Restore**.
- 14 Enter a file name and click **Save > Yes**.
- 15 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, you can select the current primary server by selecting **File > Set NetBackup client properties**.

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 [“Performing a SQL Server database move”](#) on page 77.

Performing a restore of a remote SQL Server installation

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 Select the host and instance you want to access.
See [“Selecting the SQL Server host and instance”](#) on page 68.
- 2 Browse for the backup images you want to restore.
See [“Browsing for SQL Server backup images”](#) on page 69.

- 3 Select the options for the operation.

See [“Options for NetBackup for SQL Server restores”](#) on page 71.

Save is enabled in the restore dialog box. **Launch immediately** is disabled because the generated script must be executed on the remote host that you are logged on to.

- 4 Click **Restore**.
- 5 Navigate to the `install_path\NetBackup\DbExt\MsSql\` folder on the remote host, and save the batch file there.
- 6 Run the operation from the local installation of NetBackup for SQL Server.

Restoring multistreamed SQL Server backups

When you use the NetBackup MS SQL Client, backups using multiple stripes are automatically restored 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 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 86.

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 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 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.JUNE BERRY.db.pubs.~.7.001of003.20140920101716..C:\
juneberry.MSSQL7.JUNE BERRY.db.pubs.~.7.002of003.20140920101716..C:\
juneberry.MSSQL7.JUNE BERRY.db.pubs.~.7.003of003.20140920101716..C:\
juneberry.MSSQL7.JUNE BERRY.fil.pubs.pubsnew.7.001of001.20140919175149..C:\
juneberry.MSSQL7.JUNE BERRY\NEWINSTANCE.trx.abc.~.7.001of001.20140902170920..C:\
juneberry.MSSQL7.JUNE BERRY\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 87.

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 6-1 Backup image name for a database filegroup





Figure 6-2 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>.

Protecting SQL Server data 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 the NetBackup services for a VMware backup that protects 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](#)
- [Restore SQL Server databases from a VMware backup](#)

About protecting an application database with VMware backups

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

VMware application backups let you:

- 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.

- Restore and recover databases from VMware backups to alternate clients. The target destination client can be a physical computer or a virtual machine.

Supported environments and configuration

See the following information on virtual systems compatibility:

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

Veritas VSS provider

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

See “[Installing the Veritas VSS provider for vSphere](#)” on page 18.

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

Using NetBackup Accelerator to increase speed of full VMware backups

Select the **Use Accelerator** option to use NetBackup Accelerator to potentially increase the speed of full VMware backups. 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](#).

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 an MS-SQL-Server policy for SQL Server incremental backups.

- SQL Server clusters or SQL Server availability groups.
- 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.

About configuring NetBackup for VMware backups that protect SQL Server

Table 7-1 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 “Configuring the NetBackup services for SQL Server backups and restores” on page 19. See “Configure local security privileges for SQL Server” on page 21.
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 94. See “Configuring a VMware policy to protect SQL Server using Replication Director to manage snapshot replication” on page 96.

Table 7-1 Steps to configure VMware backups that protect SQL Server
(continued)

Step	Action	Description
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 27.
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 in Host Management” on page 22.

Configuring the NetBackup services for a VMware backup that protects SQL Server

NetBackup uses the NetBackup Client Service and the NetBackup Legacy Network Service to access the SQL Server when it performs VMware backups and restores. The logon account for the NetBackup services must meet the following requirements:

- You must use an account other the Local System account as the logon account. Both services must use the same logon account. To change the logon account, you must have administrator group privileges.
- For VMware backups with Replication Director, the account has access to the CIFS shares on the NetApp disk array.
- 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.
- 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 configure the NetBackup services for a VMware backup that protects SQL Server

- 1 Log on to the Windows host with the account that necessary role and privileges.
- 2 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 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** service.
- 6 Configure the same logon account for this service as you did for the NetBackup Client Service.
- 7 Stop and start the services.
- 8 Configure the services for each host that you use to browse for backups and perform restores.

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:

- Configure all storage options.
- Configure the logon account for the NetBackup services.

See [“Configuring the NetBackup services for SQL Server backups and restores”](#) on page 19.

See [“Configure local security privileges for SQL Server”](#) on page 21.

- Review the auto-discovered mappings for the hosts in your environment.

More information on Accelerator is available:

See [“About policy attributes”](#) on page 49.

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

To configure a VMware backup policy to protect SQL Server

- 1** Create a new policy or open the policy you want to configure.
- 2** Click 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, click **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.
- 3** On the **Schedules** tab, create a schedule for full backups.
- 4** On the **Clients** tab, do the following:
 - Click **Select automatically through query**.
 - Select **NetBackup host to perform automatic virtual machine selection** and the host you want to use.
 - Use the Query Builder to create the rules that select the virtual machines you want to back up.
- 5** Select the **VMware** tab:
 - Select the **Primary VM identifier** to use to catalog the backups.
 - Select **Enable file recovery from VM backup**.
 - Select **Enable SQL Recovery**.

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.
- 6 If you want to exclude certain disks from the VMware backup, click 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.
 - 7 Click **OK** 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 1.
 - 9 Click the **VMware** tab and under **Enable SQL Recovery**, select **Truncate logs**.

For SQL Server, this option truncates the transaction logs when the VMware snapshot of the virtual machine is complete.
 - 10 Click **OK** to save the policy.
 - 11 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** Log on to the primary server as administrator.
- 2** Start the NetBackup Administration Console.
- 3** Create a policy or open the policy you want to configure.
- 4** 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**.
- 5** On the **Schedules** tab, create a schedule for full backups.
- 6** 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.
- 7** 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.
 - **Enable SQL Recovery**.
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.
- 8** Click **OK** 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 Click the **VMware** tab and under **Enable SQL Recovery**, select **Truncate logs**.
- 11 Click **OK** to save the policy.
- 12 Perform a full VMware backup.

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 Open the NetBackup MS SQL Client.
- 2 Browse for the backup images you want to restore.
See [“Browsing for SQL Server backup images”](#) on page 69.
- 3 Expand the database instance and the database.
- 4 Select the database image that you want to restore.
Only the **Recovered** recovery option is available for VMware backups of SQL Server.
- 5 Click **Restore**.

Configuring backups with Snapshot Client

This chapter includes the following topics:

- [About NetBackup Snapshot Client for SQL Server](#)
- [How SQL Server operations use Snapshot Client](#)
- [Configuration requirements for SQL Server snapshot and Instant Recovery backups](#)
- [Configuring a snapshot policy for SQL Server](#)
- [Configuring 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 backups \(legacy SQL Server policies\)](#)

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.

Although all of these features are provided through Snapshot Client support for SQL Server, not all snapshot methods are supported. For information on how to select a method, see the [NetBackup Snapshot Client Administrator's Guide](#). For a description of snapshot methods available for use with NetBackup for SQL Server, see the NetBackup Snapshot Client [compatibility list](#).

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 195.


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. To determine what backup method and policy were used when a SQL Server backup was created, right-click the backup image and select **Properties**.

Figure 8-1 Backup method that appears in the backup image properties



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 Snapshot Client Administrator's Guide](#) for details on the hardware requirements and software requirements for the snapshot method that you want to use.
- Go to the Veritas Support website for details on the snapshot methods and platforms that are supported for NetBackup for SQL Server.
- The volume(s) which contains the SQL Server databases and log files should be dedicated to SQL Server only. Other types of databases (e.g., Exchange) should not reside on the volume(s).

- NetBackup Snapshot Client is installed and configured correctly and you have a the license for this option. See the [NetBackup Snapshot Client Administrator's Guide](#) for details.
- 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.

Configuring 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 48.

See [“Adding a new SQL Server legacy policy”](#) on page 153.

To configure a snapshot policy for SQL Server

- 1 For SQL Server legacy policies, create a backup script (.bch file) using the NetBackup MS SQL Client.
- 2 Open the policy you want to configure.
- 3 Click 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, click **Snapshot Client Options**. Then from the **Snapshot method** list, select **auto**.

- To use a specific snapshot method, click **Snapshot Client Options**. From the **Snapshot method** list, select the method you want to use for this policy.

See the [NetBackup Snapshot Client Administrator's Guide](#) for details on how to select the snapshot method and automatic snapshot selection.

- 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. See the [NetBackup Snapshot Client Administrator's Guide](#).
 - Select **Perform off-host backup**.
 - Click **Use alternate client** and enter the name of the alternate client.

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 **Protect Instances** or **Protect instance groups**.
 If you choose the instances option, you can select either individual instances or databases.
 See ["Adding instances to a policy"](#) on page 53.
 See ["Adding databases to a policy"](#) on page 55.
 See ["Adding instance groups to a backup policy"](#) on page 59.
 - (SQL Server legacy policies) Choose **Clients for use with batch files**.
- 10 (SQL Server Intelligent Policy) Add other policy information as follows:
 - Add schedules.
 See ["About schedule properties"](#) on page 50.
 - (Optional) Select specific filegroups or files that you want to back up. By default, NetBackup backs up an entire database.
 See ["Adding filegroups or files to the backup selections list"](#) on page 57.
 - (Optional) Make changes to any tuning parameters.
 See ["About tuning parameters for SQL Server backups"](#) on page 60.
- 11 (SQL Server legacy policies) Add other policy information as follows:
 - Add schedules.
 See ["About schedule properties"](#) on page 154.

- Add clients.
See [“Adding clients to a policy”](#) on page 159.
- Add batch files to the backup selections list.
See [“Adding batch files to the backup selections list ”](#) on page 160.

12 Click **OK** to save the policy.

Configuring 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 48.

See [“Adding a new SQL Server legacy policy”](#) on page 153.

To configure a policy for Instant Recovery

- 1 For SQL Server legacy policies, create a backup script using the NetBackup MS SQL Client interface.
- 2 Open the policy you want to configure.
- 3 Click 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 Click **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.
- To use a specific snapshot method, click **Snapshot Client Options** and select it from the **Snapshot method** list.

See the [NetBackup Snapshot Client Administrator's Guide](#) for details on how to select the snapshot method and automatic snapshot selection.

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 **Protect Instances** or **Protect instance groups**.
 If you choose the instances option, you can select either individual instances or databases.
 See ["Adding instances to a policy"](#) on page 53.
 See ["Adding databases to a policy"](#) on page 55.
 See ["Adding instance groups to a backup policy"](#) on page 59.
- (SQL Server legacy policies) Choose **Clients for use with batch files**.

10 To configure schedules, click the **Schedules** tab.

- (SQL Server Intelligent Policies) Configure a full backup schedule.
 See ["About schedule properties"](#) on page 50.
- (Legacy policies) Follow the instructions to configure an Application and a full backup schedule.
 See ["About schedule properties "](#) on page 154.

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 (legacy 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 specific filegroups or files that you want to back up. By default, NetBackup backs up an entire database.
 See [“Adding filegroups or files to the backup selections list”](#) on page 57.
 - (Optional) Make changes to any tuning parameters.
 See [“About tuning parameters for SQL Server backups”](#) on page 60.
- 13** (SQL Server legacy policies) Add other policy information as follows:
- Add clients.
 See [“Adding clients to a policy”](#) on page 159.
 - Add batch files to the backup selections list.
 See [“Adding batch files to the backup selections list”](#) on page 160.
- 14** Click **OK** 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.

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.

Creating a copy-only backup (legacy SQL Server policies)

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

See [“About tuning parameters for SQL Server backups”](#) on page 60.

To create a copy-only backup

- 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 (legacy SQL Server 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 backups (legacy SQL Server policies)

Note: This feature is only available with legacy SQL Server backup policies.

The SQL Server agent provides a method in which multiple databases can be quiesced together and split-off to form a single snapshot. This method minimizes the usage of system resources if the databases exist on a single volume. This happens because the aggregation of constituent files uses one snapshot volume

instead of one per database. The method for aggregating database Snapshot Client backups is called backup "grouping".

When databases are backed up in a group, all of the databases are quiesced simultaneously. 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 backup

Certain requirements must be met for a grouped backup 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 are not supported.
- The master database cannot be included in a grouped backup.
- The same policy must be specified for each backup operation in the group.
- The same NetBackup server must be specified for each backup operation in the group.


Viewing the progress of a grouped backup

You can determine that a grouped backup is underway from the progress report.

See [Figure 8-2](#).

The keyword GROUPSIZE appears at the beginning of the batch file. This keyword indicates that NetBackup uses grouping to back up the selected SQL Server databases. If the appropriate conditions apply all operations are full database backups. Then all of the databases are snapped and backed up as a group. When this action happens, the progress log displays the backup image name as well as the storage image for each database in the group.

Figure 8-2 Progress report for a grouped backup operation



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 [“Restoring a SQL Server database backup”](#) on page 73.

When you launch the restore operation, note that the batch file specifies the storage image name and the backup image name.

See [Figure 8-3](#) on page 112.

Figure 8-3 Batch file shown in the progress report for the restore operation



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 legacy policies](#)
- [Protect a SQL Server availability group that crosses NetBackup domains](#)
- [Browsing for SQL Server availability group backup images](#)
- [Restoring a SQL Server availability database to a secondary replica](#)
- [Restoring a SQL Server availability database to the primary and the secondary replicas](#)
- [Restoring an availability database when an availability group crosses NetBackup domains](#)

About protecting SQL Server availability groups

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

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

- With a 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”](#) on page 129.

Note the following before you configure the policy:

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

Limitations

Note the following limitations for backups of availability groups:

- NetBackup does not support the following types of backups for availability databases:
 - Snapshot backups of filegroups or files
 - Instant Recovery backups
 - VMware backups
 - (Legacy policies) A grouped snapshot backup that includes databases in more than one availability group or in both availability databases and standard databases
 - (Intelligent Policies) A grouped snapshot backup
 - Backups of non-readable secondary replicas
NetBackup can only back up databases in a replica when you allow user connections for the replica.
If a secondary replica is the preferred replica and it is non-readable, the backup fails. If a secondary replica is not the preferred replica, NetBackup skips the backup of that replica.

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

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

Backups of this type result in a failed backup.

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”](#) on page 129.
- NetBackup supports backups of availability groups in multi-NIC environments. For more information, see the following topic:
See [“About configuration of SQL Server backups with multiple NICs”](#) on page 176.

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 9-1](#)

Table 9-1 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 “Registering a SQL Server instance or availability replica” on page 38.
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 Host Management properties on the primary server. See “Reviewing the auto-discovered mappings in Host Management” on page 22.

Table 9-1 Prerequisites for protecting the preferred or the primary replica in an availability group (*continued*)

Step	Action	Description
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 27.

Configuring 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 Open the NetBackup Administration Console.
- 2 Create a new policy.
- 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 49.
- 4 Click on the **Instances and Databases** tab.
- 5 Select **Protect availability groups**.

See [“About tuning parameters for SQL Server backups”](#) on page 60.
- 6 Click **New**.
- 7 Select the availability groups or availability databases that you want to protect.

See [“Adding an availability group to a policy”](#) on page 117.

See [“Adding availability databases to a policy”](#) on page 118.
- 8 Add schedules.

See [“About schedule properties”](#) on page 50.

- 9** Click the **Microsoft SQL Server** tab.
- 10** From the **Availability Database Backup Preference** list, choose one of the following:
 - **Protect primary replica**
 - **Protect preferred replica**See [“About tuning parameters for SQL Server backups”](#) on page 60.
- 11** (Optional) Make any other changes to the tuning parameters.
See [“About tuning parameters for SQL Server backups”](#) on page 60.
- 12** Click **OK** to save the policy.


Adding 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 **New**.
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 **OK**.

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.



Adding 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, click **Protect availability groups**.
- 2 Click **New**.
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, check 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 Click **OK**.

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.




Adding 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 Open the policy you want to edit or create a new policy.
- 2 On the **Backup Selections** tab, select **Filegroups** or **Files**.
- 3 In the left pane, expand the availability group and select the replica.

- 4 In the right pane select the filegroups or files.



- 5 Click **OK** 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.



Protecting SQL Server availability groups with legacy policies

You can use legacy 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”](#) on page 129.
- NetBackup supports backups of availability groups in multi-NIC environments.
For more information, see the following topic:
See [“About configuration of SQL Server backups with multiple NICs”](#) on page 176.

About protecting the preferred replica in a SQL Server availability group (legacy backup 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 113.
- Review the prerequisites for protecting the availability group. See [“Prerequisites for protecting SQL Server availability groups”](#) on page 115.

Prerequisites for protecting SQL Server availability groups

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

See [Table 9-2](#)

Table 9-2 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 “Configuring the NetBackup services for SQL Server backups and restores (legacy SQL Server policies)” on page 141.

Table 9-2 Prerequisites for protecting the preferred replica in an availability group *(continued)*

Step	Action	Description
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 27.
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 Host Management properties on the primary server. See “Reviewing the auto-discovered mappings in Host Management” on page 22.

Configuring 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 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 Log on to the primary server as administrator (Windows) or root (UNIX).
- 2 Open the NetBackup Administration Console.
- 3 Select **Actions > New > Policy**.
- 4 In the **Policy name** box, type a unique name for the new policy and click **OK**.
- 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 49.

- 6 On the **Instances and Databases** tab, select **Clients for use with batch files**.
The tab name changes to **Clients** and 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 the [NetBackup Administrator's Guide](#) for more information.

See [“About schedule properties ”](#) on page 154.

- 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 virtual cluster instance name.

- 9 Repeat step 3 to [Configuring an automatic backup policy for the preferred or the primary replica of a SQL Server availability group](#) 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 [“Creating batch files for the policy that protects the preferred or the primary replica”](#) on page 124.

Creating 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` `PREFERRED` 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 [“Configuring an automatic backup policy for the preferred or the primary replica of a SQL Server availability group”](#) on page 122.
- 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 Click **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 [“About using batch files with NetBackup for SQL Server”](#) on page 143.
- 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 [“Adding the batch files to the policy that protects the preferred or the primary replica”](#) on page 125.

Adding 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 [“Configuring an automatic backup policy for the preferred or the primary replica of a SQL Server availability group”](#) on page 122.

See [“Adding the batch files to the policy that protects the preferred or the primary replica”](#) on page 125.
- 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 Click **OK** to save the policy.
- 5 Repeat the steps in this procedure for each policy that you created.

About protecting a specific node in a SQL Server availability group (legacy backup policies)

This topic describes how to protect a specific node in a SQL Server availability group using a legacy SQL Server 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 113.

Configuring 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 Open the NetBackup Administration Console.
- 2 Select **Actions > New > Policy**.
- 3 In the **Policy name** box, type a unique name for the new policy and click **OK**.
- 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 49.

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

The tab name changes to **Clients** and 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](#) for more information.

See [“About schedule properties”](#) on page 154.

- 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 virtual cluster instance name.

- 8 Click **OK** to save the policy.
- 9 Repeat the step 2 through step 8 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 [“Creating a batch file for the policy that protects a specific availability replica in an availability group”](#) on page 127.

Creating 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 [“Configuring an automatic backup policy for a specific replica of a SQL Server availability group”](#) on page 126.
- 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** Click **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 [“About using batch files with NetBackup for SQL Server”](#) on page 143.
- 13** When you have created all the batch files, add these files to the policies that you created previously.

See [“Adding the batch files to the policy that protects a specific replica in the availability group”](#) on page 129.

Adding 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 [“Configuring an automatic backup policy for a specific replica of a SQL Server availability group”](#) on page 126.

See [“Creating a batch file for the policy that protects a specific availability replica in an availability group”](#) on page 127.
- 2 Open the policy that you created.
- 3 On the **Backup Selections** tab, add the batch file(s) that you created.

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

Protect a SQL Server availability group that crosses NetBackup domains

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

- Configure the storage in the NetBackup source and target domains:
 - For OpenStorage, a disk appliance of the same type in each domain. The disk appliance type must support NetBackup Auto Image Replication (A.I.R.).
 - For NetBackup deduplication, the storage that NetBackup can use for a Media Server Deduplication Pool in each domain.
- Configure the domain where the backups occur as the source domain. Then configure the domain where you want to restore the backups as the target domain.

Additional resources

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

[NetBackup Deduplication Guide](#)

NetBackup OpenStorage Solutions Guide

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

Browsing for SQL Server availability group backup images

This procedure describes how to browse for backup images of availability groups. When you have displayed the backup images you want, then follow the instructions for restoring a specific SQL Server object.

To browse for availability group backup images

- 1 Select **File > Restore SQL Server objects**.
- 2 Select the **SQL Host** whose backup images you want to browse, or type its name.
- 3 Select or type the full qualified domain name for the **Source Client**.
 - For an advanced or a basic availability group, provide the name of the Windows Server Failover Clustering (WSFC) cluster.
You can find the cluster name in Failover Cluster Manager or the job details for the backup.
 - For a read-scale availability group, provide the host name of the replica.
- 4 (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.
- 5 Select the date range to search.
- 6 Click **OK**.
- 7 Continue with the applicable instructions for how to restore the object(s).

See “[Restoring a SQL Server availability database to the primary and the secondary replicas](#)” on page 132.

See “[Restoring a SQL Server availability database to a secondary replica](#)” on page 131.

Restoring 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 any system databases or user databases in the backup, perform a separate browse and restore operation using the replica name.

To restore a SQL Server availability database to a secondary replica

- 1 Log on to the node that hosts the secondary replica.
- 2 Close any connections to the database on the secondary replica.
- 3 Remove the secondary database from the availability group.
- 4 In the NetBackup MS SQL Client, select **File > Set SQL Server connection properties**.
- 5 From the **Instance** list, select the instance that hosts the availability group.
- 6 Browse for the backup images you want to restore. Select the latest full backup image and transaction log backups.

See [“Browsing for SQL Server availability group backup images”](#) on page 130.

- 7 Select the following settings:
 - From the **Recovery** list, select **Not recovered**.
 - Select **Use replace option**.
- 8 If the replicas in the availability group use different paths for the database file, you need to create a move template to restore to a secondary replica. From the **Scripting** list, choose **Create a move template**.

See [“Performing a SQL Server database move”](#) on page 77.

- 9 Click **Restore**.
- 10 When the restore completes, join the database to the availability group.

Restoring 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 any system databases or user databases in the backup, perform a separate browse and restore operation using the replica name.

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

- 1 Log on to the host of the primary replica.
- 2 Open SQL Server Management Studio and perform the following tasks:
 - Suspend data movement on the database.
 - Remove the database from the availability group.
- 3 Close any connections to the database.
- 4 Remove the primary database from SQL Server.
- 5 In the NetBackup MS SQL Client, select **File > Set SQL Server connection properties**.
- 6 From the **Instance** list, select the instance that hosts the availability group.
- 7 Browse for the backup images you want to restore. Select the latest full backup image and transaction log backups.

See [“Browsing for SQL Server availability group backup images”](#) on page 130.

- 8 Select the following settings:
 - Select **Use replace option**.
 - From the **Recovery** list, select **Recovered**.
- 9 Click **Restore**.

- 10 When the restore completes, add the database to the availability group using the **Skip initial data synchronization** option.
- 11 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.
- 12 In the NetBackup MS SQL Client, select **File > Set SQL Server connection properties**.
- 13 From the **Instance** list, select the instance that hosts the availability group.
- 14 Browse for the backup images you want to restore. Select the same set of images that you restored to the primary replica.
See [“Browsing for SQL Server availability group backup images”](#) on page 130.
- 15 Select the following settings:
 - From the **Recovery** list, select **Not recovered**.
 - Select **Use replace option**.
- 16 If the replicas in the availability group use different paths for the database file, you need to create a move template to restore to a secondary replica. From the **Scripting** list, choose **Create a move template**.
See [“Performing a SQL Server database move”](#) on page 77.
- 17 Click **Restore**.
- 18 When the restore completes, join the database to the availability group.
- 19 Repeat step 11 through step 18 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”](#) on page 129.

See [“Restoring a SQL Server availability database to a secondary replica”](#) on page 131.

See [“Restoring a SQL Server availability database to the primary and the secondary replicas”](#) on page 132.

Protecting SQL Server in a cluster environment

This chapter includes the following topics:

- [Configuring backups of clustered SQL Server instances \(SQL Server Intelligent Policy\)](#)
- [Configuring backups of clustered SQL Server instances \(legacy SQL Server policies\)](#)
- [Performing a restore of a virtual SQL Server instance](#)

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

This procedure describes how to protect SQL Server clustered instances 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 [“Configuring backups of a SQL Server cluster when you have multiple NICs \(SQL Server Intelligent Policies\)”](#) on page 182.

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

- 1 Open the NetBackup Administration Console.
- 2 Create a policy (for example, VIRTSQLPOLICY).
- 3 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.
- 4 On the **Instances and Databases** tab, select **Protect instances**.
See [“Adding instances to a policy”](#) on page 53.
- 5 Add the instances or databases that you want to protect.
See [“Adding instances to a policy”](#) on page 53.
See [“Adding databases to a policy”](#) on page 55.
See [“Adding instance groups to a backup policy”](#) on page 59.
- For a clustered instance, the host name is the virtual name of the SQL Server cluster.



- 6 Add other policy information as follows:
- Add schedules.
See [“About schedule properties”](#) on page 50.
 - (Optional) Select the specific filegroups or files that you want to back up.
By default, NetBackup backs up an entire database.
See [“Adding filegroups or files to the backup selections list”](#) on page 57.
 - (Optional) Make changes to any tuning parameters.
See [“About tuning parameters for SQL Server backups”](#) on page 60.
- 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 27.
- 8 Configure the **Mapped Host Names** for the SQL Server hosts in your environment.
Configure this property in Host Management on the primary server.
See [“Reviewing the auto-discovered mappings in Host Management”](#) on page 22.

Configuring backups of clustered SQL Server instances (legacy SQL Server policies)

This procedure describes how to protect SQL Server clustered instances with a legacy policy that uses batch files and clients. 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 [“Configuring backups of a SQL Server cluster when you have multiple NICs \(legacy SQL Server policies\)”](#) on page 183.

To configure backups of clustered SQL Server instances

- 1 Open the NetBackup Administration Console.
- 2 Create a policy (for example, VIRTSQLPOLICY).
- 3 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.
- 4 On the **Instances and Databases** tab, select **Clients for use with batch files**.
- 5 On the **Schedules** tab, add an automatic backup schedule.
- 6 On the **Clients** tab, add the virtual SQL Server name (VIRTUALSERVER).
- 7 On the **Backup Selections** tab, add one or more script names (batch files).
- 8 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 27.

- 9 Configure the **Mapped Host Names** for the SQL Server hosts in your environment.

Configure this property in Host Management on the primary server.

See [“Reviewing the auto-discovered mappings in Host Management”](#) on page 22.

Performing a restore of a virtual SQL Server instance

This procedure describes how to perform a restore of a virtual SQL Server instance.

To perform a restore on a virtual SQL Server instance

- 1 Open the NetBackup MS SQL Client on the active node.
- 2 Select **File > Restore SQL Server objects**.
- 3 In the **SQL Host** list, select the virtual server name (VIRTUALSERVER) of the SQL Server.
- 4 Click **OK**.
- 5 Select a backup image or staged image list.
- 6 Click **OK**.

Configuring backups with legacy SQL Server policies using clients and batch files

This chapter includes the following topics:

- [About legacy SQL Server policies](#)
- [About configuring backups with legacy SQL Server policies](#)
- [Configuring the NetBackup services for SQL Server backups and restores \(legacy SQL Server policies\)](#)
- [About SQL Server security with NetBackup legacy backup policies](#)
- [About using batch files with NetBackup for SQL Server](#)
- [Adding a new SQL Server legacy policy](#)
- [About schedule properties](#)
- [Adding clients to a policy](#)
- [Adding batch files to the backup selections list](#)
- [Selecting the SQL Server host and instance](#)
- [Options for SQL Server backup operations](#)
- [About viewing the properties of the objects selected for backup](#)

- [Performing user-directed backups of SQL Server databases](#)
- [Performing a backup of a remote SQL Server installation](#)
- [About file checkpointing with NetBackup for SQL Server](#)
- [About automatic retry of unsuccessful SQL Server backups](#)

About legacy SQL Server policies

A legacy NetBackup for SQL 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 legacy SQL Server 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

About configuring backups with legacy SQL Server policies

Table 11-1 Steps to configure SQL Server backups that use legacy SQL Server 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 “Configuring the NetBackup services for SQL Server backups and restores” on page 19.
Step 2	Configure the batch files for the policy.	See “About using batch files with NetBackup for SQL Server” on page 143.

Table 11-1 Steps to configure SQL Server backups that use legacy SQL Server policies (*continued*)

Step	Action	Description
Step 3	Configure a legacy SQL Server policy.	See “Adding a new SQL Server legacy policy” on page 153.
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 27.
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 in Host Management” on page 22.

Configuring the NetBackup services for SQL Server backups and restores (legacy SQL Server 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 [“Configuring the NetBackup services for a VMware backup that protects SQL Server”](#) on page 93.

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 NetBackup legacy backup 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\
```

About using 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 executed from the `install_path\DbExt\MsSql\` directory.

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

- NetBackup MS SQL Client
- `dbbackex` command line
- Automatically scheduled backups 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 221.
- Batch files are in Unicode text.
- A batch file consists of a series of operations that run in sequence. For legacy SQL Server backup 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 executed. 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 [“Creating a batch file”](#) on page 152.

See [“About using batch files with NetBackup for SQL Server”](#) on page 143.

[Table 11-2](#) describes the keywords and values that can be used in batch files.

Table 11-2 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_</code> <code>SNAPSHOT</code> Default: <code>BACKUPMODEL_</code> <code>CONVENTIONAL</code> Required: No

Table 11-2 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 156.	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 108.	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 11-2 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 2-31. (Legacy policies) For availability group backups, all databases in the grouped backup must be part of the availability group. NetBackup does not support any grouped snapshot backups that include both standard databases and availability databases. (Intelligent Policies) NetBackup does not support grouped snapshot backups.	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 - (NetBackup 9.1 and later clients) 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

Table 11-2 Keywords and values used in batch files (*continued*)

Keyword and description	Type and values
MAXRESTARTSETS - Use MAXRESTARTSETS to enable file checkpointing. This parameter specifies the number of separate streams into which the backup request is sub-divided. Range is 2–32.	Integer Default: None Required: No
MAXTRANSFERSIZE - Maximum transfer size is calculated as 64 KB * 2^MAXTRANSFERSIZE. Range is 0–6.	Integer Default: 0 Required: No
MOVE - Specifies a filegroup name. Used for the MOVE restore type. For any backups that were made with a SQL Server legacy 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*

Table 11-2 Keywords and values used in batch files (*continued*)

Keyword and description	Type and values
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
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
PREFERREDREPLICA - 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. (NetBackup 8.2 and later clients) 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

Table 11-2 Keywords and values used in batch files (*continued*)

Keyword and description	Type and values
RESTARTTYPE Available only for backups. Use <code>AUTO</code> to automatically retry backup of failed objects. Use <code>MANUAL</code> 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.	<code>AUTO</code>, <code>MANUAL</code> Default: None Required: No
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 (<code>STOPAFTER</code>).	String Default: None Required: No
RESTORECOPYNUM - (NetBackup 9.1 and later clients) 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.	<code>REPLACE</code> 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.	<code>TRUE</code>, <code>FALSE</code> Default: <code>FALSE</code> Required: No
RESTORETOMARK - Recovers the transaction log to a transaction log mark.	String Default: None Required: No

Table 11-2 Keywords and values used in batch files (*continued*)

Keyword and description	Type and values
RESTORETOMARK AFTERTIME - Recovers the transaction log to a transaction log mark, but after a point in time (STOPAFTER).	String Default: None Required: No
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 path> sequences. (SQL Server legacy 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

Table 11-2 Keywords and values used in batch files (*continued*)

Keyword and description	Type and values
STOPAFTER - Specifies datetime for RESTORETOMARK options. The datetime string is formatted as YYYY/MMDDHH:MM:SS .	Datetime string Default: None Required: No
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

Table 11-2 Keywords and values used in batch files (*continued*)

Keyword and description	Type and values
VERIFYONLY - Tells SQL Server to verify a backup image but not to restore it.	TRUE, FALSE Default: FALSE Required: No
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

Creating 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 `dbbackex` 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 Select **File > Backup SQL Server objects** or **File > Restore SQL Server objects**.
- 2 Select the object you want to back up or restore.
- 3 Select the backup or restore options.
See [“Options for SQL Server backup operations”](#) on page 162.
See [“Options for NetBackup for SQL Server restores”](#) on page 71.
- 4 In the **Backup script** or **Restore script** group, click **Save**.
- 5 Click **Backup** or **Restore**.
- 6 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 221.

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

- 8 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.

Running 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 Log on to the host and instance you want to access.
See [“Selecting the SQL Server host and instance”](#) on page 161.
- 2 Select **File > Manage script files**.
- 3 Double-click the batch file.
- 4 Click **Start**.
- 5 To monitor the operation, select **File > View status**.

Adding a new SQL Server legacy policy

This topic describes how to create a SQL Server legacy 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 [“Configuring multistriped backups of SQL Server”](#) on page 64.

See [“About configuration of SQL Server backups with multiple NICs”](#) on page 176.

To add a new SQL Server legacy policy

- 1 Log on to the primary server as administrator (Windows) or root (UNIX).
- 2 Open the NetBackup Administration Console.
- 3 If your site has more than one primary server, choose the one on which you want to add the policy.
- 4 In the left pane, expand **NetBackup Management** and select **Policies**.
- 5 Select **Actions > New > Policy**.

- 6 In the **Add a New Policy** dialog box, in the **Policy name** box, type a unique name for the new policy.
- 7 Click **OK**.
- 8 In the **Add New Policy** dialog box, in the **Policy type** list, select **MS-SQL-Server**.

The database agent policy type does not appear in the drop-down list unless your primary server has a license for the database agent.
- 9 Complete the entries on the **Attributes** tab.

See [“About policy attributes”](#) on page 49.
- 10 On the **Instances and Databases** tab, select **Clients for use with batch files**.

The tab name changes to **Clients** and the **Backup Selections** tab now lets you specify and browse for scripts.
- 11 Add other policy information as follows:
 - Add schedules.
 See [“About schedule properties”](#) on page 154.
 - Add clients.
 See [“Adding clients to a policy”](#) on page 159.
 - Add batch files to the backup selections list.
 See [“Adding batch files to the backup selections list”](#) on page 160.
- 12 When you have added all the schedules, clients, and backup selections you need, click **OK**.

About schedule properties

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 11-3 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 “Legacy policy backup types” on page 155.

Table 11-3 Description of schedule properties (*continued*)

Property	Description
Schedule type	You can schedule an automatic 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. ■ 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.
Multiple copies	If you want to specify multiple copies of a backup for the policy, configure Multiple copies on the application backup schedule. If using Snapshot Client, also specify Multiple copies on the automatic schedule.

Legacy policy backup types

[Table 11-4](#) shows that the backup types you can specify for a NetBackup for SQL Server legacy policy that uses clients and batch files. Intelligent Policies have a different set of backup types.

Table 11-4 Legacy policy backup types

Backup type	Description
Application Backup	The application backup schedule enables user-controlled NetBackup operations from the client. These operations include those initiated from the client and those initiated 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.
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 144. See “Converting differential backups to full backups” on page 156.

Converting differential backups to full backups

If a differential backup runs and a full backup does 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. Enable this behavior with the keyword `CONVERTBACKUP`.

See [“Keywords and values used in batch files”](#) on page 144.

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 MS SQL Client. Or, if NetBackup expired the backup, you can import the full backups into the NetBackup catalog. Then you can restore both the full and the differential backups with the NetBackup MS SQL Client.

The agent checks to determine if a full backup was ever performed for each database. If no previous full backup exists, the backup is converted to a full as follows:

- If you select a database for backup, the backup is converted to a full database backup.
If you select **Read-write filegroups** for the **Type of Backup**, the backup is converted to a full read/write filegroup backup.
- If you select a filegroup for 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.
- If you perform a partial differential backup, NetBackup does the following:
 - If no previous full backup exists for the default filegroup, NetBackup adds the filegroup to the backup and converts the operation to a full partial backup.
 - If a previous full backup exists for the default filegroup but a secondary filegroup in the files list does not have a full backup, NetBackup converts the operation to a full partial backup.

- The `CONVERTBACKUP` 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.

Configuring 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. The NetBackup for SQL Server agent 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 In the **Policy** dialog box, click the **Schedules** tab.
To access the **Policy** dialog box, double-click the policy name in the **Policies** list in the NetBackup Administration Console.
- 2 Double-click the schedule that is named **Default-Application-Backup**.
- 3 Specify the other properties for the schedule.
See [“About schedule properties”](#) on page 154.

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 11-5 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

Configuring 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 On the **Policy** dialog box, click the **Schedules** tab.
- 2 Click **New**.
- 3 Specify a unique name for the schedule.
- 4 Select the **Full Backup** schedule.
See “[Legacy policy backup types](#)” on page 155.
- 5 Specify the other properties for the schedule.
See “[About schedule properties](#) ” on page 154.
- 6 Click **OK**.

Example automatic backup schedule

[Table 11-6](#) shows example settings for an automatic backup schedule.

Table 11-6 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

Adding 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.

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

- SQL Server
- NetBackup client or server
- The backup or restore batch files

Note: Each batch file must be present on each client.

To add clients to a policy

- 1 Open the policy you want to edit or create a new policy.
To access the **Policy** dialog box, double-click the policy name in the **Policies** list in the NetBackup Administration Console.
- 2 Before you can add clients, you must select **Clients for use with batch files** on the **Instances and Databases** tab.
- 3 Click the **Clients** tab and click **New**.

- 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 in Host Management”](#) on page 22.

See [“Configuring mappings for restores of a distributed applications, clusters, or virtual machines ”](#) on page 27.

- 5 Choose one of the following:
 - To add another client, click **Add**.
 - If this client is the last client you want to add, click **OK**.
- 6 In the **Policy** dialog box, click **OK**.

Adding 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. Create the schedules on the **Schedules** tab. NetBackup runs the batch files in the order that the batch files appear in the backup selections list.

Note: Specify the correct batch file names in the backup selections list to prevent an error or possibly a wrong operation.

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 221.
- 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 Click the **Backup Selections** tab.
- 5 Click **New**.
- 6 In the **Add Backup Selection** dialog box, specify the names of the batch files that you want to use. Specify the file name in one of the following ways:
 - Click **Browse**. Navigate to and select the batch file, then click **OK**.
 - 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 Click **OK** to add the batch files to the backup selections list.
- 9 Click **OK**.

Selecting the SQL Server host and instance

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 In the **SQL Server connection properties** dialog box, from the **Host** drop-down list, select the SQL Server host.
- You can type a host name if it does not appear in the drop-down list. If you select a remote host and click **Apply**, the **Host type** is shown as "remote".
- 4 From the **Instance** drop-down list, select the SQL Server instance.
- You can type an instance name if it does not appear in the drop-down list. You can designate the default instance either by setting the Instance box to <default> or to empty (no spaces).
- 5 Click **Apply** to save your changes.
- 6 Click **Close**.

Options for SQL Server backup operations

Table 11-7 describes the options that are available when you perform backups. 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 11-7 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.

Table 11-7 Options for SQL Server backup operations (*continued*)

Option	Description
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.
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.

Table 11-7 Options for SQL Server backup operations (*continued*)

Option	Description
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 “Configuring multistriped backups of SQL Server” on page 64.
Resume options for this selection	■ Do not resume unsuccessful backups ■ Retry from the beginning Restart failed backups after waiting 60 seconds. ■ Save work and restart at point of failure Divide the backup into multiple streams and back up separately. Any streams that fail are restarted after 60 seconds. This option is available when the following conditions are met: ■ Exactly one object has been selected, ■ The object that is selected for backup is a database or filegroup and the backup type is full, ■ The SQL Server object uses the “full” or “bulk-logged” recovery method.
NetBackup policy	If this host is the NetBackup master 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.

Table 11-7 Options for SQL Server backup operations (continued)

Option	Description
Backup	Start a database backup or generate a database backup script. This option is enabled only when you select an object to back up.

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 11-8](#) 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 11-8 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.

Table 11-8 Properties of the objects that are selected for backup (*continued*)

Property	Description
Path	(Database files only) The absolute path of the database file.

Performing user-directed backups of SQL Server databases

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 [“Selecting the SQL Server host and instance”](#) on page 161.
- 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 162.
- 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**.

Performing user-directed backups of SQL Server transaction logs

This procedure describes how to perform a transaction log backup.

Caution: Ensure that the entire sequence of transaction logs 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 [“Selecting the SQL Server host and instance”](#) on page 161.
- 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 drop-down 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 disk.

- 9 Select the backup options.
See [“Options for SQL Server backup operations”](#) on page 162.
- 10 Click **Backup**.
To view the progress of the backup, select File > **View status**.

Performing user-directed backups of SQL Server database filegroups

More information is available on how to use read-write and read-only filegroups in your backup strategy.
See [“Performing user-directed backups of read-write filegroups”](#) on page 169.

See [“Performing user-directed backups of read-only filegroups”](#) on page 168.

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 [“Selecting the SQL Server host and instance”](#) on page 161.
- 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 162.
- 8 Click **Backup**.

To view the progress of the backup, select File > **View status**.

Performing user-directed backups 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 [“Viewing SQL Server read-only backup sets”](#) on page 169.

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 Administration Console, 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 all read-only groups are backed up by viewing the read-only backup set.

See [“Viewing SQL Server read-only backup sets”](#) on page 169.

Viewing SQL Server read-only backup sets

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”](#) on page 69.
- 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**.

Performing user-directed backups 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 [“Performing user-directed backups of read-write filegroups”](#) on page 169.

See [“Performing user-directed backups of read-only filegroups”](#) on page 168.

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 **Backup Microsoft SQL Server Objects** dialog box, 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 162.
- 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 Administration Console.
- 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 user-directed backups of SQL Server database files

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 [“Selecting the SQL Server host and instance”](#) on page 161.
- 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 162.
- 8 Click **Backup**.
To view the progress of the backup, select File > **View status**.

Performing partial database backups

This procedure describes how to create a script for to perform a partial database backup. This type of back is only available for SQL Server legacy backup 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 [“Selecting the SQL Server host and instance”](#) on page 161.
- 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 162.
- 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**.

Performing a backup of a remote SQL Server installation

You can use NetBackup for SQL Server 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 perform a backup of a remote SQL Server installation

- 1 Select the host and instance you want to access.
See [“Selecting the SQL Server host and instance”](#) on page 161.
- 2 Select **File > Backup SQL Server objects**.
- 3 Select the options for the operation.

See [“Options for SQL Server backup operations”](#) on page 162.

Save is enabled in the backup dialog box. **Launch immediately** is disabled because the generated script must be executed 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 file checkpointing with NetBackup for SQL Server

Use file checkpointing if you need to perform a large backup and want to save completed work in case the operation fails before it completes. When file checkpointing is enabled, the database or filegroup is divided into file sets and backed up as separate units. The following batch file command initiates file checkpointing:

`MAXRESTARTSETS integer`

The backup operation is split into the number of operations equal to the *integer* value. If the number of files is less than the *integer* value, then the number of separate operations is equal to the number of files.

File checkpointing is available for databases and filegroups that are backed up as streams or with the snapshot option. However, the following restrictions exist:

- The backup object must contain at least two files.
- The recovery model of the database cannot be “simple”.
- If the snapshot option is used for backup, then the method cannot be Instant Recovery. However, file checkpointing that uses Instant Recovery to a storage unit is supported.
- The batch file that you use for a file checkpoint backup can specify only one database or filegroup. You cannot use the `DATABASE $ALL` option.

When you use file checkpointing for backing up a full database, NetBackup for SQL Server automatically splits the database into fileset components. Recovering the database from components requires a restore of the transaction log. NetBackup for SQL Server automatically includes a backup log directive in the generated batch file when you choose file checkpointing from the backup dialog box.

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 <i>install_path\dbext\mssql\temp</i> 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
```

Retry may also be used with file checkpoints. Any parts of the operation that fail can be written to a new batch file that can be launched either automatically or manually.

See [“About file checkpointing with NetBackup for SQL Server ”](#) on page 173.

You can enable file checkpointing with automatic retry in the backup dialog in the NetBackup for SQL Server Client. Select a single database (or filegroup), then from the **Resume options for this selection** list, select **Save work and restart at point of failure**.

This action creates a batch file that contains the following scripting:

```
MAXRESTARTSETS 32  
RESTARTWAITSECONDS 60  
NUMRESTARTS 1
```

`MAXRESTARTSETS 32` means that up to 32 pieces are backed up independently. The keywords `RESTARTWAITSECONDS` and `NUMRESTARTS` are synonymous with the following:

```
RETRYWAITSECONDS 60  
NUMRETRIES 1
```

These keywords indicates the following things: first, that an automatic retry is launched after 60 seconds for all of the pieces that failed to get backed up on the first time. Second, the restart is attempted only one time. You can manually change either of these parameters.

In addition, you can choose to not have the retry script automatically launched. Replace the `NUMRETRIES` command with `RETRYTYPE MANUAL`. For example, replace the following:

```
NUMRETRIES 1
```

with

```
RETRYTYPE 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.

Using NetBackup for SQL Server with multiple NICs

This chapter includes the following topics:

- [About configuration of SQL Server backups with multiple NICs](#)
- [Configuring the NetBackup client with the private interface name](#)
- [Configuring backups of SQL Server when you have multiple NICs \(SQL Server Intelligent Policies\)](#)
- [Configuring backups for SQL Server when you have multiple NICs \(legacy SQL Server policies\)](#)
- [Performing restores of SQL Server when you have multiple NICs](#)
- [Configuring backups of a SQL Server cluster when you have multiple NICs \(SQL Server Intelligent Policies\)](#)
- [Configuring backups of a SQL Server cluster when you have multiple NICs \(legacy SQL Server policies\)](#)
- [Creating a batch file for backups of a SQL Server cluster when you have multiple NICs \(legacy SQL Server policies\)](#)
- [Performing restores of a SQL Server cluster when you have multiple NICs](#)

About configuration of 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 and 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 27.

See [“Reviewing the auto-discovered mappings in Host Management”](#) on page 22.

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 [“Configuring the NetBackup client with the private interface name”](#) on page 178.
- For intelligent policies, configure a backup policy that includes the private interface name of the host or client.
See [“Configuring backups of SQL Server when you have multiple NICs \(SQL Server Intelligent Policies\)”](#) on page 179.
See [“Configuring backups of a SQL Server cluster when you have multiple NICs \(SQL Server Intelligent Policies\)”](#) on page 182.
- For legacy SQL Server policies, configure a backup policy that includes the private interface name of the host or client.
See [“Configuring backups for SQL Server when you have multiple NICs \(legacy SQL Server policies\)”](#) on page 180.

See [“Configuring backups of a SQL Server cluster when you have multiple NICs \(legacy SQL Server policies\)”](#) on page 183.

Note that if you want to protect a SQL Server cluster with a legacy SQL Server policy, you must edit the backup batch file. The `BROWSECLIENT` parameter must indicate the private name of 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 81.

- For restores in a multi-NIC environment, refer to the following topic:

See [“Performing restores of SQL Server when you have multiple NICs”](#) on page 181.

If you want to perform a restore from a SQL Server cluster, you must edit the restore batch file. In the batch file, you must change the `BROWSECLIENT` parameter to indicate the private name of virtual SQL Server.

See [“Performing restores of a SQL Server cluster when you have multiple NICs”](#) on page 185.

Configuring 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 [“Configuring backups of clustered SQL Server instances \(legacy SQL Server policies\)”](#) on page 137.

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 Click 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`.

Configuring 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 [“Configuring the NetBackup client with the private interface name”](#) on page 178.
- 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 [“Configuring the NetBackup client with the private interface name”](#) on page 178.
- 2 Open the NetBackup Administration Console.
- 3 Expand **NetBackup Management > Applications > Microsoft SQL Server**.
- 4 Click **All Instances**.
- 5 Find and register the instance that has the private interface name of the SQL Server host (sqlhost1-NB).
- 6 Create a new policy or open an existing policy.
- 7 On the **Instances and Databases** tab, select **Protect instances**.
- 8 Click **New**.

- 9 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 [“Adding instances to a policy”](#) on page 53.
 See [“Adding databases to a policy”](#) on page 55.
- 10 Add other policy information as follows:
 - Add schedules.
 See [“About schedule properties”](#) on page 50.
 - Add database objects to the backup selections list.
 See [“Adding filegroups or files to the backup selections list”](#) on page 57.
 - (Optional) Make changes to any tuning parameters.
 See [“About tuning parameters for SQL Server backups”](#) on page 60.

Configuring backups for SQL Server when you have multiple NICs (legacy SQL Server policies)

This topic describes how to configure a legacy backup 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 [“Configuring the NetBackup client with the private interface name”](#) on page 178.
- 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 (legacy backup 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 [“Configuring the NetBackup client with the private interface name”](#) on page 178.
- 2 Open the NetBackup Administration Console.
- 3 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 Add other policy information as follows:

- Add schedules.
See [“About schedule properties”](#) on page 154.
- Create and add batch files to the backup selections list.
See [“About using batch files with NetBackup for SQL Server”](#) on page 143.
See [“Adding batch files to the backup selections list”](#) on page 160.

Performing restores of SQL Server when you have multiple NICs

To perform restores of a SQL Server in a multi-NIC environment, you need to do the following:

- Connect to SQL Server host using the public name of the host.
- To browse for backup images, specify 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 [“Performing restores of a SQL Server cluster when you have multiple NICs”](#) on page 185.

To perform SQL Server restores 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 Click **OK**.
- 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 Click **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 [“Restoring a SQL Server database backup”](#) on page 73.

Configuring 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 Open the NetBackup Administration Console.
- 2 Expand **NetBackup Management > Applications > Microsoft SQL Server**.
- 3 Click **All Instances**.
- 4 Manually add a new instance and register it. For the **Host**, provide the private interface name of the virtual SQL Server (`virtsql-NB`).
- 5 Create a new policy or open an existing policy.
- 6 On the **Instances and Databases** tab, select **Protect instances**.
- 7 Click **New**.
- 8 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`).

See [“Adding instances to a policy”](#) on page 53.

See [“Adding databases to a policy”](#) on page 55.

- 9 Add other policy information as follows:
 - Add schedules.
See [“About schedule properties”](#) on page 50.

- Add database objects to the backup selections list.
 See [“Adding filegroups or files to the backup selections list”](#) on page 57.
- (Optional) Make changes to any tuning parameters.
 See [“About tuning parameters for SQL Server backups”](#) on page 60.

Configuring backups of a SQL Server cluster when you have multiple NICs (legacy SQL Server policies)

This topic describes how to create a legacy SQL Server backup 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 (legacy backup policies)

- 1 Open the NetBackup Administration Console.
- 2 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`.

- 4 Add other policy information as follows:
 - Add schedules.
 See [“About schedule properties”](#) on page 154.
 - 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 [“Creating a batch file for backups of a SQL Server cluster when you have multiple NICs \(legacy SQL Server policies\)”](#) on page 183.
 See [“Adding batch files to the backup selections list”](#) on page 160.

Creating a batch file for backups of a SQL Server cluster when you have multiple NICs (legacy SQL Server policies)

This topic describes how to create a batch file for a legacy backup 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 Click **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 162.

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, click **Save**.
- 9 Click **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
UPDATABSE "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.

Performing restores of a SQL Server cluster when you have multiple NICs

To perform restores of a SQL Server cluster in a multi-NIC environment, 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 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 manually edit it to include the private name of the virtual SQL Server.

If you do not have a cluster environment, you must follow a different procedure:

See [“Performing restores of SQL Server when you have multiple NICs”](#) on page 181.

To perform restores of a cluster when you have multiple NICs

- 1** On a specific node in the 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** Click **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** Click **OK**.
- 8** Select the databases to restore.

See [“Options for NetBackup for SQL Server restores”](#) on page 71.

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** Click **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 click **Start**.

Performance and troubleshooting

This chapter includes the following topics:

- [What are the components of NetBackup for SQL Server?](#)
- [How does NetBackup for SQL Server back up a database?](#)
- [How does NetBackup for SQL Server recover a database?](#)
- [Performing a manual backup](#)
- [About debug logging for SQL Server troubleshooting](#)
- [NetBackup for SQL Server performance factors](#)
- [About monitoring NetBackup for SQL Server operations](#)
- [Setting the maximum trace level for NetBackup for SQL Server](#)
- [Troubleshooting credential validation](#)
- [Reporting of unsuccessful filegroup or file backups](#)
- [About minimizing timeout failures on large SQL Server database restores](#)
- [Troubleshooting VMware backups](#)
- [SQL Server log truncation failure during VMware backups of SQL Server](#)
- [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
- About disaster recovery of SQL Server

What are the components of NetBackup for SQL Server?

Table 13-1 describes the components of NetBackup for SQL Server.


Table 13-1 Components of NetBackup for SQL Server

Component	Filename	Description
graphical user interface (GUI)	dbbackup.exe	You use this interface to: ■ Browse database objects and backup images. ■ Create restore scripts and launch restore operations. ■ (Legacy SQL Server policies) Create backup scripts and launch backup operations.
driver	dbbackex.exe	Launches backup and restore operations.
library	dbbackmain.dll	Facilitates backup and restore activities, access to SQL Server, and other operations that NetBackup for SQL Server performs.

These components also interface with `VxBSA.dll`, which is a common NetBackup client module that connects NetBackup for SQL Server to the NetBackup server.

Figure 13-1 shows the relationships of NetBackup for SQL Server with other software components.

Figure 13-1 NetBackup for SQL Server components



The following interactions occur between NetBackup for SQL Server and other software components:

- Every backup or restore operation is initiated through `dbbackex.exe`, in one of the following ways:
 - Scheduled backups
The NetBackup scheduler calls `bphdb` (1), which calls `dbbackex` (2).
 - GUI-initiated backups
`dbbackup.exe` invokes `dbbackex.exe` (3).
 - Command line
`dbbackex.exe` is invoked directly from a command line or third-party tool.
- `Dbbackex.exe` makes function calls to `dbbackmain.dll` (4) to facilitate a backup or a restore operation. The operation is carried out as `dbbackmain.dll` facilitates one or more data streams between SQL Server and NetBackup server. The data stream (7) is established through VDI (5) and the XBSA interface (6). VDI interacts with SQL Server whereas XBSA interacts with the NetBackup database client.

- (Legacy SQL Server policies) The NetBackup for SQL Server GUI (`dbbackup.exe`) lets you browse for SQL Server objects, normally, databases, filegroups, and database files. `dbbackup.exe` invokes `dbbackmain.dll` (8) for accessing the SQL Server master database. NetBackup for SQL Server accesses information about SQL Server through ODBC.
- The NetBackup for SQL Server GUI (`dbbackup.exe`) also lets you browse for SQL Server backup images. The NetBackup catalog contains the images you can browse. To access the contents of the catalog the GUI invokes `dbbackmain.dll`, which uses VxBSA function calls to access the NetBackup server database manager.

How does NetBackup for SQL Server back up a database?

When a backup is executed, NetBackup for SQL Server does the following: creates a backup script, generates an SQL Server backup statement, logs into SQL Server, and delivers the SQL statement to SQL Server through ODBC. Next, the database agent connects to SQL Server through one or more VDI objects. One virtual device is created per backup stripe. In addition, a VxBSA session is initiated for each stripe. These separate sessions allow NetBackup to start a backup job for each stream that is generated from SQL Server.

When the backup completes, the database agent obtains detailed properties of the object that was backed up, including its relationships to other objects. The agent writes this information to the NetBackup catalog and associates it with the backup image. If there are multiple stripes, then the metadata is associated with the first backup image. The adjunct stripes are associated with one another based upon a common naming convention.

How does NetBackup for SQL Server recover a database?

The NetBackup MS SQL Client displays backup images in a logical hierarchy that mirrors the composition of the database. If you select a transaction log or differential image, then NetBackup examines the metadata that is stored with the images for the selected database. It then determines the most efficient recovery set. Then the agent generates a batch file that includes a sequence of scripted restores. When the scripts are executed, the database is recovered.

The individual restore operations work in a similar manner to backups. A SQL Server restore statement is generated and provided to SQL Server by ODBC. A VDI

connection is made. Then a VxBSA session is initiated that starts the data flow between the media manager and SQL Server. NetBackup determines the number of streams (and the corresponding virtual devices and VxBSA sessions) by the number of stripes that were generated during backup.

After all of the recovery operations have completed, the SQL Server agent takes the final step that sets the database into the recovered state. The database goes back online and becomes available for use.

Performing a manual backup

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

To perform a manual backup

- 1 In the left pane, click **Policies**.
- 2 In the **All Policies** pane, select the policy you want to test.
- 3 Select **Actions > 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 legacy SQL Server policies, select the clients that you want to include for the manual backup.

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 Veritas 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 Veritas VSS provider logs. See “Veritas VSS provider logs” on page 194.

Create all debug logs

To create all debug logs

- ◆ Run the following batch file:

```
install_path\NetBackup\logs\mklogdir.bat
```

Setting the debug level

To control the amount of information that is written to the debug logs, change the Database debug level. Typically, the default value of 0 is sufficient. However, technical support may ask you to set the value higher to analyze a problem.

The debug logs are located in *install_path*\NetBackup\logs.

Information is also available about the **Client Trace Level**. See [“Setting the maximum trace level for NetBackup for SQL Server”](#) on page 200.

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.

Veritas VSS provider logs

The Veritas 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 Veritas VSS provider logging in the registry

Enable the Veritas VSS provider logging on the NetBackup computer where SQL Server is installed.

To enable Veritas 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 Veritas 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 Veritas 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
```

NetBackup for SQL Server performance factors

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

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

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

SQL Server buffer space parameters

The **Maximum transfer size**, **Backup block size**, and **Client buffers per stripe** can increase buffer space in SQL Server. SQL Server must have the available resources to support the increase of these values. Buffer space parameters are applicable for stream-based backups only.

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

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

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

Stripes and parallel backup operations

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

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

See [“Configure the number of jobs allowed for backup operations”](#) on page 29.

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

Shared memory usage

For optimal performance, install NetBackup server on the same host as NetBackup for SQL Server. Also use shared memory for data transfer instead of sockets.

Shared memory is the default unless you create a `install_path\NetBackup\NOSHM` file.

Alternate buffer method

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

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

About alternate buffer method with backup operations

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

- NetBackup shared memory is in use.
- The backup is stream-based.
- The backup is not multiplexed.
- The backup policy does not specify either NetBackup compression or NetBackup encryption.
- The NetBackup buffer size equals the SQL Server block size.

The default NetBackup buffer size is 64 KB, but this value can be overridden in the following settings:

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

`install_path\NetBackup\db\config\SIZE_DATA_BUFFERS_DISK` (for disk backups)

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

The backups that are initiated from an automatic backup policy are started with the NetBackup Client Service so the same account is already in use. However,

you can start a SQL Server backup through NetBackup for SQL Server or through `dbbackupex`. In this case, your logon account must be the same as the NetBackup Client Service account. Then your backups can be candidates for the alternate buffer method.

About alternate buffer method with restore operations

Conditions for backups require that you use the alternate buffer method. Restores also require that backups have been made with the alternate buffer method. You can verify that the alternate buffer method was used. Look for the words `Using alternate buffer method`, which appear in the `dbclient` log and the progress report.

SQL Server checksum

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

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

Instant data file initialization

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

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 in the NetBackup Administration Console 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"
NBSCHEM "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:\NBU\Veritas\NetBackup\dbext\mssql\temp\__01_35_42_508_00.bch>:
<0> operations succeeded. <1> operations failed.

INF - The following object(s) were not backed up successfully.

INF - Archive

```

Setting the maximum trace level for NetBackup for SQL Server

Note: For SQL Server backups, this feature is only available with legacy SQL Server backup 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.

Troubleshooting credential validation

Table 13-2 describes the reasons that validation can fail for an instance, replica, or instance group.

Table 13-2 Reasons for credential validation failure

Status code or error	Description	Explanation
40	Could not validate credentials. Failed to connect to client: <client>.	The host name is invalid.
46	The validation operation timed out waiting for a response from the client	You cannot connect to the host because the host is down.
41	Validation of operating system user/password failed for client: <client>.	■ The host name is correct, but the user name or password is invalid. ■ The credentials use have the setting Use these specific credentials, but the user account does not have the required the local security privileges Impersonate a client after authentication and Replace a process level token. See “Configure local security privileges for SQL Server” on page 21.
1939	The specified user does not have SQL Server System Administrator privileges.	The credentials do not have the “sysadmin” role and the validation fails.
Invalid configuration detected.	Invalid configuration detected. The service user for the NetBackup Client and NetBackup Legacy Network services must be the same user. Change the service users in the Windows Service Manager and try again.	The NetBackup Client Service or the NetBackup Legacy Network Service requires but does not use the same user for the logon account. See “Configuring the NetBackup services for SQL Server backups and restores” on page 19.

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 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. Use the NetBackup Administration Console on the server to change the properties of each client that contains a database 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 195.

Troubleshooting VMware backups

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

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

Table 13-3 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 91. 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.

Table 13-3 Issues with using a VMware policy to protect databases
(continued)

Issue	Explanation
The Application State Capture (ASC) job fails and the databases are not protected.	When the ASC job fails, the VMware snapshot or backup continues. Application-specific data cannot be restored.
	When you query the SQL Server Management Studio (SSMS), it may show that the database was backed up. In this case, though the database was skipped, the snapshot was still successful.
	You disabled the Virtual Machine quiesce option.
	Database objects are on a VHD disk. No objects in the backup are not cataloged, including those that do not exist on the VHD.
	You excluded any data disks from the VMware policy, on the Exclude Disks tab. Be sure that any disks that you exclude do not contain database data.
	The VMware disk layout has changed since the last discovery. In this situation, you must force NetBackup to rediscover virtual machines by lowering the value of the Reuse VM selection query results for option. See the NetBackup for VMware Administrator's Guide .
	You cannot use a VMware incremental policy to protect SQL Server. However, the VMware backup job is successful.
You can recover the entire virtual machine from the backup, but you cannot recover the databases individually.	You did not select Enable SQL Recovery , 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.

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

SQL Server restore fails when you restore a SQL Server compressed backup image as a single stripe or with multiple stripes

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.

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 Veritas Support website:

<http://www.veritas.com/docs/000099850>

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. `Exit status 114` is received in the NetBackup Administration Console when you browse for databases from a SQL Server intelligent policy. In the web UI, a read-scale availability group is not discovered, but no error message is given.

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 208.
- 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 the NetBackup MS SQL Client 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, Veritas 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](#)
- [Configuring 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 85.

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 legacy SQL Server 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 legacy SQL Server 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, Veritas does not recommend multiplexing multiple SQL Server streams from the same backup to a single tape. However, you may want to do this if you vault or export backup images. During the restore of this type of multiplexed backup, NetBackup may time out while trying to synchronize access to data blocks from the backup tape. To prevent this time out, change the stripes parameter in the recovery batch file from `STRIPE N` to `STRIPE 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 [“Creating a batch file for database backups”](#) on page 212.

See [“Creating a batch file for transaction log backups”](#) on page 212.

Creating 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.

Creating 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 198.
- 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 214.

- If the R/3 database disk system is damaged or the transaction log disk system is damaged, follow the instructions for that scenario.
See [“Restoring the R/3 database after a disk crash”](#) on page 214.
- To perform a regular restore of the R/3 database, follow the instructions for that type of restore.
See [“Restoring the database backups and transaction log backups”](#) on page 215.

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.

Restoring 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.

Restoring 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 212.

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 Administration Console 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](#).

Configuring 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. You must disable log truncation for the transaction log backups.

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 that is used to back 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 217.

- 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 81.

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 must *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> CODBAccess::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.>
```

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*.

Configuring 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 [“Performing simultaneous backups for mirrored partners”](#) on page 219.

- 3 On the primary server, configure permissions for a redirected restore for both mirroring partners.

See [“Configuring permissions for redirected restores”](#) on page 81.

- 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.

Performing 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 `SQLHOST` and `SQLINSTANCE` are different.

Restoring 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://www.veritas.com/content/support/en_US/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\Veritas\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 execute on a client. The `none` entry is useful if an administrator wants to completely lock down a server from executing 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.