

Dell EMC NetWorker Module for Microsoft for SQL VDI

Version 18.1

User Guide

302-004-267

REV 01

Copyright © 2007-2018 Dell Inc. or its subsidiaries. All rights reserved.

Published July, 2018

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS-IS.” DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.
Published in the USA.

Dell EMC
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.DellEMC.com

CONTENTS

Figures	7
Tables	9
Preface	11
Chapter 1	Overview 15
	Using NMM with Virtual Device Interface API..... 16
	NMM backup and recovery workflow..... 18
	Traditional backup workflow (stand-alone backup over a storage node)..... 18
	Federated backup workflow (AlwaysOn Availability Group configuration)..... 19
	Recovery workflow..... 20
	SQL Server cluster environments..... 22
	SQL Server backups..... 23
	Types of supported backups..... 23
	Setting backup levels..... 25
	Backup promotion..... 27
	SQL Server recovery..... 29
	Types of supported recovery..... 29
	Recovery modes..... 31
	Recovery time..... 31
	Recovery window restrictions..... 32
	The recovery process..... 32
	SQL Server instance and database names for backup and recovery..... 33
	Named and default instances of SQL Server..... 33
	Supported special characters in database names for NMM backup and recovery 34
Chapter 2	Configuration 37
	Configuring NMM in a SQL VDI environment..... 38
	Supported Windows Server and SQL Server versions..... 38
	Migrating from VSS solution to VDI solution for SQL Server data protection..... 38
	Multi-stream Data Domain Boost..... 38
	Microsoft SQL Server Always On Availability Group feature 38
	Availability group listeners..... 39
	SQL Client Direct to AFTD or DD devices..... 40
	Microsoft SQL Server named log marks..... 40
	Database consistency checks..... 41
	Microsoft hybrid cloud environments..... 42
	Transparent data encryption..... 42
	Setting the MAXTRANSFERSIZE environment variable..... 43
	NMM 18.1 compatibility with NetWorker 8.2.3 or 8.2.4 servers 44

	Configuring permissions to perform NMM backup and recovery of Microsoft SQL Server.....	44
	Access privileges for backup and recovery.....	44
	Assign SQL server roles for backup and recovery operations.....	45
	Assigning Windows user privileges for backup and recovery operations.....	46
Chapter 3	Graphical User Interfaces	49
	User interfaces for backup and recovery.....	50
	NetWorker User for SQL Server GUI.....	51
	Using the nwmssql command.....	51
	Views.....	52
	Display conventions for SQL data types.....	53
	Marking semantics and restrictions.....	54
	Fake objects.....	54
	NMM Microsoft SQL Server Management Studio plugin GUI.....	55
	Views.....	55
Chapter 4	Manual Backups	59
	Manual backup overview.....	60
	Federated backup preferences for Availability Group databases.....	60
	Specifying a retention policy for manual backups.....	62
	Performing manual backups from the NMM SSMS plugin GUI.....	63
	Performing manual backups from the command prompt.....	68
	Command syntax for nsrsqlsv.....	69
	Command options for nsrsqlsv.....	70
	Backup and recovery command syntax for SQL Server data.....	75
	Example backup command syntax.....	79
	Performing manual backups from the NetWorker User for SQL Server GUI... 80	
Chapter 5	Scheduled Backups	87
	Overview of scheduled backup.....	88
	Prerequisites.....	88
	Federated backup preferences for Availability Group databases.....	88
	Excluding incompatible databases in backups.....	90
	Configuring scheduled backups.....	91
	Setting up backup levels.....	92
	Configuring a client resource.....	92
	Setting data protection policies.....	103
	Monitoring scheduled backups	108
Chapter 6	Data Recovery	111
	Overview.....	112
	Prerequisites.....	112
	Performing a recovery from the NMM SSMS plugin GUI.....	113
	Performing a recovery from the command prompt.....	120
	Command syntax for nsrsqlrc.....	121
	Command options for nsrsqlrc.....	122
	Backup and recovery command syntax for SQL Server data.....	131
	Example recovery command syntax	135
	Performing a recovery from the NetWorker User for SQL Server GUI.....	135
	Setting the browse time.....	136

	Setting up the recovery.....	138
	Viewing the required volumes.....	152
	Setting the restore properties.....	152
	Starting the recovery.....	163
	Performing recovery of federated backups with NetWorker User for SQL Server	164
Chapter 7	Granular Level Recovery	171
	Overview.....	172
	Considerations.....	172
	Performing granular-level recovery.....	177
	Dismounting backups after performing granular-level recovery.....	180
	Using the Mount system tray icon.....	181
Chapter 8	Striped Backup and Recovery	183
	Overview of striped backups.....	184
	Performance considerations.....	184
	Performing striped backups.....	185
	Performing striped recovery.....	186
	Optimal striped recovery operations.....	186
	Fail-safe striped recovery operations.....	186
Chapter 9	Disaster Recovery	189
	Overview of disaster recovery.....	190
	Disaster recovery features.....	190
	Performing disaster recovery.....	191
	When not to reinstall the SQL Server.....	191
	Recovering a damaged primary disk.....	192
	Recovering a damaged binary disk.....	193
	Recovering SQL Server and NetWorker server.....	194
	Recovering the SQL Server without reinstalling the SQL Server..	195
	Recovering the SQL Server.....	197
Chapter 10	Bare Metal Recovery	199
	Planning bare-metal recovery.....	200
	Overview.....	200
	System requirements.....	200
	Protecting an environment before a disaster.....	202
	BMR by using NetWorker and NMM.....	203
	SQL Server in a cluster environment.....	204
	Backing up a SQL Server for BMR.....	204
	Performing BMR of a SQL Server cluster.....	206
	SQL Server in a stand-alone environment.....	208
	Backing up a SQL Server for BMR.....	208
	Performing BMR of a standalone SQL Server.....	209
Chapter 11	Troubleshooting	211
	Error logs for backup and recovery.....	212
	SQL savegroup notifications.....	212
	Troubleshooting general issues.....	213
	Troubleshooting GLR.....	213

FIGURES

1	VDI backup process between NMM and SQL Server.....	17
2	VDI recovery process between NMM and SQL Server.....	18
3	Traditional backup workflow.....	19
4	Federated backup workflow.....	20
5	Traditional recovery workflow.....	21
6	Message showing DBCC was successful.....	42
7	Assigning SQL Server privileges.....	46
8	Adding a user to Windows User Groups.....	48
9	Backup window.....	52
10	Restore Operation window.....	52
11	Change Server window.....	53
12	Select SQL Instance window.....	53
13	Marking a fake object.....	55
14	The NetWorker window in the NMM SMSS plugin GUI	56
15	NetWorker Backup General page.....	63
16	NetWorker Backup Options page.....	65
17	NetWorker Backup Monitor page.....	68
18	The Backup window with a hierarchical list of SQL Server data objects.....	81
19	The General tab in the Backup Options dialog box.....	82
20	Properties dialog box.....	84
21	Starting the Client Configuration wizard.....	94
22	Specify Client Information page.....	95
23	Specify Backup Configuration Type page.....	96
24	Specify the Backup Options page.....	98
25	Viewing clients in the Protection tab.....	99
26	Create Policy window.....	104
27	Creating a workflow for the policy.....	105
28	Specifying action information in the Policy Action wizard.....	106
29	Adding a group to a policy from the Protection pane.....	107
30	Specifying workflow and client in the Create Group window.....	108
31	Successful backup messages.....	109
32	Failed backup messages.....	109
33	NetWorker dialog box—Database Restore tab General page.....	114
34	NetWorker dialog box—Database Restore Files/Filegroups page.....	116
35	NetWorker dialog box—Database Restore Options page.....	117
36	NetWorker dialog box—Database Restore Monitor page.....	120
37	Change Browse Time dialog box.....	137
38	Specify Browse Time dialog box.....	137
39	Restore window.....	139
40	Properties dialog box—General tab.....	139
41	Specify the Standby Undo File dialog box.....	140
42	Properties dialog box—Files tab.....	141
43	Specify the File Destination dialog box.....	142
44	Properties dialog box—Restore Time tab.....	143
45	Select the SQL Server dialog box.....	147
46	NetWorker User for SQL Server dialog box.....	148
47	Restore window.....	148
48	Properties dialog box.....	148
49	Restore window.....	149
50	Restore (Verify Only) window.....	150
51	Restore (Piecemeal) window.....	151
52	Restore Options dialog box.....	152
53	Properties dialog box.....	153

54	Properties dialog box.....	155
55	Specify the File Destination dialog box.....	158
56	Read File Configuration dialog box.....	159
57	Point in time restore dialog box.....	162
58	Backed-up availability group databases displayed in the NetWorker User for SQL Server page.....	165
59	Remove Database from Availability Group dialog box.....	167
60	Change Browse Time window.....	167
61	Restore Status window.....	168
62	Add Database to Availability Group window - Select Initial Data Synchronization view	168
63	Add Database to Availability Group window - Validation view.....	169
64	NetWorker dialog box—Table Restore General page.....	178
65	NetWorker dialog box—Table Restore Options page.....	179
66	NetWorker dialog box—Table Restore Monitor page.....	180
67	Mount system tray menu.....	181
68	Mount Details window.....	181
69	SQL Server Properties dialog box.....	207
70	Example of SQL savegroup notification.....	213

TABLES

1	Revision history.....	12
2	Style conventions.....	12
3	Backup levels for SQL Server data.....	25
4	Full backup every 1 to 2 weeks.....	26
5	Logs-only backup after a full backup.....	26
6	Backup level advantages and disadvantages.....	27
7	Creating additional backup levels with data objects.....	27
8	Backup level promotion process	28
9	Types of recovery for SQL Server VDI	29
10	Recovery modes	31
11	Supported special characters in database names.....	34
12	Access privileges required for backup and recovery operations.....	45
13	Where to start backup operations.....	50
14	Where to start recovery operations.....	51
15	Command options for the nwmssql command.....	51
16	SQL Server storage hierarchy display conventions.....	53
17	Command options for nsrsqlsv	70
18	Command syntax for SQL Server data.....	75
19	Command syntax for names containing a period.....	78
20	Command syntax for names containing a backslash.....	78
21	Command syntax for names containing a colon.....	78
22	Tasks for configuring SQL Server VDI backups	91
23	Backup levels for SQL Server data.....	92
24	Save sets.....	100
25	Application Information field values.....	101
26	Advanced recovery options.....	118
27	Command options for nsrsqlrc	122
28	Command syntax for SQL Server data.....	131
29	Command syntax for names containing a period.....	134
30	Command syntax for names containing a backslash.....	134
31	Command syntax for names containing a colon.....	134
32	Modify the destination for the files in list options.....	142
33	Fields on the General tab of the Restore Options dialog box.....	145
34	Advanced recovery options.....	162
35	Guidelines for fail-safe striped recovery.....	186
36	Disaster recovery features.....	190
37	Rebuilding SQL Server system databases.....	192
38	Rebuilding SQL Server system databases.....	193
39	Rebuilding SQL Server system databases.....	194
40	Rebuilding SQL Server system databases.....	196
41	Rebuilding SQL Server system databases.....	207
42	Rebuilding SQL Server system databases.....	210
43	Program and log file names without debug logging enabled.....	212
44	Program and log file names with debug logging enabled.....	212

Preface

As part of an effort to improve product lines, periodic revisions of software and hardware are released. Therefore, all versions of the software or hardware currently in use might not support some functions that are described in this document. The product release notes provide the most up-to-date information on product features.

If a product does not function correctly or does not function as described in this document, contact a technical support professional.

Note

This document was accurate at publication time. To ensure that you are using the latest version of this document, go to the Support website at <https://support.emc.com>.

Purpose

This guide contains information about using the NetWorker Module for Microsoft (NMM) 18.1 software to back up and recover SQL Server using the Virtual Device Interface (VDI) technology.

Note

The *NetWorker Module for Microsoft Administration Guide* supplements the backup and recovery procedures described in this guide and must be referred to when performing application-specific tasks. Ensure to download a copy of the *NetWorker Module for Microsoft Administration Guide* from the Support website at <https://support.emc.com> before using this guide.

Audience

This guide is part of the NetWorker Module for Microsoft documentation set and is intended for use by system administrators during the setup and maintenance of the product. Readers should be familiar with the following technologies used in backup and recovery:

- NetWorker software
- Microsoft Virtual Device Interface (VDI) technology

Revision history

The following table presents the revision history of this document.

Table 1 Revision history

Revision	Date	Description
01	July 07, 2018	First release of this document for NMM 18.1.

Related documentation

The NMM documentation set includes the following publications:

- *NetWorker Module for Microsoft Release Notes*
- *NetWorker Module for Microsoft Administration Guide*
- *NetWorker Module for Microsoft Installation Guide*
- *NetWorker Module for Microsoft for SQL and SharePoint VSS User Guide*
- *NetWorker Module for Microsoft for SQL VDI User Guide*
- *NetWorker Module for Microsoft for Exchange VSS User Guide*
- *NetWorker Module for Microsoft for Hyper-V User Guide*
- *ItemPoint for Microsoft SQL Server User Guide*
- *ItemPoint for Microsoft Exchange Server User Guide*
- *ItemPoint for Microsoft SharePoint Server User Guide*
- NetWorker documentation set

Special notice conventions that are used in this document

The following conventions are used for special notices:

NOTICE

Identifies content that warns of potential business or data loss.

Note

Contains information that is incidental, but not essential, to the topic.

Typographical conventions

The following type style conventions are used in this document:

Table 2 Style conventions

Bold	Used for interface elements that a user specifically selects or clicks, for example, names of buttons, fields, tab names, and menu paths. Also used for the name of a dialog box, page, pane, screen area with title, table label, and window.
<i>Italic</i>	Used for full titles of publications that are referenced in text.
Monospace	Used for: <ul style="list-style-type: none"> • System code • System output, such as an error message or script • Pathnames, file names, file name extensions, prompts, and syntax

Table 2 Style conventions (continued)

	<ul style="list-style-type: none"> • Commands and options
<i>Monospace italic</i>	Used for variables.
Monospace bold	Used for user input.
[]	Square brackets enclose optional values.
	Vertical line indicates alternate selections. The vertical line means or for the alternate selections.
{ }	Braces enclose content that the user must specify, such as x, y, or z.
...	Ellipses indicate non-essential information that is omitted from the example.

You can use the following resources to find more information about this product, obtain support, and provide feedback.


Where to find product documentation

- <https://support.emc.com>
- <https://community.emc.com>

Where to get support

The Support website at <https://support.emc.com> provides access to licensing information, product documentation, advisories, and downloads, as well as how-to and troubleshooting information. This information may enable you to resolve a product issue before you contact Support.

To access a product specific Support page:

1. Go to <https://support.emc.com/products>.
2. In the **Find a Product by Name** box, type a product name, and then select the product from the list that appears.
3. Click .
4. (Optional) To add the product to **My Saved Products**, in the product specific page, click **Add to My Saved Products**.

Knowledgebase

The Knowledgebase contains applicable solutions that you can search for by solution number, for example, 123456, or by keyword.

To search the Knowledgebase:

1. Go to <https://support.emc.com>.
2. Click **Advanced Search**.
The screen refreshes and filter options appear.
3. In the **Search Support or Find Service Request by Number** box, type a solution number or keywords.
4. (Optional) To limit the search to specific products, type a product name in the **Scope by product** box, and then select the product from the list that appears.
5. In the **Scope by resource** list box, select **Knowledgebase**.
The **Knowledgebase Advanced Search** panel appears.
6. (Optional) Specify other filters or advanced options.

7. Click .

Live chat

To participate in a live interactive chat with a support agent:

1. Go to <https://support.emc.com>.
2. Click **Chat with Support**.

Service requests

To obtain in-depth help from Licensing, submit a service request. To submit a service request:

1. Go to <https://support.emc.com>.
2. Click **Create a Service Request**.

Note

To create a service request, you must have a valid support agreement. Contact a sales representative for details about obtaining a valid support agreement or with questions about an account. If you know the service request number, then directly enter the service request number in the `Service Request` field to get the valid details.

To review an open service request:

1. Go to <https://support.emc.com>.
2. Click **Manage service requests**.

Online communities

Go to the Community Network at <https://community.emc.com> for peer contacts, conversations, and content on product support and solutions. Interactively engage online with customers, partners, and certified professionals for all products.

How to provide feedback

Feedback helps to improve the accuracy, organization, and overall quality of publications. You can send feedback to DPAD.Doc.Feedback@emc.com.

CHAPTER 1

Overview

This chapter includes the following sections:

- [Using NMM with Virtual Device Interface API](#) 16
- [NMM backup and recovery workflow](#) 18
- [SQL Server cluster environments](#) 22
- [SQL Server backups](#) 23
- [SQL Server recovery](#) 29
- [SQL Server instance and database names for backup and recovery](#) 33

Using NMM with Virtual Device Interface API

You can use the NetWorker Module for Microsoft (NMM) software to back up and recover Microsoft SQL Server data. NMM uses Virtual Device Interface (VDI), an API that Microsoft SQL Server provides, to integrate with the SQL Server and enable the NetWorker software to back up and recover SQL Server data.

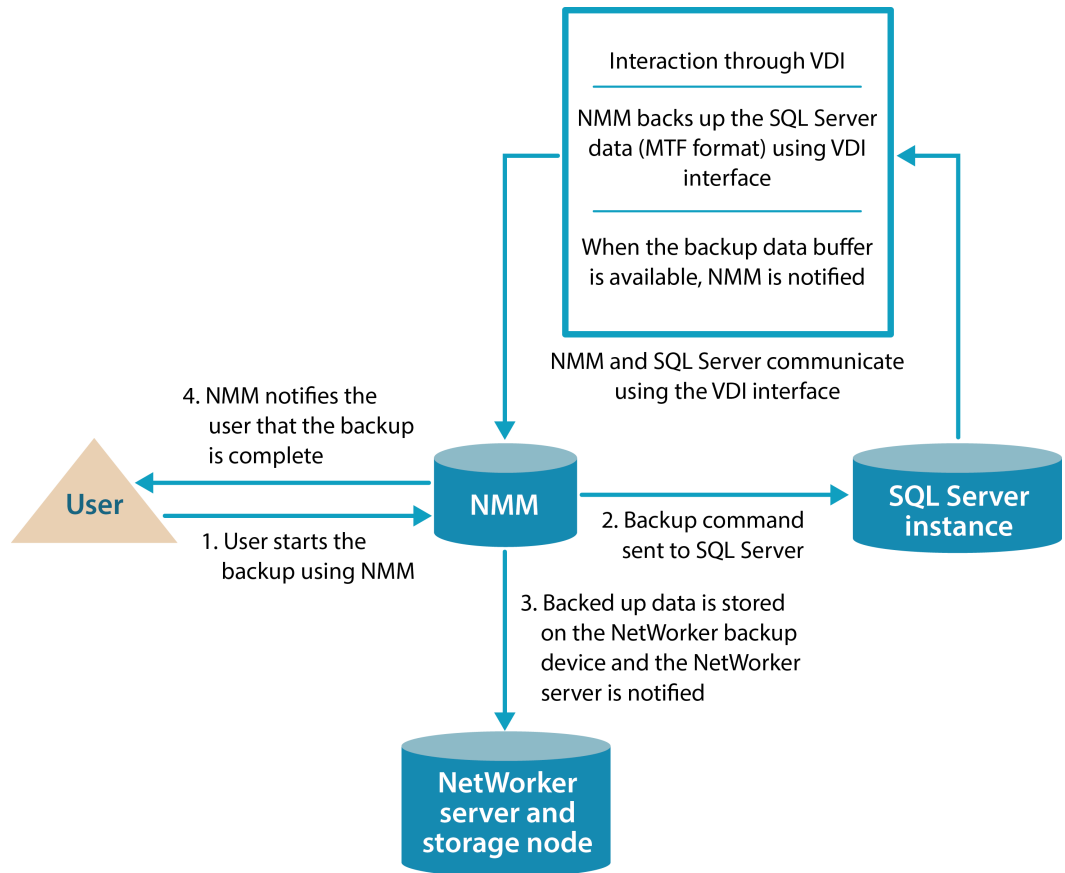
When you install NMM, you can run the System Configuration Checker from the Installation wizard. It is recommended that you run the System Configuration Checker to ensure that the setup is correctly configured for backup and recovery. The *NetWorker Module for Microsoft Installation Guide* provides details.

Note

If you are a NetWorker Module for SQL Server (NMSQL) user and are migrating to NMM VDI, perform a full backup of the SQL Server data after you install NMM VDI. NMM VDI cannot recover SQL snapshot data backed up with NMSQL.

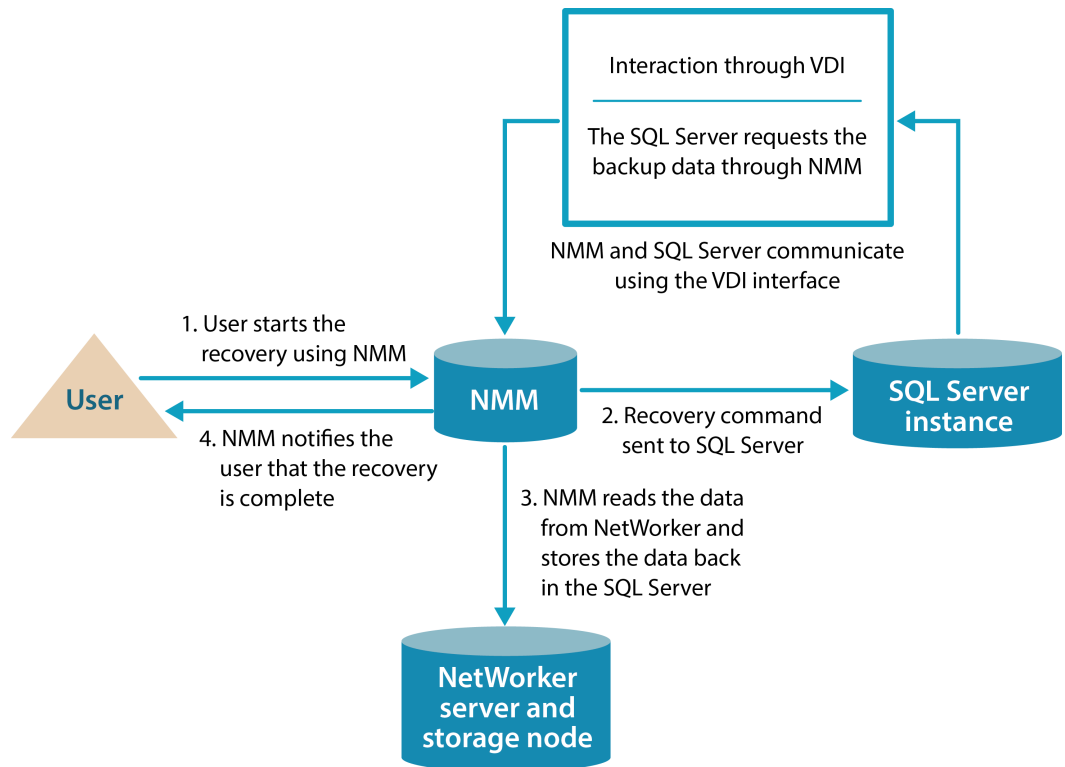
The following figure describes the backup process that takes place between NMM and the SQL Server using VDI:

1. The user starts the backup process with NMM.
2. The backup command is sent to the SQL Server. This interaction is performed through the VDI API.
3. NMM reads the data from SQL Server and stores it on the NetWorker server.
4. NMM sends the backup status to the NetWorker Server and notifies the user when the backup is complete.

Figure 1 VDI backup process between NMM and SQL Server

The following figure describes the recovery process that takes place between NMM and the SQL Server using VDI:

1. The user starts the recovery process with NMM.
2. The restore command is sent to the SQL Server. This interaction is performed through the VDI API.
3. NMM reads the data from the NetWorker server and passes the data to the SQL Server using VDI.
4. NMM notifies the user when the recovery is complete.

Figure 2 VDI recovery process between NMM and SQL Server

NMM backup and recovery workflow

This section provides information about backup and recovery workflows.

Traditional backup workflow (stand-alone backup over a storage node)

During a backup in a traditional environment, processes interact between the NetWorker client (that is, NMM), the NetWorker server, and the SQL Server.

During a traditional backup, the backup starts from the `nsrsqlsv` program, which is started by using one of the following:

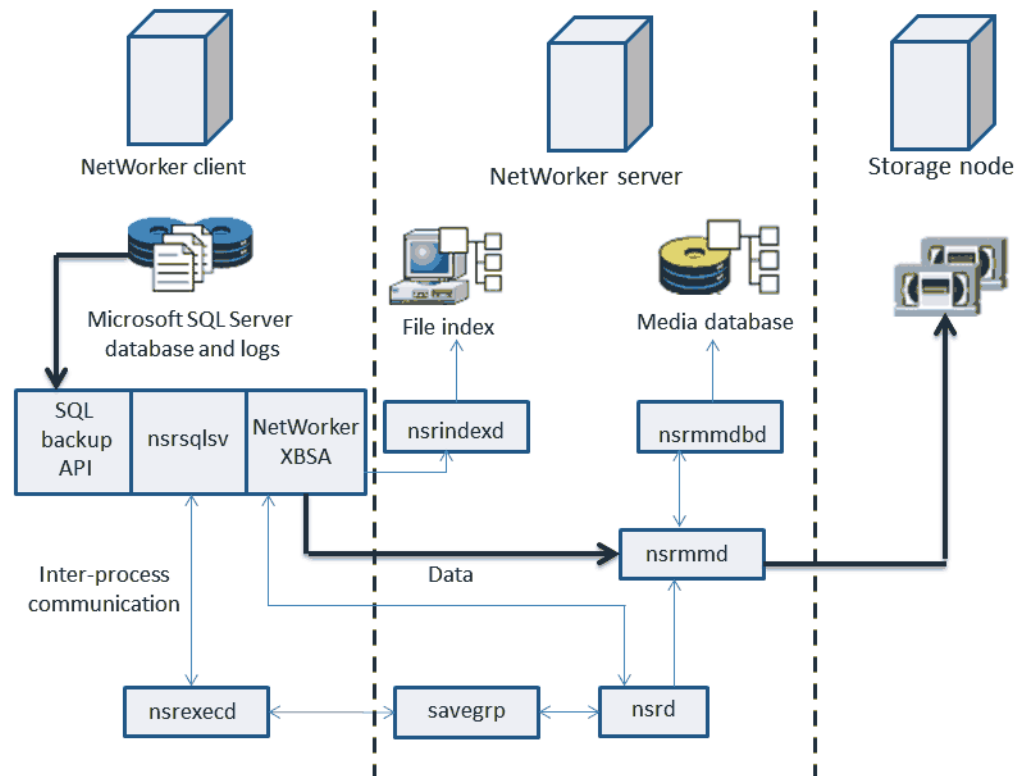
- Command prompt
- Microsoft SQL Server Management Studio GUI
- NetWorker User for SQL Server GUI

The following process occurs in a traditional backup:

1. The `nsrd` program starts the backup `nsrworkflow` on the NetWorker server.
2. The `nsrworkflow` starts the `savegrp` program.
3. The `savegrp` program runs the NMM backup command (`nsrsqlsv`) on the client instead of performing a standard NetWorker save.
4. The `nsrsqlsv` program passes the backup data from SQL Server to the NetWorker server through an X-Open Backup Services application programming interface (XBSA).

The NetWorker server schedules and performs all storage management tasks.

The following figure shows the traditional backup workflow.

Figure 3 Traditional backup workflow

The *NetWorker Administration Guide* provides information about the NetWorker services and operations.

Federated backup workflow (AlwaysOn Availability Group configuration)

During a federated backup, processes interact between the NetWorker client (that is, NMM), the NetWorker server, and the SQL Server.

Note

NMM supports federated backups of any SQL Server with support for AlwaysOn Availability Groups (SQL Server 2012 and later).

The backup starts from the `nsrsqlsv` program, which is started by using one of the following:

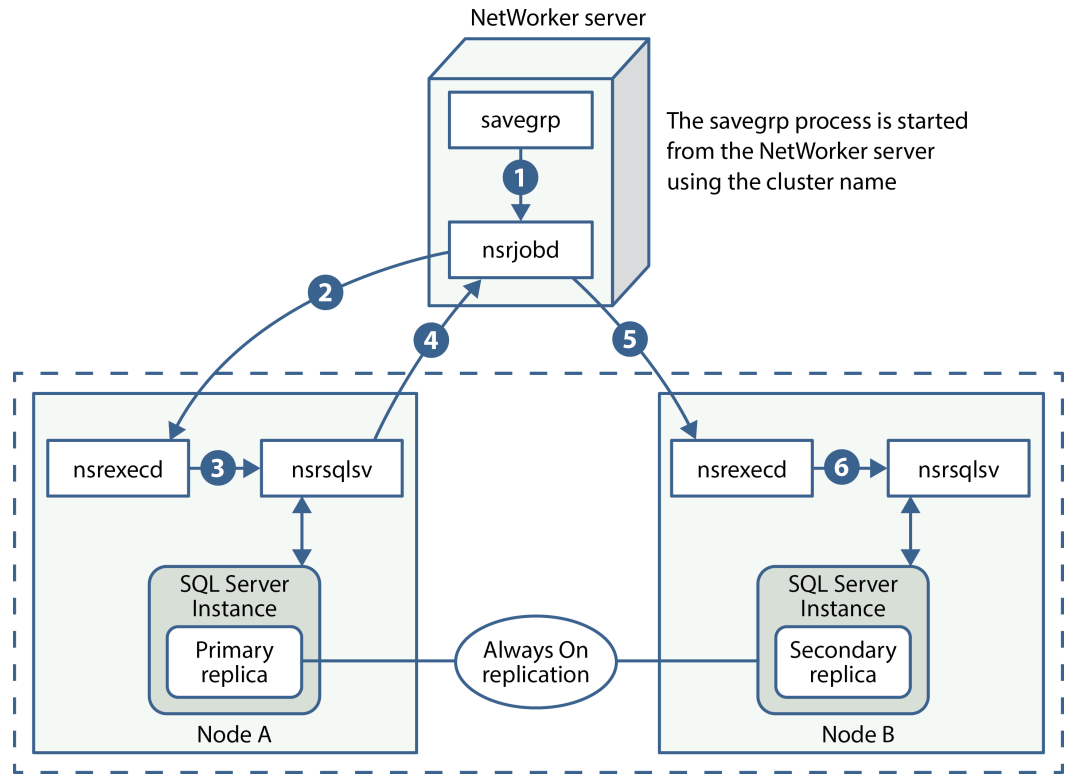
- Command prompt
- Microsoft SQL Server Management Studio GUI
- NetWorker User for SQL Server GUI

The following process occurs during a federated backup:

1. The NetWorker server starts the `nsrsqlsv` program in the active node of the Windows cluster (called the coordinator process).
2. The coordinator process queries the SQL Server and detects the Backup Preference and priority from the Availability group, and starts the worker process on the detected preferred node.
3. The backup is configured with and stored under the Windows cluster name on the Availability Group.

The following figure shows the federated backup workflow.

Figure 4 Federated backup workflow



Note

The coordinator process always goes through the `nsrjobd` service to start the worker process on the secondary nodes to ensure that the NetWorker server and `nsrjobd` can communicate with secondary nodes.

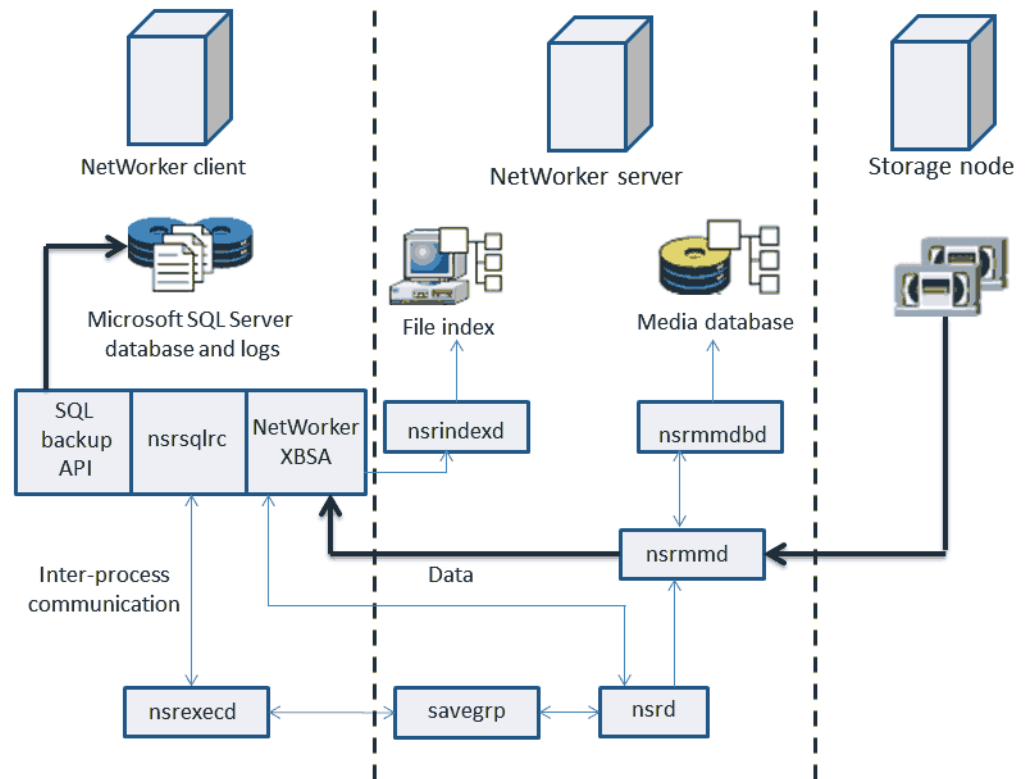
Recovery workflow

During a traditional recovery operation that uses a storage node without the Client Direct feature, process interactions occur between the NetWorker client (that is, NMM), the NetWorker server, and the SQL Server.

The following process occurs in a traditional recovery operation:

1. The `nsrsqlrc` program starts the recovery.
2. The NetWorker XBSA API translates the object names that NMM requests into a format that NetWorker understands and forwards the translated object names to the NetWorker server `nsrd` service.
3. The `nsrmmnd` media service, contacts the `nsrmmdbd` service to search the NetWorker server's media database for the volumes that contain the requested objects.
4. After the media is mounted, the `nsrmmnd` program sends the data through the NetWorker XBSA API to `nsrsqlrc`, which then sends data to the SQL Server.

The following figure shows the traditional recovery workflow.

Figure 5 Traditional recovery workflow

NMM recovery interactions with the SQL Server

When a SQL instance-level recovery occurs, NMM stops and starts the SQL Server and dependent services.

When you want to recover the SQL Server system database types like master and msdb, the `nsrsqlrc` program automatically stops and restarts the SQL Server services appropriately, as follows:

1. Before the recovery process begins, NMM stops the SQL Server and other dependent services.
When the SQL Server Analysis Services (SSAS) is running, it might use the only available database connection if the SQL Server is in a single-user mode. Stop the Analysis Services before restoring the master database.
2. NMM starts the SQL Server in single-user mode.
3. NMM performs the recovery.
4. After the recovery process finishes, NMM waits for the SQL Server to shut down.
5. For stand-alone and cluster environments, NMM restarts the SSAS.

When you recover a master database, there can be timing issues related to stopping and starting of services. If you are recovering a master database, it is recommended that before you start the recovery, you manually stop all SQL Server services except for the SQL Server.

SQL Server cluster environments

NMM can back up or recover data from an SQL server that is running as a SQL virtual server in a Windows Server Failover Cluster (WSFC). The SQL Server instance that is running in this configuration is called Failover Cluster Instance (FCI).

NMM requires the SQL virtual server name so it can perform the following tasks:

- Connect to the SQL Server instance.
- Accept data from or deliver data to the SQL Server in the cluster, and to initialize the SQL Server VDI.
- Create entries in the NetWorker client file index.

NMM creates index entries under the virtual server name in the NetWorker client file index.

Note

Backup and recovery of SQL Server data in a cluster with NMM requires Cluster Client Connection licenses on the NetWorker server host. A separate Cluster Client Connection license is required for each node in the cluster.

NMM detects SQL Server instances

NMM automatically detects all the SQL Servers in a WSFC, including the SQL virtual servers.

NMM detects the SQL Servers only on the active nodes, whenever the following conditions occur:

- The NetWorker User for SQL Server GUI is started.
- The NetWorker window in the NMM Microsoft SQL Server Management Studio (NMM SSMS) plugin GUI is opened.
- The Select SQL Instance option is selected.
- A backup or recovery is started.

Named instances in failover cluster configurations

NMM provides failover cluster support by using the multiple instance features provided in the SQL Server. In a failover configuration, the SQL virtual servers run as either default instances or as named instances. Only one default instance of a SQL virtual server can be installed. Additional SQL virtual servers might be installed as named instances, where each instance name must be unique within the cluster.

Multiple named instances are supported as SQL virtual servers in a cluster configuration. The number of instances that are supported depends on the SQL Server version that is used in the setup. The *NetWorker Module for Microsoft Installation Guide* and the Microsoft SQL Server documentation provide more information.

Each named SQL virtual server instance has the following qualities:

- A unique IP address, network name, and instance name
- Datafile files that are installed on a clustered drive that belongs to the same cluster group as the associated virtual server for that named instance

Active and passive cluster configurations

When the NetWorker User for SQL Server GUI or the NMM SSMS plugin GUI is started on the active node, NMM automatically uses the SQL virtual server as the

client name. The client name is used for reading or writing to the NetWorker media database and client file index. The NetWorker User for SQL Server GUI or the NMM SSMS plugin GUI can be used only on the active node.

When you use the command line interface, use the `nsrsqlsv -A <SQL virtual server>` backup command when the target database is a SQL cluster.

If the SQL Server fails over to the secondary cluster node and the NetWorker User for SQL Server program is started on the secondary node, NMM automatically uses the SQL virtual server name. The NMM SSMS plugin GUI shows all active clusters that are running on the host in the **SQL Server Instance** list.

If one SQL virtual server is running on each physical node in the cluster, an active or active cluster configuration exists, and the following events occur:

- If no failover occurs and each SQL virtual server is running on a different physical node, NMM automatically communicates with the SQL virtual server that is running on the same physical node.
- The **Select SQL Instance** option from the **Operation** menu on the NetWorker User for SQL Server GUI is disabled.

However, if one of the physical cluster nodes goes offline (for example, if a failover occurs) and both SQL virtual servers are running on the same physical cluster node when the NetWorker User for SQL Server GUI is started, the **Select SQL Instance** dialog box appears. The **Select SQL Instance** dialog box lists the SQL Server instances. After you select an instance from the **Select SQL Instance** dialog box, NMM tries to validate the instance as a NetWorker client.

SQL Server backups

This section introduces NMM as a tool to back up an SQL Server.

Types of supported backups

NMM supports manual and scheduled backups of SQL Server data.

Manual (traditional) backup

NMM supports traditional backups, which are often referred to as manual backups. A traditional backup of SQL data can be performed at any time and is independent of any scheduled backup.

NMM supports traditional backup of the following items:

- Database
- File
- File group
- Filestream data
- Transaction log

Also, NMM for SQL Server supports file group differential, file differential, and copy-only backups.

A file group differential backup can reduce both media requirements and recovery time because data is stored across more than one disk or disk partition, so recovery time is reduced. A differential backup can substitute for any logs-only backups performed between the full and differential backups. A full backup must be performed first.

Note

NMM can recover a full backup of SQL Server data (including files and file groups) that were created with NMSQL. However, NMM cannot recover snapshot (PowerSnap based) backups that were created with NMSQL.

Scheduled backup

The most reliable way of protecting SQL data is to ensure that backups of the SQL Server are run at regular intervals, that is, scheduled backups. Scheduled backups ensure that all SQL Server data, including the NetWorker server's client indexes and bootstrap file, is automatically saved. If a disaster occurs, the client indexes and bootstrap file are vital for restoring data to the SQL Server.

Backup limitations

Due to SQL Server behavior, the following limitations apply to SQL Server backups:

- For simple recovery model databases, only full backups (including copy-only full backups) are supported.
- For full backups of a secondary replica, SQL Server supports only copy-only backups.

Backup levels

Learn about the backup levels that NMM supports

Note

NetWorker 8.2.x defines backup levels differently than NMM 18.1. If you are using NetWorker 8.2.x server with the NMM 18.1 client, refer to the *NetWorker Module for Microsoft for SQL VDI User Guide* version 8.2 SP1 for information about backup levels.

NMM supports the following three levels of backup.

Full backups

Entire database backup, including all file groups or files in the database.

Cumulative incremental backups

A cumulative incremental backup captures all changes since the last full backup.

Logs-only backups

A logs-only backup that corresponds to a SQL Server transaction log backup. A log file backup cannot be used to recover a database.

A log file backup is used after a database recovery to restore the database to the point of the original failure.

The logs only backup option appears in the NetWorker Management Console. When using the command prompt to perform a logs-only backup, use the `txnlogs -l` command.

Note

When a logs-only backup is taken, NMM records two entries in the media database. One entry is the database save set, which lists the backup level as `txnlog`. The second entry is a cover set entry, which lists the backup level as `incr`. Both entries are required for cloning and restore operations.

Copy-only backups

You can take a manual SQL Server copy-only backup at any time without affecting the backup schedule or log chain. You can perform copy-only backups at the full or logs-only backup level.

Copy-only backups are not promoted to a different backup level, which allows other backups to run without disrupting the backup chain.

Copy-only full backups are not considered level full backups during promotion. Having a copy-only full backup does not prevent subsequent backups from being promoted to a full backup.

Setting backup levels

NMM enables you to specify backup levels to logs only, cumulative incremental, and full.

The availability of a backup level depends on the type of data selected for backup and any SQL Server settings on those objects, as listed in the following table.

Note

NetWorker 8.2.x defines backup levels differently than NMM 18.1. If you are using NetWorker 8.2.x server with the NMM 18.1 client, refer to the *NetWorker Module for Microsoft for SQL VDI User Guide* version 8.2 SP1 for information about backup levels.

Table 3 Backup levels for SQL Server data

SQL Server data objects	Supported SQL Server backup levels		
	Full	Cumulative incremental	Logs only
All databases of SQL default or named instances	Yes	Yes	Yes
Specified databases	Yes	Yes	Yes
All filegroups in specified databases	Yes	Yes	Not applicable
Filestream data in specified databases	Yes	Yes	Yes
Specified filegroups in specified database	Yes	Yes	Not applicable
Specified files in filegroups in specified databases	Yes	Yes	Not applicable

When you perform a logs-only backup for SQL Server data objects, ensure that the SQL Server database options are correctly configured. The Microsoft SQL Server documentation provides more information. Individual items are subject to promotion.

Example strategies for backing up SQL Server data

This section describes example strategies for backing up SQL Server data.

Example 1 Backup strategy one

Example 1 Backup strategy one (continued)

If the SQL Server manages a significant amount of data, schedule a backup of the databases every 1 to 2 weeks, as shown in the following table.

Table 4 Full backup every 1 to 2 weeks

Fri	Sat	Sun	Mon	Tues	Wed	Thurs
Full	Logs-only	Logs-only	Logs-only	Logs-only	Cumulative incremental	Logs-only
Logs-only	Logs-only	Logs-only	Cumulative incremental	Logs-only	Logs-only	Logs-only
Full	Repeat					

Example 2 Backup strategy two

Another backup strategy is to schedule logs-only backups on several successive days immediately following a full backup, as shown in the following table. This schedule backs up all data that has changed since the previous logs-only backup.

A level 1 cumulative incremental backup can also be scheduled after several days of logs-only backups, as shown in the following table. This schedule backs up all data since the previous full backup.

NOTICE

If a database is read-only, perform a full backup of the database. A read-only database cannot be restored from an existing transaction log backup.

Table 5 Logs-only backup after a full backup

Fri	Sat	Sun	Mon	Tues	Wed	Thurs
Full	Logs-only	Logs-only	Logs-only	Cumulative incremental	Logs-only	Logs-only
Repeat						

Using backup levels

Because it is not practical or efficient to run full backups every day, you can specify different backup levels for scheduled backups. Limiting the frequency of full backups can decrease server load while ensuring complete data protection.

Differences between backup levels

The following table outlines the differences between backup levels.

Table 6 Backup level advantages and disadvantages

Backup level	Advantages	Disadvantages
Full	Fastest restore time.	<ul style="list-style-type: none"> • Slow backup. • Increases load on client, server, and network. • Uses the most volume space.
Logs only	<ul style="list-style-type: none"> • Faster backup time than a full backup. • Decreases the load on server and uses the least volume of space. • Enables point-in-time restore. 	<ul style="list-style-type: none"> • Slow restore. • Data can spread across multiple volumes. • Multiple transaction logs can spread across multiple volumes.
Cumulative incremental	<ul style="list-style-type: none"> • Faster backup time than a full backup. • Captures all changes since the last full backup. 	Generally more time-consuming than a logs-only backup (depending on the backup schedule strategy).

Combining data objects to create backup levels

NMM enables the selection of SQL Server data objects in various combinations to create scheduled backups of different levels, as shown in the following table.

Table 7 Creating additional backup levels with data objects

Backup level	Database objects
Full database	To create a level full database backup of the selected databases and their transaction log files, select one or more databases.
Full file or file group	To create a level full file or file group backup of the selected files or file group, but not their transaction logs, select one or more files or one or more filegroups.
Database logs only	<p>To create a database logs only level backup of only the logs only for the selected database, select one or more databases.</p> <p>The SQL database must be previously configured to enable logs only backups.</p>
Database cumulative incremental (level 1)	To create a database level cumulative incremental backup of only the changes that are made to the selected databases since the last full level backup was created, select one or more databases.
File or file group cumulative incremental	To create a file or file group level cumulative incremental backup, select one or more files or one or more filegroups. This backup only includes the changes that are made to the selected files or filegroups since the last full level backup.

Backup promotion

Guidelines for Microsoft SQL Server best practices indicate that a full database backup should be the first step in implementing a recovery strategy for a database. In adhering to these guidelines, NMM supports backup level promotion. Backup level promotion is based on data on the NetWorker server and the SQL Server.

Note

NMM does not support backup promotion during copy-only backups.

The following table explains the scenarios which cause backup promotion to occur.

Table 8 Backup level promotion process

Backup item	Requested backup level	Level of promotion	Reason for promotion
Database	Cumulative incremental	Database full	<ul style="list-style-type: none"> • A full database backup does not exist. • A restore was done after the most recent full database backup. • The last database backup was not performed with NMM. • The database name is "master" or "msdb." • For Always On Availability Groups only, the last database backup was taken with a different save set. • Either the mirror partner has a more recent backup or its backup status cannot be determined.
	Logs only	Database full	<ul style="list-style-type: none"> • A full database backup does not exist. • A restore was done after the most recent full database backup. • The last database backup was not performed with NMM. • The database name is "master" or "msdb." • For Always On Availability Groups only, the last database backup was taken with a different save set. • Either the mirror partner has a more recent backup or its backup status cannot be determined. • The database model changed from simple to full or to bulk logged. • The database is using the simple recovery model. • The database is in emergency mode.
File/Filegroup	Full	Database full	<ul style="list-style-type: none"> • A full database backup does not exist. • A restore was done after the most recent full database backup. • The last database backup was not performed with NMM. • The database name is "master" or "msdb." • For Always On Availability Groups only, the last database backup was taken with a different save set. • Either the mirror partner has a more recent backup or its backup status cannot be determined.
	Cumulative incremental	Database full	
	Logs only	Database full	
	Logs only	Filegroup/file full	Logs-only backups of files or filegroups are not supported.

Toggle backup promotion functionality

You can toggle backup promotion for both scheduled and manual backups using the following tools:

- Scheduled backup:
 - In the Client Backup Configuration wizard, select **Turn off backup promotion** to turn off backup promotion.

- In the Client Properties dialog box, type the `NSR_BACKUP_PROMOTION` application information variable with a valid value.
- Manual backup:
 - At a command prompt, use the `BACKUP_PROMOTION` flag with a valid value with the `nsrsqlsv` command.
 - In the NMM SSMS plugin, on the **Options** page, select an option from the **Backup Promotion** list.

Note

Consider the following when you disable backup promotion:

- When backing up the 'msbd' and 'master' databases, the backup level is always set to full and the backup promotion setting is ignored.
 - The first backup of a database is always set to level full and the backup promotion setting is ignored.
 - If you disable backup promotion, data loss may occur if the backup chain is broken. For example, if a transaction logs backup is taken with a third party software or SQL native backup tools in between two NMM logs only backups for the same database, NMM may be unable to restore the database using the logs only backup.
-

The Manual Backups chapter and the Scheduled Backups chapter provide more information about changing backup promotion settings.

SQL Server recovery

This section introduces NMM as a tool to recover an Microsoft SQL Server.

Types of supported recovery

This section lists the types of supported recovery for SQL Server VDI.

NOTICE

NMM supports recovery of a SQL Server 2012 or later database only after the Always On Availability Group replication has been removed for the corresponding database.

The following table lists the types of recovery for SQL Server VDI in NMM.

Table 9 Types of recovery for SQL Server VDI

Type of recovery	When used	Description
Traditional recovery	For data that was backed up by traditional backup, NMM supports traditional recovery.	<p>Data recovery from a traditional backup can be performed:</p> <ul style="list-style-type: none"> • At any time with NMM. • By running NMM recover command (<code>nsrsqlrc</code>) from the command prompt. <p>Traditional recovery operations recover files, file groups, databases, and transaction log backups.</p>

Table 9 Types of recovery for SQL Server VDI (continued)

Type of recovery	When used	Description
Normal recovery	NMM uses the normal restore type as the default.	<p>The normal restore type restores:</p> <ul style="list-style-type: none"> • The entire set of data that is associated with one or more SQL Server backups, including full, logs only, and cumulative incremental backups. • A file, file group, or a database to the database originally backed up. • Level full, level 1 (cumulative incremental), and level logs only backups in the order required by SQL Server. <p>NMM can back up and restore specified files and file groups. Also, a single file group, or multiple file groups or files, can be restored from a full database backup.</p>
Copy recovery	A copy recovery is an operation in which data is recovered to a SQL Server host other than the one from which it was backed up. A copy restore from and to the same SQL Server instance can also be done.	<p>The copy recovery type creates a copy of a database by restoring a SQL Server database to a new location, or with a new database name. The copy recovery type makes it easy to duplicate a database that was previously backed up. You can only mark a single item for this operation. Also, you can copy a system database, but you cannot overwrite it. NMM does not support copy recovery of filestream databases.</p>
Flat file recovery	Flat file recovery allows you to restore the backup to a file.	<p>When you perform a flat file recovery, NMM writes the backup to files instead of directly to the SQL database. This feature allows you to restore the recovery files later without access to the network by using the standard T-SQL file restore command.</p>
Granular-level recovery (GLR)	Granular-level recovery allows you to recover SQL server data at the table level.	<p>By using GLR with SQL Server, you can recover individual tables to the production database. This feature reduces the space requirements on local system storage during a restore operation. Depending on the size of the content database, GLR may also reduce recovery time. Granular-level recovery is performed using NMM and ItemPoint for Microsoft SQL Server.</p>

Recovery modes

To recover a database, you must specify a recovery mode. A recovery mode instructs the SQL Server how to interact with the database after the recovery operation completes. For instance, recovery modes can leave the database in an intermediate state, so additional transaction logs can be applied.

The following table shows how the recovery modes correspond to SQL Server database restore options.

Table 10 Recovery modes

Types of recovery mode	Description
Normal	The normal restore mode instructs SQL Server to leave the database in an operational state after the restore completes. This state then enables database reads and writes. The normal restore mode is the default mode NMM uses when restoring a database.
No-recovery	The no-recovery restore mode activates the SQL Server NORECOVERY database restore option for the last stage that was restored. The no-recovery restore mode places the database in a state that cannot be loaded after the restore, but it is still able to process additional transaction log restore operations.
Standby	The standby restore mode activates the SQL Server STANDBY database restore option for the last stage that is restored, which forces the database to be in a read-only state between transaction log restore operations. The standby restore mode provides an undo file for SQL Server to use when rolling back the transactions.
Online	SQL Server provides the ability to perform a restore operation while a SQL Server database is active. The database is completely offline only while the primary file group is being restored. After the primary file group is restored, the database can be brought online while the rest of the file groups are being restored, and then only the data that is being restored is unavailable. The rest of the database remains available during this type of restore. Earlier versions of SQL Server require that you bring a database offline before you restore the database.

Recovery time

Backups can be recovered to a specific time. The recovery time controls which backup data should be reinstated when a database is recovered. The recovery time also

controls the portions of a logs only backup that must be recovered when NMM is instructed to discard transactions that are performed after a particular time.

The default or current recovery time for each database comes from the create time of the marked item. By default, the most recent backup is recovered. If the most recent backup is logs only level or 1, dependent backups are recovered first. User-specified recovery times can restore older backup versions or perform point-in-time recovery operations. For example, a point-in-time recovery is specified by using a recovery time that is earlier than the create time of the transaction log backup, but later than the create time of the previous backup.

NMM provides three methods for recovering to a specific time:

- Database backup versions
- Point-in-time recovery of a transaction log (level logs only) backup
- Recovering to a named log mark

Recovery window restrictions

In the **Recovery** window, the rules for marking an item are based on the selected restore type. The normal restore type does not restrict marking in any way. All restorable objects (file, file group, filestream data, database) can be marked. When the copy restore type is chosen, only one database object can be marked. Marking the root SQL Server item is not permitted, the file groups and files of the selected database are automatically marked and restored as part of the full database restore.

The recovery process

A recovery uses the following process.

1. NMM recovers the most recent full backup, and then recovers the most recent differential (level 1) backup (if any).
If a full database backup is removed from the NetWorker server, and an incremental backup is tried, the recovery fails. NMM checks the SQL Server instance to determine if a full database backup was performed, but does not verify if a full backup still exists on the NetWorker server.
2. NMM recovers all the transaction log backups that ran after the most recent differential backup (or that ran after the last full backup, if there was no differential backup). To correctly recover uncommitted transactions, the SQL Server NORECOVERY mode is specified for all intermediate transaction logs. The recovery of the final transaction log specifies the restore mode if STANDBY or NORECOVERY is selected. The default selection is Normal.

For example, if you selected a restore mode of NORECOVERY, that specification appears in the output for a database restore as follows:

```
nsrsqlrc -s NetWorker_server_name my_database

nsrsqlrc: Restoring database my_database...
nsrsqlrc: RESTORE database my_database FROM
virtual_device='BSMSQL' WITH norecovery, stats
nsrsqlrc: RESTORE database my_database from
virtual_device='BSMSQL' WITH norecovery (differential)
nsrsqlrc: RESTORE transaction my_database FROM
virtual_device='BSMSQL' WITH norecovery
nsrsqlrc: RESTORE transaction my_database FROM
virtual_device='BSMSQL' WITH norecovery
Received 1.0 MB 4 files from NSR server
```


Database file relocation restrictions

NMM imposes the following restrictions on database file relocation:

- Only database backups can be relocated. Individual file and file group backups cannot be relocated without relocating the database that contains those files.
- If the configuration of a database has changed since the most recent, level full database back up was created, you cannot relocate the database. Configuration changes include the deletion or addition of files, file groups, or transaction log files.
- A system database might not be the destination database of a relocation.
- The relocation fails if the destination does not have sufficient space to create a database.

SQL Server instance and database names for backup and recovery

SQL Server instance, database, and filegroup names are not case-sensitive. In other words, NMM does not distinguish the difference between upper and lowercase letters in SQL Server instance and database names.

Note

If there are two or more databases with the same name but with different capitalization, such as DB1 and db1, NMM views these databases as the same and by default backs up only one of the databases.

Named and default instances of SQL Server

NMM supports backup and recovery of named and default instances.

NMM supports recovery from the SQL Server default instance or named instances by using a copy restore to any instance of a SQL Server. This support includes recovery operations when the destination server is different from the source.

Each named instance has a unique instance name in the following format:

computerName\instanceName

where:

- *computerName* is the network name of the computer.
- *instanceName* is the logical name of the named instance.

Note

Consider the following information when naming a SQL database or an instance:

- Use unique names.
 - Do not use the name that the SQL Server uses to identify itself (MSSQL:).
 - Do not use the names of SQL instances that you have installed.
-

Use the following syntax to specify a SQL stand-alone named instance of SQL Server at a command prompt:

```
MSSQL$Standalone_Named_Instance:[dbName ...][.fgName ...]
[.fileName ...]
```

An entry of `mssql:` for the Save Set attribute during the client resource configuration yields a backup of all databases for the default instance.

When running multiple instances, the `nsrsqlsv` and `nsrsqlrc` commands support specification of only one instance at a time. If save sets for more than one instance are specified, the backup or restore operation fails.

Index entries for stand-alone named instances are created by using the local host on which the instance is running. Index entries for clustered named instances are created with the SQL virtual server name. To differentiate backups for the default instance and named instances, the index name has been extended to logically partition the index for a client.

All running named instances, except clustered instances and the default instance, are maintained in the client file index directory. This named instance directory is created at the end of each traditional backup. Run the `nsrinfo` command after a backup to verify that this directory was created. For example, type:

```
%SystemDrive% nsrinfo -V -L -n mssql client_name
```

Supported special characters in database names for NMM backup and recovery

The following table lists the special characters that may be used in SQL Server database names when performing backup and recovery in stand-alone, cluster, Always On Availability Group, and federated configurations.

Table 11 Supported special characters in database names

Special characters		Stand-alone and cluster configurations (Database-level and Instance-level backup and recovery)	Always On Availability Group and federated configurations (availability group level backup and recovery)
~	Tilde	Successful	Successful
-	Hyphen	Successful	Successful
!	Exclamation mark	Successful	Successful
{	Open curly bracket	Successful	Successful
%	Percentage	Successful	Successful
}	Close curly bracket	Successful	Successful
)	Close parenthesis	Successful	Successful
(Open parenthesis	Successful	Successful
`	Accent grave	Successful	Successful
@	At the rate	Successful	Successful
#	Hash	Successful	Fails

Table 11 Supported special characters in database names (continued)

Special characters		Stand-alone and cluster configurations (Database-level and Instance-level backup and recovery)	Always On Availability Group and federated configurations (availability group level backup and recovery)
			Note Hash is an availability group identifier.
_	Underscore	Successful	Successful
&	Ampersand	Successful	Successful
^	Caret	Successful	Successful
.	Period	Successful	Successful
\	Backslash	Successful	Successful
'	Apostrophe	Successful Note Backup and recovery operations are successful but warnings are displayed when performing recovery.	Successful Note Backup and recovery operations are successful but warnings are displayed when performing recovery.
\$	Dollar	Fails Note Dollar is a SQL instance identifier.	Fails Note Dollar is a SQL instance identifier.
:	Colon	Fails Note Colon is a database name identifier.	Fails Note Colon is a database name identifier.

CHAPTER 2

Configuration

This chapter contains the following sections:

- [Configuring NMM in a SQL VDI environment.....](#) 38
- [NMM 18.1 compatibility with NetWorker 8.2.3 or 8.2.4 servers 44](#)
- [Configuring permissions to perform NMM backup and recovery of Microsoft SQL Server..... 44](#)

Configuring NMM in a SQL VDI environment

This section provides information about configuring NMM in a SQL VDI environment.

Supported Windows Server and SQL Server versions

The *NetWorker Online Software Compatibility Matrix* at <http://compatibilityguide.emc.com:8080/CompGuideApp/> provides the most up-to-date information about supported Windows Server and SQL Server versions.

Migrating from VSS solution to VDI solution for SQL Server data protection

Review this section if you are using an NMM version earlier than NMM 2.4 to recover SQL Server data with VSS technology and want to use NMM 8.2 or later to back up and recover SQL Server data with VDI technology.

- If you are using NetWorker 8.2.x server with NMM 18.1 client, in the **Backup Group Properties** window, clear the **Snapshot** option.
- Assign a backup schedule to the backup.
- In the **Client Properties** dialog box for each client resource:
 - In the **Backup Command** field, type `nsrsqlsv`. For SQL virtual server in a cluster environment, specify `nsrsqlsv -A` SQL virtual server.
 - Leave the **Application Information** field empty.
 - Change the save set to `MSSQL:` for SQL default instance-level backup or `MSSQL:dbname` for database-level backup on a SQL default instance.

Multi-stream Data Domain Boost

NMM supports multi-stream backups for a SQL Server to a Data Domain device. This support leverages the Data Domain Boost feature. SQL Server multi-stream backups over Data Domain Boost enhance the performance by running the backups much faster.

Microsoft SQL Server Always On Availability Group feature

NMM supports federated backups, during which NMM detects the SQL Server preferred backup setting for the Availability Group and performs the backup at the preferred node.

The Microsoft website describes the Always On Availability Group feature in SQL Server 2012 or later and provides detailed information about how to configure the setup for this support. This feature allows multiple replicas of a database.

An availability group is a logical group of databases that has the Always On capability. An availability group is failed over to other nodes as a group, meaning that all the databases that are part of the availability group are failed over together during a failover.

An availability replica hosts each set of availability databases.

Two types of availability replicas exist:

- A single primary replica, which hosts the primary databases.
- One or more secondary replicas, each of which hosts a set of secondary databases and serves as a potential failover target for the availability group. You can

configure the following number of secondary replicas depending on the SQL Server version:

- One to four secondary replicas for SQL Server 2012.
- One to eight secondary replicas for SQL Server 2014 or later.

The secondary replicas can be configured in either synchronous or asynchronous mode.

Consider the following points when you use Always On Availability Groups with NMM:

- To back up secondary replicas with NMM, in the **Availability Group Properties** window, set the **Readable Secondary** configuration option to either **Yes** or **Read-intent only** for each of the primary and secondary replicas. This option allows NMM to connect to the secondary replica to gather information about the secondary replica (for example, database file location, which can be different from the other replicas).
- To recover an Always On Availability Group database, suspend replication before you recover the database. You can use either the NMM Microsoft SQL Server Management Studio plugin GUI or the T-SQL query window for this task.

It is recommended that you back up an Always On Availability Group with a federated backup workflow, and that you use the Windows cluster name as the client name.

Always On Availability Group failover cluster instance support

NMM supports Always On Availability Groups over two normal SQL clusters. This configuration protects data with both SQL cluster instance-level failover and availability group database-level failover.

Note

When you configure a federated backup of an Availability Group that is configured with multiple SQL server instances and one of these instances is a failover cluster instance, use the **Client Properties** dialog box to configure the backup. This situation exists when there is a SQL virtual server instance in the cluster. This configuration is not available from the Client Backup Configuration wizard. See the Federated Backup and Recovery chapter for more information about configuring an availability group over two SQL clusters.

Availability group listeners

NMM supports scheduled and manual federated backups of SQL availability groups through the availability group listener.

Availability group listeners are virtual networks that provide access to an Always On availability group and all of its databases. You can use NMM to backup and recover an availability group using the listener name instead of the cluster name.

Unlike a cluster, when you use the listener name as the backup client name, NMM backs up only the availability group that the listener is configured for.

NOTICE

When you configure backups of an availability group using the listener name, if you later decide to configure a client for the availability group using either the cluster name or a different listener name, take a full backup with the new client. The same is true in reverse if you decide to change a client using the cluster name to the listener name. NMM cannot restore log-only backups that are taken after the client name changes.

SQL Client Direct to AFTD or DD devices

The NetWorker client software enables clients with network access to Advanced File Type Device (AFTD) or Data Domain devices to send their backup data directly to the devices. This functionality uses the Client Direct feature to allow the client to bypass the NetWorker storage node.

The Client Direct feature is enabled by default, but it can be disabled on each client by clearing the Client Direct attribute on the client resource in NMC. When a Client Direct backup is not performed, a traditional storage node backup is performed instead.

The `nsrsqlsv.log` backup log file displays details about the Client Direct activity for the SQL Server.

Microsoft SQL Server named log marks

Microsoft SQL Server enables enhanced point-in-time restore operations because it allows named log marks to be specified during transactions.

Database applications create named log marks when transactions are performed. The marks enable you to access to specific transaction points in a database transaction log backup. NMM allows you to select a named log mark during a restore operation and restores to the beginning or end of a named log mark during a database restore operation. Restoring data through named log marks is an improvement over point-in-time restore because the time that is associated with restoring to a specific transaction can be more accurately determined.

When a named log mark is created in the SQL Server database, the log mark time is saved to the millisecond. However, NetWorker's time format, which is used to specify point-in-time restore, supports granularity only to the second. If log marks with duplicate names are created within a second of each other, NMM restores to the most recently named log mark.

Managing SQL Server database transaction logs

NMM provides implicit management of the SQL Server database transaction logs.

Implicit management uses SQL database transaction log backups to manage log space. This management can occur when:

- A backup schedule is implemented that is level logs only.
- You run the `nsrsqlsv` command with the `-l txnlog` option.

Note

If you are using NetWorker server 8.2.x, the backup level is incremental and the `-l incr` option is used with the `nsrsqlsv` command.

Prevent log overflow

In Windows, prevent database logs from overflowing available log space by creating an alert in the SQL Server Performance Monitor. The alert forces a log to backup only when the database's log space reaches a certain capacity (for example, 80% full). An alert is a user-defined response to a SQL Server event. Backups truncate the logs and clear disk space.

Database consistency checks

NMM can provide database consistency checks (DBCC) before a backup operation is performed.

A DBCC examines all tables in the database to detect whether index and data pages are correctly linked and indexes are in proper-sorted order. A DBCC also checks that pointers are consistent and that the data information on each page and page offsets are reasonable. It helps recognize problems early, which prevents problem escalation and possible data loss.

Note

DBCC can be performed for database-level, multiple database-level, and instance-level save sets on cluster and stand-alone configurations. For federated configurations, DBCC can be performed for database-level save set, but DBCC cannot be performed for Always On Group level backups.

You can configure a DBCC through the following methods:

- For manual backups—From the NMM Microsoft SQL Server Management Studio (SSMS) plugin GUI, in the **NetWorker** window. The topic [Performing manual backups from the NMM SSMS plugin GUI](#) on page 63 provides more information.
- For manual backups—From the NetWorker User for SQL Server GUI:
 - On the **Database Properties** page, select the **Perform DBCC consistency checks** option by right-clicking a database in the **Backup** page.
 - In the **Backup options** window, on the **General** tab, select the **Perform DBCC consistency checks** option.
The topic [Performing manual backups from the NetWorker User for SQL Server GUI](#) on page 80 provides more information.
- For scheduled and federated backups—In the NetWorker Management Console, in the **Client Properties** dialog box for the SQL Server host client resource:
 - In the **Application Information** field, type `NSR_CONSISTENCY_CHECKS=database1,database2` where *database1* and *database2* are the database names for which you want to run a consistency check, for example: `NSR_CONSISTENCY_CHECKS=testdb3,CopyOfDB010,test4`. Ensure that there are no spaces between the database names. For example, `NSR_CONSISTENCY_CHECKS=testdb3, CopyOfDB010,test4` is incorrect.
 - To specify that you want to run a consistency check on all the databases in an instance, in the **Application Information** field, type `NSR_CONSISTENCY_CHECKS=ALL`.
- In the **Client Backup Configuration** wizard, select **Perform DBCC consistency check** when you configure the client resource, and then select one of the following:
 - To perform DBCC for all the databases in an instance, select **All**.
 - To perform DBCC for specific databases in an instance, specify the databases for which you want to run the DBCC checks.
- From the command prompt, type the `nsrsqlsv` command with the various backup command options as follows:
 - To perform DBCC checks for databases in an instance, type the command option `-j testdb3,CopyOfDB010,test4`.

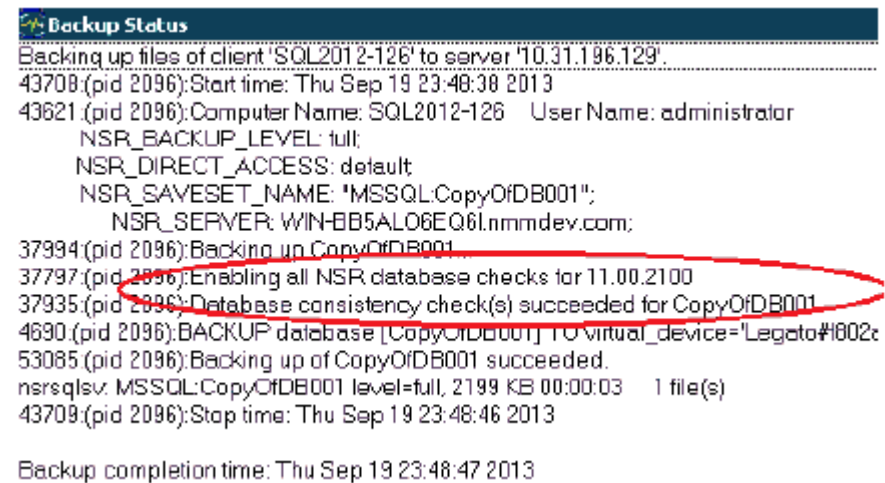
Where *testdb3*, *CopyOfDB010*, and *test4* are the names of the databases in the instance. Ensure that there are no spaces between the database names. For example, using *testdb3*, *CopyOfDB010*, *test4* is incorrect.

- To perform DBCC checks for all the databases in an instance, type the command option `-j ALL`.

If the DBCC runs successfully for a database, a message appears in the NMM logs or backup output page.

The following figure shows an example of the message that appears when the DBCC is successful.

Figure 6 Message showing DBCC was successful



```

Backup Status
Backing up files of client 'SQL2012-126' to server '10.31.196.129'.
43708 (pid 2096): Start time: Thu Sep 19 23:48:38 2013
43621 (pid 2096): Computer Name: SQL2012-126  User Name: administrator
    NSR_BACKUP_LEVEL: full;
    NSR_DIRECT_ACCESS: default;
    NSR_SAVESET_NAME: 'MSSQL:CopyOfDB001';
    NSR_SERVER: WIN-BB5AL06EQ6L.nmmdev.com;
37994 (pid 2096): Backing up CopyOfDB001.
37797 (pid 2096): Enabling all NSR database checks for 11.00.2100
37935 (pid 2096): Database consistency check(s) succeeded for CopyOfDB001
4690 (pid 2096): BACKUP database [CopyOfDB001] TO virtual_device='Legato#002:
53085 (pid 2096): Backing up of CopyOfDB001 succeeded.
nrsqslv: MSSQL:CopyOfDB001 level=full, 2199 KB 00:00:03  1 file(s)
43709 (pid 2096): Stop time: Thu Sep 19 23:48:46 2013

Backup completion time: Thu Sep 19 23:48:47 2013

```

If DBCC for a database fails, the failure status appears as part of the policy notification and the database backup is omitted. The DBCC failure is reported in the NMM logs or backup output page.

Microsoft hybrid cloud environments

NMM supports SQL Server 2014 and later VDI backups and recoveries in Microsoft hybrid cloud environments.

NMM supports stand-alone, and federated backup workflows for both private and hybrid cloud environments.

Backup and recovery operations in a cloud environment are only supported for onsite databases. The onsite backup workflow is the same as the SQL Server 2012 or later Always On Availability Group workflow.

Note

NMM does not support backup and recovery of SQL Server data directly from the hybrid cloud. If you are using the SQL Server 2016 Stretch Database feature, ensure the data that you must back up and recover is onsite, otherwise the operation fails.

Transparent data encryption

Microsoft SQL transparent data encryption (TDE) is a feature that performs realtime I/O encryption and decryption of the data and log files.

TDE uses a database encryption key (DEK), which is stored in the database boot record for availability during recovery. Encryption of the database file is performed at

the page level. The pages in an encrypted database are encrypted before they are written to disk and decrypted when read into memory. When using this feature, ensure that the certificate and private key are backed up with the encrypted data.

Note

NMM does not support third party transparent data encryption for SQL VDI.

Microsoft SQL Server 2008 and later introduce the TDE database-level encryption feature. This feature provides protection for the entire database at rest, without affecting existing applications. NMM supports SQL data encryption at the cell level, at the full database level by using TDE, or at the file-level with encryption options provided by Microsoft.

The Microsoft SQL Server product documentation provides more information about TDE, enabling data encryption, and protecting the encryption keys.

Note

When enabling TDE, back up the certificate and the private key associated with the certificate. If the certificate becomes unavailable or if the database is restored on another server, backups of both the certificate and the private key must be available to open the database.

Setting the MAXTRANSFERSIZE environment variable

`MAXTRANSFERSIZE` is a data transfer option that specifies the largest unit of transfer in bytes to be used between SQL Server and the backup media or VDI backup application. This option applies to both backup and recovery operations.

The `MAXTRANSFERSIZE` value is set through the `NSR_VDI_MAXTRANSFERSIZE` environment variable. Setting the `NSR_VDI_MAXTRANSFERSIZE` variable determines the `MAXTRANSFERSIZE` value as follows:

$MAXTRANSFERSIZE = 1024 \times NSR_VDI_MAXTRANSFERSIZE$

The valid values of the `NSR_VDI_MAXTRANSFERSIZE` variable range from 64 to 4096. The resulting values of the `MAXTRANSFERSIZE` variable range from 65,536 bytes to 4,194,304 bytes.

Use the following procedure to change the `MAXTRANSFERSIZE` from the default value of 4,194,304 bytes.

Procedure

1. Open **Control Panel > System**.
2. Click **Advanced system settings**.
The **System Properties** window opens.
3. Click **Environment Variables...**
The **Environment Variables** window opens.
4. In the **User variables** section, perform one of the following:
 - If an `NSR_VDI_MAXTRANSFERSIZE` variable is not listed, click **New** to create and define the `NSR_VDI_MAXTRANSFERSIZE` variable.
 - If an `NSR_VDI_MAXTRANSFERSIZE` variable is listed, click **Edit** to edit the variable.
5. Click **OK**.

Results

After this parameter is set in the registry, subsequent NMM backups use the registry setting. For SQL cluster environments, set the key on all cluster nodes where backups might run.

NMM 18.1 compatibility with NetWorker 8.2.3 or 8.2.4 servers

NMM supports backup and recovery with NetWorker client version 18.1 and NetWorker server version 8.2.3 or 8.2.4.

The *NetWorker Module for Microsoft Installation Guide* contains the NMM support matrix for NetWorker server and client versions. Refer to the individual NMM release sections of the *NetWorker Online Software Compatibility Matrix* at <http://compatibilityguide.emc.com:8080/CompGuideApp/> for more details.

Note the following limitations when you configure NMM backup and recovery with an NMM 18.1 client and a NetWorker 8.2.3 or 8.2.4 server:

- **Dedicated Storage Node:** NetWorker 8.2.3 and 8.2.4 servers do not support NetWorker storage node 18.1. As a result, you cannot configure a dedicated storage node when you use NetWorker 18.1 client with NetWorker 8.2.3 or 8.2.4 server.
- **Backup levels:** NetWorker 8.2.3 and 8.2.4 servers use NetWorker server 8.x backup-level definitions, and do not support the NetWorker server version 9.x and later backup levels.

Configuring permissions to perform NMM backup and recovery of Microsoft SQL Server

Review the privileges that are required to perform Microsoft SQL Server backup and recovery operations with NetWorker Module for Microsoft (NMM).

Access privileges for backup and recovery

NMM for Microsoft SQL Server VDI requires that the user starting backup and recovery operations is assigned certain privileges from the Microsoft SQL Server, the Windows application host, and the NetWorker server.

These privileges are required for both scheduled and manual backups:

- For manual backups, the user starting the backup must be granted the required privileges.
- For scheduled backups, the remote user who is assigned to the client resource must be granted the required privileges.

Microsoft documentation provides additional information on how to configure user accounts.

Table 12 Access privileges required for backup and recovery operations

Microsoft SQL Server user roles	Windows user groups	NetWorker user roles
<p>Assign the user to the following server roles with the Microsoft SQL Server Studio Management (SSMS) GUI:</p> <ul style="list-style-type: none"> • sysadmin • public <p>Assign the NT AUTHORITY \SYSTEM user to the following server roles with the Microsoft SSMS GUI:</p> <ul style="list-style-type: none"> • sysadmin • public 	<p>Assign the user to the following user groups on the Windows application host on all SQL nodes:</p> <ul style="list-style-type: none"> • Backup Operators • Administrators <hr/> <p>Note</p> <p>The user must be a domain user for failover cluster or availability group backups. You can use a local user for standalone backups.</p>	<p>Assign the user to the Operators role in the NetWorker Administration GUI.</p> <p>The <i>NetWorker Administration Guide</i> provides more information about assigning NetWorker User roles.</p>

Assign SQL server roles for backup and recovery operations

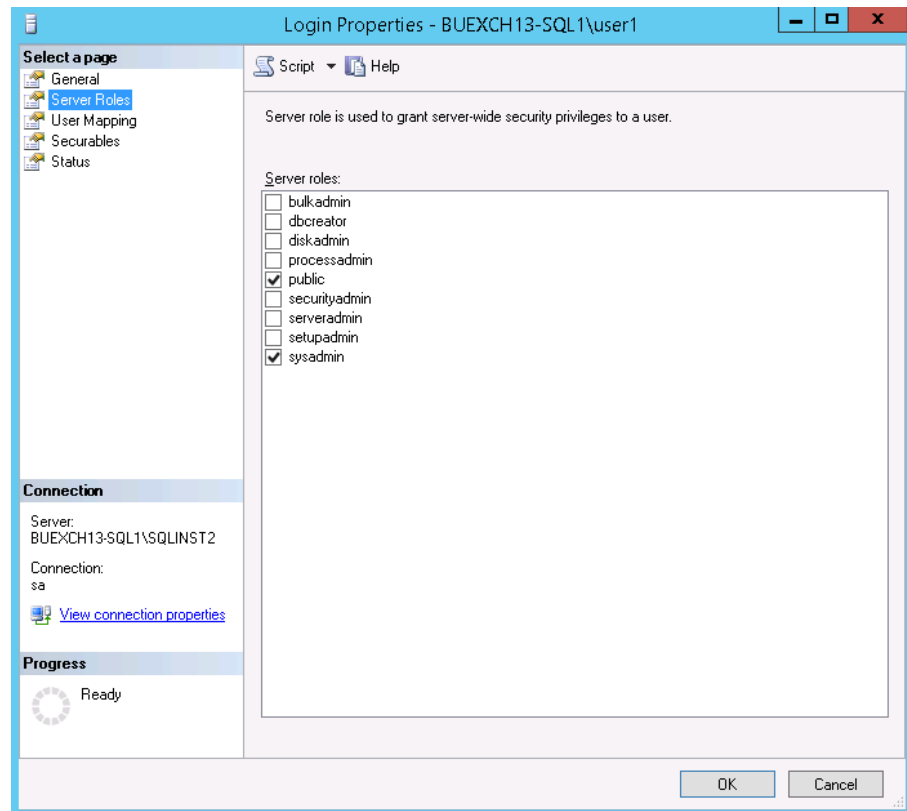
Assign the required privileges on the SQL Server to perform SQL Server backup and recovery operations.

Before you begin

A user account must exist on the SQL Server.

Procedure

1. Launch the SQL Server Management Studio.
2. In the **Connect to Server** window, specify the details for the SQL Server that you want to backup and recover. Ensure that you select **Windows Authentication**.
3. From the **Object Explorer** pane, expand the folder for the SQL Server that you want to backup and recover, then expand the **Security** folder, and then expand the **Logins** folder.
4. Right-click the name of the user that you want to assign backup and recovery privileges to, and then click **Properties**.
The **Login Properties** window appears.
5. On the **Server Roles** page, ensure that **sysadmin** and **public** are selected, and then click **OK**.

Figure 7 Assigning SQL Server privileges

6. Return to the **Logins** folder.
7. Right-click the "NT AUTHORITY\SYSTEM" user, and then click **Properties**.
The **Login Properties** window appears.
8. On the **Server Roles** page, ensure that **sysadmin** and **public** are selected, and then click **OK**.

Assigning Windows user privileges for backup and recovery operations

Assign the required user privileges on the Windows application host to perform SQL Server backup and recovery operations with NMM.

There are two types of users of which access privileges can be modified:

- **Local User**
This type of user is used for standalone server and Always On Availability Group databases.
- **Domain User**
This type of user is used for Always On Availability Group and Failover Cluster Instance databases.

Assign user privileges to a Local User

Local user privileges are modified using the **Local Users and Groups** window.

Procedure

1. On each SQL Server node that you want to back up and recover, click **Start > Programs > Administrative Tools > Computer Management > Local Users and Groups**.

2. In the left pane, under **Local Users and Groups (Local)**, perform one of the following action sequences.
 - To assign privileges to an existing user, use the following steps:
 - a. Click **Users**.
 - b. From the list of Local Users, right-click the user, and then click **Properties**.
 - To assign privileges to a new user, use the following steps:
 - a. Select and right-click **Users**, and then click **New User**.
 - b. In the **New User** window, specify the details for the new user, and then click **Create**.
The user appears in the list of Local Users in the **Users** folder.
 - c. Right-click the newly created user, and then click **Properties**.

The **User Properties** window appears.
3. On the **Member Of** tab, add the user to the **Backup Operators** and **Administrators** user groups, and then click **Apply**.
4. Click **OK**.

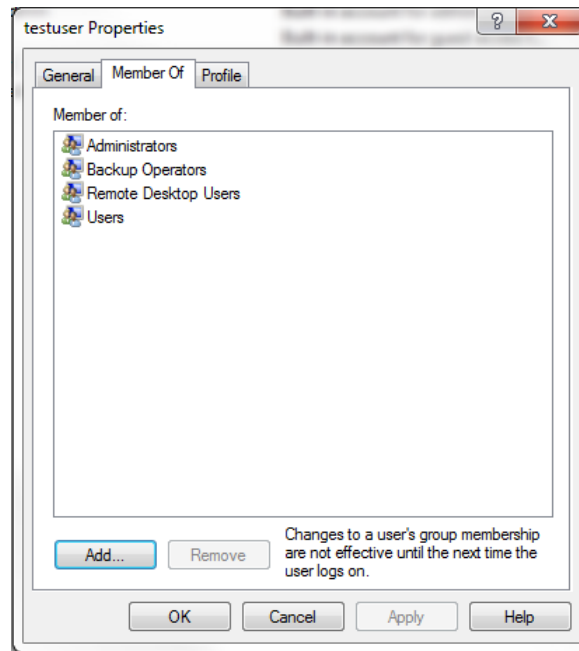
Assign user privileges to a Domain User

Domain user privileges are modified using the **Active Directory Users and Computers** window.

Procedure

1. On the domain controller, click **Start > Programs > Administrative Tools > Active Directory Users and Computers**.
2. In the left pane, under **Active Directory Users and Computers**, expand the Active Directory domain name, and then perform one of the following action sequences:
 - To verify that an existing user is a domain user, use the following steps:
 - a. Click **Users**.
 - b. From the list of Active Directory users, right-click the user, and then click **Properties**.
 - c. On the **Member Of** tab, verify that the user is listed as a member of the Active Directory domain.
 - To create a new domain user, use the following steps:
 - a. Select and right-click **Users**, and then click **New > User**.
The **New Object - User** window appears.
 - b. In the **New Object - User** window, select the Active Directory domain for the account and specify the details for the new user, and then click **Next**.
The user appears in the list of Active Directory users in the **Users** folder.
3. In the **Users** folder, right-click the user that you want to use for NMM backup and recovery, and then click **Properties**.
4. In the user **Properties** window, on the **Member Of** tab, click **Add**.
5. Add the user to the **Backup Operators** and **Administrators** user groups, and then click **Apply**.

The user groups are listed on the **Member Of** tab.

Figure 8 Adding a user to Windows User Groups

Enable user access for NMM when User Access Control (UAC) is used

Grant the Windows "Log on as a batch job" privilege to the remote user that performs NMM operations. This privilege allows the user to log in with a privileged security token

Note

These steps must be performed on each node in the SQL Server cluster.

Procedure

1. Open the Local Security Policy (secpol.msc) on the client
2. Browse to **Local Policies > User Rights Assignment**.
3. Add the Backup Operators and Administrators user groups to the **Log on as a batch job** privilege.
4. Click **OK**.

CHAPTER 3

Graphical User Interfaces

This chapter includes the following sections:

- [User interfaces for backup and recovery](#)..... 50
- [NetWorker User for SQL Server GUI](#)..... 51
- [NMM Microsoft SQL Server Management Studio plugin GUI](#)..... 55

User interfaces for backup and recovery

This section describes the graphical user interfaces (GUIs) where you can perform backup and recovery operations.

- The NetWorker Administration GUI on the NetWorker server—Start NMC on the NetWorker server and open the NetWorker Administration GUI to configure, perform, monitor, and report on scheduled backups.
- NetWorker User for SQL Server GUI—On the host where NMM is installed, select **Start > Programs > NetWorker Modules > NetWorker User for SQL Server** to start the NetWorker User for SQL Server GUI. The NetWorker User for SQL Server GUI has specific views, display conventions, and marking semantics and restrictions that you can use to perform manual backup and recovery operations.
- NetWorker window in the NMM Microsoft SQL Server Management Studio plugin GUI—Start the NMM Microsoft SQL Server Management Studio (NMM SSMS) plugin GUI to view the **NetWorker** window. You can use the **NetWorker** window to perform manual backup and recovery operations.

The following table summarizes the locations from which you can start backup operations, and the backup levels that are supported for each interface.

Table 13 Where to start backup operations

Backup type	Backup started from	Available backup levels		
		Full	txnlog	Diff
Scheduled	The NetWorker Administration GUI on the NetWorker server	Yes	Yes	Yes
Manual	Command prompt on the SQL Server, which is the client	Yes	Yes	Yes
	NetWorker User for SQL Server GUI	Yes	No	No
	In the NMM SSMS plugin GUI, in the NetWorker window	Yes	Yes	Yes

The following table summarizes the locations from which you can start recovery operations.

Table 14 Where to start recovery operations

Recovery type	Recovery started from
Full	NetWorker User for SQL Server GUI
	In the NMM SSMS plugin GUI, in the NetWorker window

NetWorker User for SQL Server GUI

This section provides information on the various views in the NetWorker User for SQL Server GUI, display conventions for SQL data types, marking semantics and restrictions, and fake objects.

Note

Support for the NetWorker User for SQL Server GUI is deprecated. It is recommended that you use the NMM SQL Server Management Studio plugin GUI for all manual backup and recovery operations.

Using the nwmssql command

The `nwmssql` command starts the NetWorker User for SQL Server GUI.

To run the NetWorker User for SQL Server GUI from the Windows Start menu, click **Programs > NetWorker > NetWorker User for SQL Server**.

To create a desktop shortcut, go to the `<install_path>\nsr\bin` directory and drag the `nwmssql.exe` file to your desktop while pressing the [Ctrl] key.

Command options for nwmssql

This section provides the list of options available for the `nwmssql` command.

Table 15 Command options for the `nwmssql` command

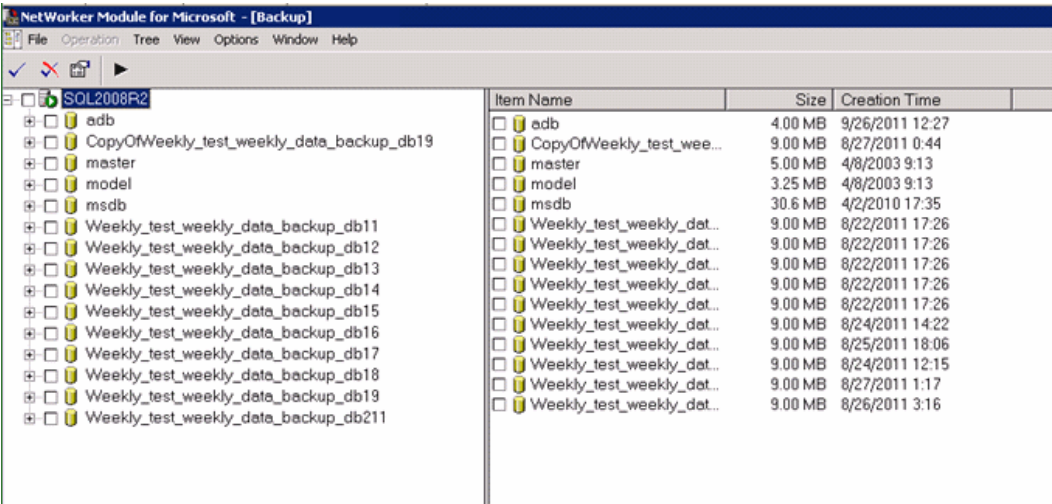
Command options for <code>nwmssql</code>	Description
<code>-s</code>	Specifies the NetWorker server to use.
<code>-t</code>	Diagnoses the current backup or recovery command issued by NMM. The NetWorker User for SQL Server GUI displays the full backup or recovery command in the operation status window, but the GUI does not run the backup or restore operation.
<code>-b</code>	Limits the browse time to what is mentioned in the <code>-b</code> command option.


Views

The NetWorker User for SQL Server GUI has four windows.

- **Backup**—All backups are performed from the **Backup** window. Select the database that must be backed up, and either right-click or select **Options > Backup Options** to view the **Backup Options** window. The following figure shows the **Backup** window.

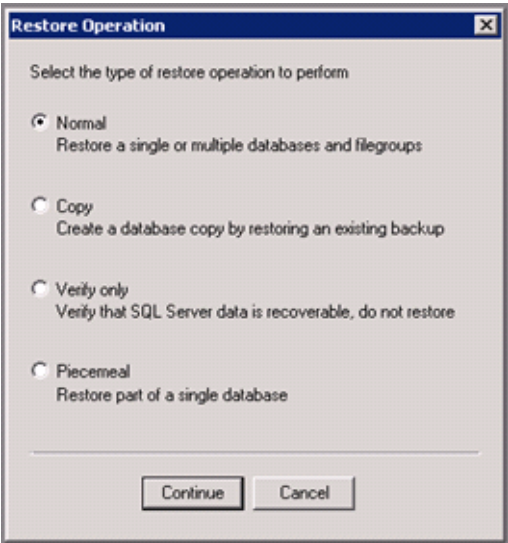
Figure 9 Backup window




- **Restore Operation**—Normal, copy, or piecemeal restores are performed from the **Restore Operation** window. Select the  icon to view the **Restore Operation** window.

The **Restore Operation** window is shown in the following figure.


Figure 10 Restore Operation window



- **Change Server**—You can choose the NetWorker server from the choices available from this window. Select the  icon to view the **Change Server** window.

The **Change Server** window is shown in the following figure.

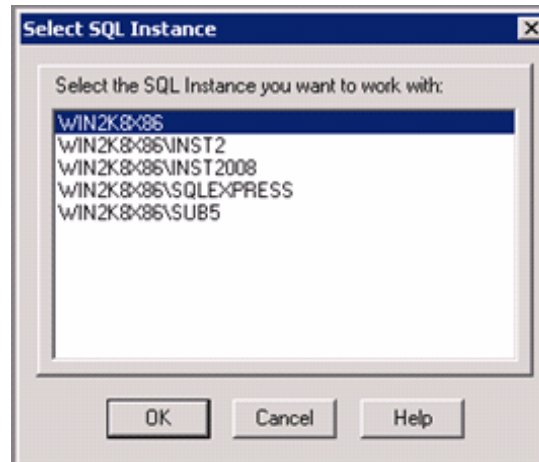
Figure 11 Change Server window



- **Select SQL Instance**—This window allows you to select the SQL instances that you want to work with. Select the  icon to view the **Select SQL Instance** window.

The **Select SQL Instance** window is shown in the following figure.

Figure 12 Select SQL Instance window



Display conventions for SQL data types

The NetWorker User for SQL Server GUI uses specific data item names, text characteristics, and icons to distinguish the variable qualities of SQL Server data.

The following table outlines these conventions.

Table 16 SQL Server storage hierarchy display conventions







Icon	Data item	Description
	SQL Server	<ul style="list-style-type: none"> • Root of the storage hierarchy • Signifies all SQL Server databases on the host

Table 16 SQL Server storage hierarchy display conventions (continued)

Icon	Data item	Description
	Database	<ul style="list-style-type: none"> • Descendant of root • Signifies a database • Might contain file groups
	File group	<ul style="list-style-type: none"> • Descendant of a database • Signifies a database file group or filestream data • Might contain files
	Fake file group	Signifies that the file group cannot be selected for backup
	File	<ul style="list-style-type: none"> • Descendant of a file group • Signifies a database file
	Fake file	Signifies that the file cannot be selected for backup

Marking semantics and restrictions

To support browsing capabilities, the NetWorker User for SQL Server GUI imposes certain semantics and restrictions on how items can be marked.

Whether an item can be marked depends on that item's predecessors and descendants mark status. Depending on what is marked, message dialog boxes might provide additional information on the current marks and the type of operation.

The NetWorker User for SQL Server GUI enables the following mark and unmark operations:

- Mark a single file, file group, or database.
- Mark multiple, heterogeneous items.
- Mark an item when any of that item's surrounding items are already marked.
- Mark or unmark all SQL Server data by right-clicking the root item and selecting Mark All Databases or Unmark All Databases from the shortcut menu.
- Unmark all databases from the SQL Server root.

The NetWorker User for SQL Server GUI includes the following restrictions:

- You cannot mark an item if any of the predecessors or descendants are already marked except in the SQL Server root.
- When a database is marked, all the item's descendants are not automatically marked.
- When a file group is marked, all the files are not automatically marked.

Fake objects

NMM enables you to browse file groups and files that are contained in a database. However, the software can display data items that are unavailable for backup or restore operations. These items are referred to as fake objects.

When a file or file group exists in the SQL Server storage hierarchy, but SQL Server settings on the database prevent its backup, the item displays in the **Backup** window with the fake file group or fake file convention. If you try to mark a fake object, an error dialog box appears.

The following figure shows the error message for a fake object.

Figure 13 Marking a fake object



NMM Microsoft SQL Server Management Studio plugin GUI

You can use the NMM Microsoft SQL Server Management Studio (SSMS) plugin GUI to perform manual backup and recovery operations.

Perform backup and recovery operations through the **NetWorker** window, with the **Backup**, **Database Restore**, and optional **Table Restore** tabs in the **NetWorker** window.

To use this feature, select the NMM SSMS Plugin GUI option during NMM installation. To enable the **Table Restore** tab, select the SQL Granular Recovery option whenever you install NMM.

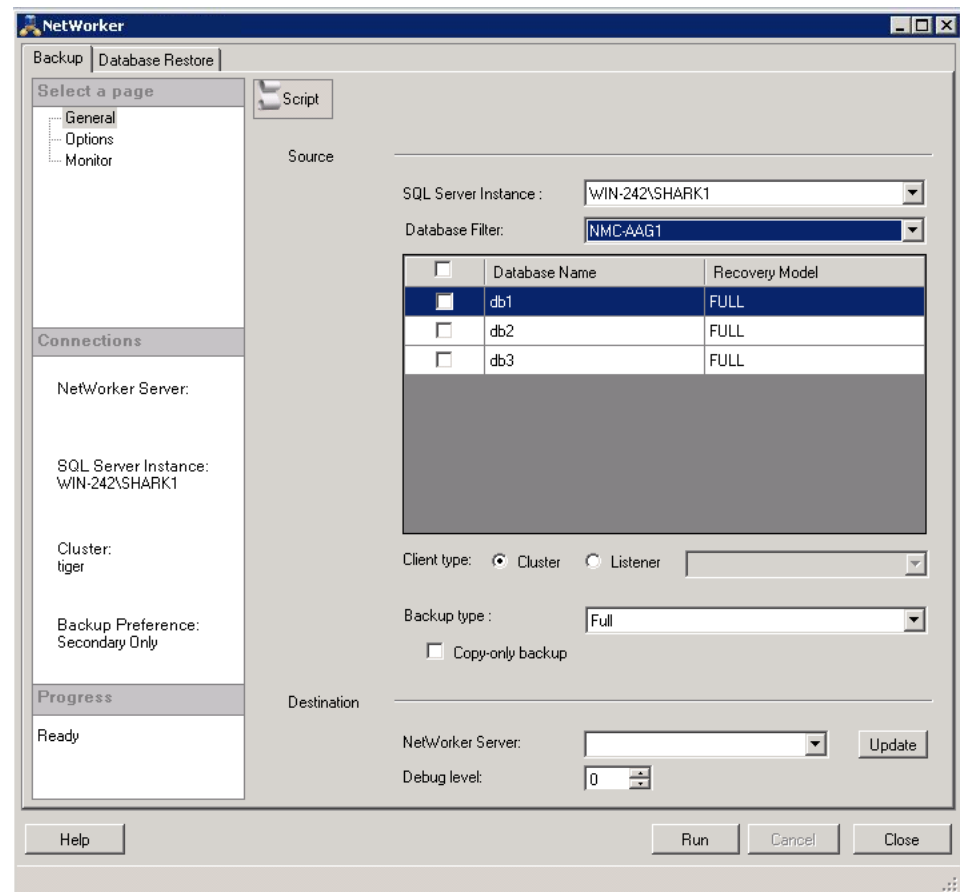
Note

If you have not installed the SSMS option and want to launch the NMM SSMS plugin, navigate to the `<install_path>\nsr\bin\` folder, and then double-click the `nwssmsaddinapp.exe` file.

Views

The **NetWorker** window in the NMM SMSS plugin GUI has two main tabs and a third optional tab with multiple pages in each.

The following figure shows the **NetWorker** window in the NMM SMSS plugin GUI.

Figure 14 The NetWorker window in the NMM SMSS plugin GUI**Backup tab**

In the **NetWorker** window, click the **Backup** tab to see the following pages which must be used to perform backup operations:

- General
- Options
- Monitor

Database Restore tab

In the **NetWorker** window, click the **Database Restore** tab to see the following pages which must be used to perform recovery operations:

- General
- Files/Filegroups
- Options
- Monitor

Table Restore tab

If you installed the SQL Granular Recovery functionality, from the **NetWorker Backup** window select the **Table Restore** tab to see the following pages which must be used to perform table-level recovery:

- General

- Options
- Monitor

CHAPTER 4

Manual Backups

This chapter includes the following sections:

- [Manual backup overview](#) 60
- [Federated backup preferences for Availability Group databases](#) 60
- [Specifying a retention policy for manual backups](#) 62
- [Performing manual backups from the NMM SSMS plugin GUI](#) 63
- [Performing manual backups from the command prompt](#) 68
- [Performing manual backups from the NetWorker User for SQL Server GUI](#) 80

Manual backup overview

You can start a manual backup of SQL data at any time. A manual (unscheduled) backup can be started immediately.

Manual backups are generally performed under special circumstances, such as NMM setup.

Before performing a manual backup, ensure that the user starting the backup is assigned the required user privileges as defined in the section [Access privileges for backup and recovery](#) on page 44.

For manual backups of SQL Server cluster environments, the backup operation must be run from the active host in the cluster.

The following combinations of data objects can be backed up manually with NMM:

- The entire SQL Server storage hierarchy
- One or more entire databases
- One or more file groups
- One or more files in a file group
- A heterogeneous collection of files, file groups, and databases
- Transaction log backups

Note

Filestream data, which are stored in SQL Server 2008 or later databases, is displayed in the backup window as a single file group folder without subordinate objects.

The storage hierarchy is defined as the database storage components the SQL Server Storage Engine exposes to third-party backup vendors. The storage components include files, file groups, databases, and transaction logs.

When performing a manual full backup of a file or file group, also perform a database logs-only level backup to maintain the validity of the transaction log. You can perform logs-only backups only through the NMM Microsoft SQL Server Management Studio plugin GUI, in the **NetWorker** window.

Note

Manual backups of SQL data do not back up the client indexes and bootstrap file. If you are using NetWorker 8.2.3 or 8.2.4 server with NMM 18.1, follow the instructions provided in the "Backing up client indexes and a bootstrap file" section of the *NetWorker Module for Microsoft for SQL VDI Release User Guide* version 8.2 SP1 to back up the client indexes and bootstrap file. NetWorker server 9.x and later have a built-in protection policy to back up the client indexes and bootstrap file.

Federated backup preferences for Availability Group databases

The SQL Server backup preference for the Availability Group can be configured through the NMM SSMS plugin GUI or the Transact-SQL command.

NMM supports federated backups, during which NMM detects the SQL Server preferred backup setting for the Availability Group and performs the backup at the

preferred node. The database administrator can set the backup priority for the Availability Group or a database in the Availability Group and nominate a particular replica for the backup. This feature improves backup performance.

SQL Server 2012 or later provides the following options that determine where backups run for a specific AG replica:

- **AUTOMATED_BACKUP_PREFERENCE**—Specify any one of the following options:
 - **PRIMARY**—Specifies that the backups will always occur on the primary replica. This option is useful if you need backup features, such as creating differential backups, that are not supported when backup operations run on a secondary replica.
 - **SECONDARY_ONLY**—Specifies that backups will never be performed on the primary replica. If the primary replica is the only replica that is online, the backup will not occur.
 - **SECONDARY (Prefer Secondary)**—Specifies that backups should occur on the secondary replica. If the primary replica is the only replica online, the backup is performed using the primary replica. This is the default option.
 - **NONE (Any replica)**—Specifies that you prefer that backup jobs ignore the role of the availability replicas when choosing the replica on which to perform backups. Note backup jobs might evaluate other factors such as backup priority of each availability replica in combination with its operational state and connected state.
- **BACKUP_PRIORITY =n**—Specifies your priority for performing backups on this replica relative to the other replicas in the same availability group. The value is an integer in the range of 0 - 100. These values have the following meanings:
 - 0 indicates that this availability replica will never be chosen for performing backups.
This choice is useful for a remote availability replica for which you never want backups to fail over.
 - 1 indicates the lowest priority that an availability replica could be chosen for a backup operation.
If **BACKUP_PRIORITY = 1**, the availability replica will be chosen only if no higher priority availability replicas are available.
 - 100 indicates the highest priority that an availability replica could be chosen for a backup operation.

Follow the procedures provided in the article “Configure Backup on Availability Replicas (SQL Server)” on the Microsoft MSDN website <http://msdn.microsoft.com/en-us/library/hh710053.aspx> to decide which replica the backups will be run on.

Note

Federated backup operations fail if a replica has the **Readable Secondary** field set to **No**. For backups to succeed, in the **Availability Group Properties** window, set the **Readable Secondary** configuration option to either **Yes** or **Read-intent only** for each of the primary and secondary replicas.

Note

The following limitations apply to SQL Server backups of a secondary SQL Server replica in an Availability Group:

- For full backups of a secondary replica, SQL Server supports only copy-only backups.
 - Backup promotion is not supported for copy-only backups, including copy-only transaction log backups.
 - SQL Server does not support differential backups on secondary SQL Server replicas.
-

Specifying a retention policy for manual backups

A retention policy defines the length of time that a backup or clone save set is retained and browsable.

Note

NetWorker 8.2.x uses separate policies for browse and retention. Refer to the *NetWorker Module for Microsoft for SQL VDI User Guide* version 8.2 SP1 for information on the browse and retention policies for NetWorker 8.2.x.

If you specify a retention policy with a manual backup from the command prompt, the retention policy takes effect for all the save sets in the manual backup. To specify a retention policy with a manual backup at the command prompt, type the following command:

```
nsrsqlsv -y
```

The retention policy value must be typed in time and date formats accepted by the `nsr_getdate` program.

The *NetWorker Command Reference Guide* or the UNIX man pages provide more information about the `nsr_getdate` program. Refer to the "Command Prompt" chapter for information about using the `nsrsqlsv` command for manual backups.

If you do not specify a retention policy for a manual backup, the save sets in a manual backup adopt the browse policy of the client resource. If multiple client resources exist for the NetWorker host, the client resource with the longest retention policy is adopted. However, if a retention policy is set up for the media pool to which the backup is directed, the retention policy is the longer of either:

- The client resource retention policy
- The media pool retention policy

The *NetWorker Administration Guide* provides more information about retention policies.

Performing manual backups from the NMM SSMS plugin GUI

It is recommended that you use the NMM SQL Server Management Studio (SSMS) plugin GUI to perform SQL Server VDI manual backup operations.

In the **NetWorker** window, the **Script** view is available in each page. Click **Script** to generate a command prompt equivalent script, which you can use to create a .bat file to perform scheduled backups, automation, and other tasks.

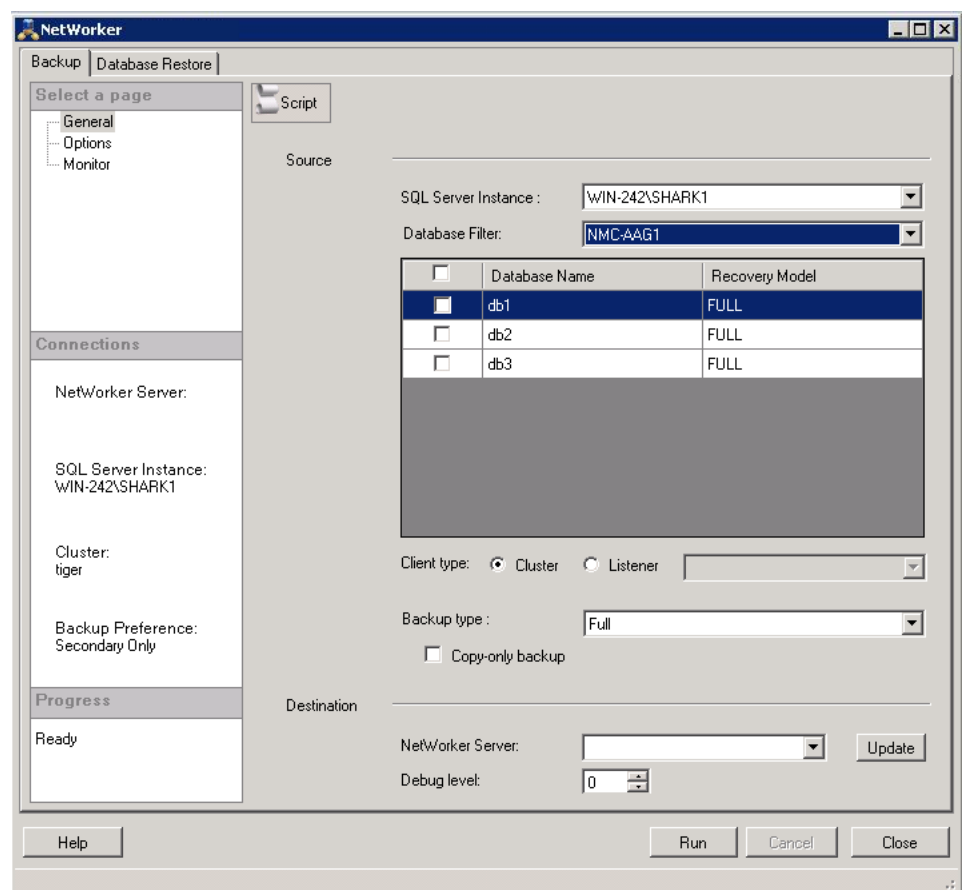
Procedure

1. In the SQL Server Management Studio, select the SQL Server instance that you want to manually back up and open the **NetWorker** window.
2. In the **NetWorker** window, on the **Backup** tab, under **Select a page**, click **General**.

The **General** page appears.

The following figure shows the **NetWorker** window open to the **Backup** tab, on the **General** page.

Figure 15 NetWorker Backup General page



3. Under **Source**, make the following selections:
 - a. In the **SQL Server Instance** field, select the SQL Server instance that you want to back up.

- b. In the **Database Filter** list, filter the databases by selecting one of the following:
- Select **Non AAG Databases** to show all databases that are not part of an Always On Availability Group (AG).
 - Select the name of an AG to show all databases that are within that AG.

The list is populated with databases available for backup. If an AG is selected from the list, under **Connections**, the **Cluster** and the **Backup Preference** appear.

- c. Select one or more databases that you want to back up.
- d. Under **Client type**, perform one of the following actions for availability groups:
- To use the cluster name, select **Cluster**.
 - To use an availability group listener, select **Listener**, and then select a listener from the list.

Note

The **Client type** setting can be modified only for availability group backups. In non-availability group environments, it is greyed out.

- e. In the **Backup type** list, select the backup type.

Note

The backup levels are the same regardless of the version of the NetWorker server you are using.

- f. (Optional) To perform a copy-only backup, select **Copy-only backup**.

A copy-only backup backs up the SQL server without interrupting the existing backup schedule or log chain.

4. Under **Destination**, make the following selections:

- a. In the **NetWorker Server** list, select the destination NetWorker Server for the backup.

You can select a NetWorker server from the available list. Click **Update** to force an update of the list of available NetWorker servers.

- b. In the **Debug level** box, select the level of troubleshooting information to be sent to the backup status window during the backup operation.

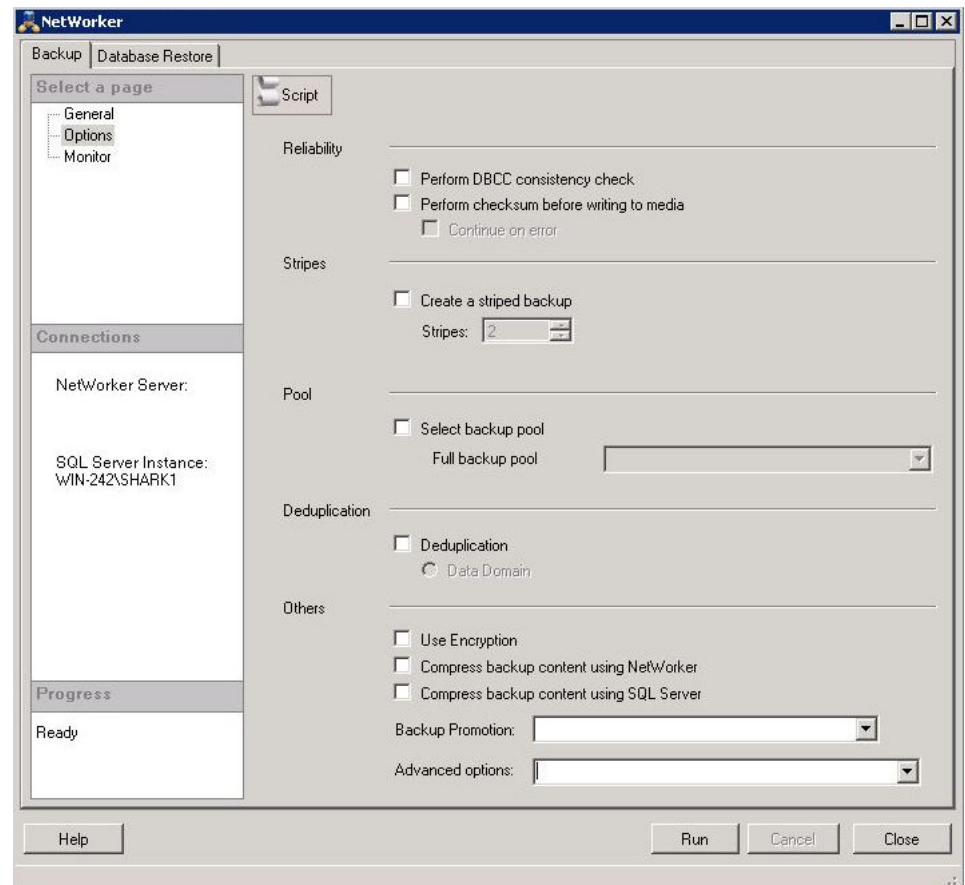
Levels range 0–9, with 0 representing no information and 9 representing the most amount of information.

5. (Optional) To specify additional backup options, under **Select a page**, click **Options**.

The **Options** page appears.

The following figure shows the **NetWorker** window open to the **Backup** tab, on the **Options** page.

Figure 16 NetWorker Backup Options page



6. (Optional) Under **Reliability**, select the following options as required:

- To perform a thorough check of the entire database before the backup is performed, select **Perform DBCC consistency check**.

Note

When **Perform DBCC consistency check** is selected in the SSMS plugin, the script generated contains `-j ALL` even when specific databases are selected. This does not affect functionality.

- To perform a checksum operation with the backup and save the information to the backup media, select **Perform checksum before writing to media**. Another checksum is performed before a restore to ensure that it matches the backup.
A checksum detects a partial backup or restore state. The SQL Server verifies the checksum by calculating a local result and comparing the local value with the stored value.
- To continue with backup and recovery even if errors are discovered during the checksum, select **Continue on error**.

7. (Optional) Under **Stripes**, select the following options as required:

- To enable the SQL Striped feature, select **Create a striped backup**.
- In the **Stripes** box, type or select the number of stripes to use during the backup.

NMM supports a maximum of 32 stripes. However, the number of stripes cannot exceed the value set for NetWorker client parallelism.

The section [Striped Backup and Recovery](#) provides more information about striping.

8. (Optional) Under **Pool**, select the following options as required:
 - a. To select the pool where the backup is stored, select **Select backup pool**.
 - b. From the **Full backup pool** list, select the pool that you want to back up to.
The pools that are listed are populated from the NetWorker server.
9. (Optional) Under **Deduplication**, select the following options as required:
 - a. To enable deduplication for the backup, select **Deduplication**.
 - b. To use a Data Domain device to store the backup, select **Data Domain**.
10. (Optional) Under **Others**, select the following options as required:

Note

It is recommended that no form of encryption or compression be used with Data Domain because these reduce the deduplication ratio.

- To back up the data with AES encryption, select **Use Encryption**.
Data is encrypted with the default or current pass phrase which is provided by the NetWorker Server. If the NetWorker Server has a different pass phrase at recovery time, you must specify the pass phrase that was used at the time of backup. The *NetWorker Administration Guide* provides complete information about AES encryption, and setting the pass phrase.
- To apply XBSA compression to all marked databases before the backup data is written to the storage device, select **Compress the backup content using NetWorker**.
In the same manual backup, you cannot back up some databases with compression and others without.

Compressing data for a backup generates less network traffic and uses less backup media space, but it consumes additional CPU resources. Most tape devices perform compression, which makes software compression unnecessary.
- To compress the backup using SQL Server, select **Compress the backup content using SQL Server**.
The Microsoft SQL Server product documentation provides more information.
- To specify backup promotion options, select one of the following values from the **Backup Promotion** list:
 - **ALL**: Enables backup promotion to occur in any applicable scenario.
 - **NONE**: Disables backup promotion. Logs a warning when backup promotion would normally occur.
 - **NONE_WITH_WARNINGS**: Disables backup promotion. Logs a warning when backup promotion would normally occur.
 - **SKIP_RECOVERY_MODEL**: Disables database recovery model change detection. Backup promotion as a result of recovery model change will not occur, but backup promotion in other scenarios will still occur.

Note

If you leave the **Backup Promotion** field on the client side blank and the client resource for the server has backup promotion set to **NONE**, then backup promotion will not occur. However, if you select an option from the **Backup Promotion** list on the client side, the option will override the server setting.

- To apply advanced options to the backup operation, in the **Advanced options** field, type or select advanced options. Separate multiple entries with a comma, for example:
`BUFFERCOUNT=2, READ_WRITE_FILEGROUPS`
 NMM supports the following advanced backup options:
 - `BUFFERCOUNT=number_of_IO_buffers`: Specifies the total number of IO buffers that can be used during the backup operation.
 - `READ_WRITE_FILEGROUPS`: Backs up only the read/write (active) filegroups within the database.

11. To start the backup, click **Run**.

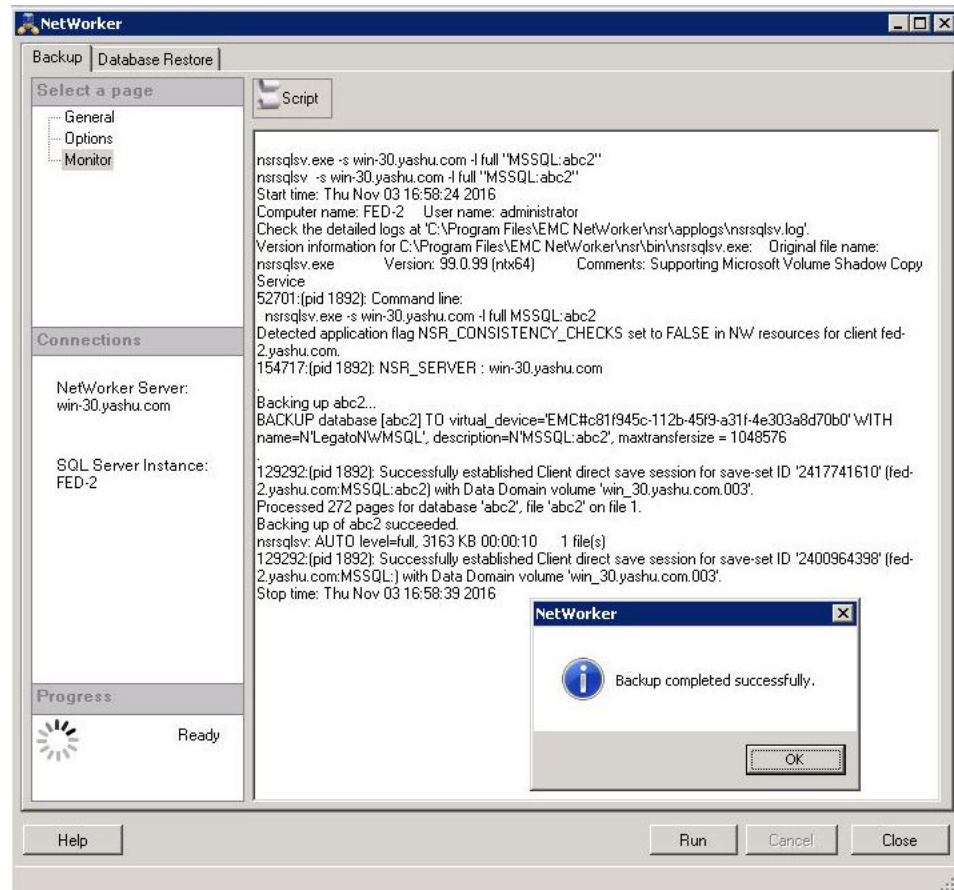
Results

The backup operation runs.

To view the status of the backup operation, open the **Monitor** page.

The following figure shows the **NetWorker** window open to the **Backup** tab, on the **Monitor** page.

Figure 17 NetWorker Backup Monitor page



Performing manual backups from the command prompt

The `nsrsqlsv` command is used to back up SQL Server data objects, which consist of files, file groups, and databases.

To start a backup, specify the `nsrsqlsv` command and its options from a Windows command prompt.

Before you use the backup command, review the following considerations:

- Ensure that each command option is either lowercase or uppercase and, frequently both the cases of a letter are in the set of command options. Case is very important when specifying command-line flags.
For example, `-c` specifies the NetWorker client name, while `-C` copies the database that is being restored to either the same SQL Server or a different SQL Server.
- Ensure that you use correct spacing. Depending on the command option, the space separator between an option and its corresponding argument can be optional, required, or not allowed.
For example, both of the following expressions are valid:

```
-l backup_level
```

```
-lbackup_level
```

The following expression is invalid because a space is not allowed between the + argument and *log_mark_name*:

```
-M + log_mark_name
```

- Ensure that you use brackets ([]) to denote the optional portions of a command. When starting an actual backup or recovery operation, do not include brackets.
- Ensure that data items follow all other command options and parameters at the command prompt.

Command syntax for nsrsqlsv

The nsrsqlsv command syntax is:

```
nsrsqlsv [-CGjqRTvkuHXZ]
[-a "option_key=value"]
[-A virtual-server]
[-b pool]
[-c client]
[-C]
[-D debug_level]
[-f aes]
[-g group]
[-G]
[-h "database_name"]
[-H]
[-I text_file]
[-k]
[-j database_names]
[-l backup_level]
[-m masquerade]
[-N name]
[-O "Option1, Option2, Option3"]
[-P password]
[-q]
[-R]
[-s networker_server]
[-S number_of_stripes]
[-T]
[-u]
[-U username]
[-v]
[-w browse]
[-X]
```

```
[-y retention_period]
[-z federated backup options]
[-Z]
```

Command options for nsrsqlsv

The following table outlines the list of options available for the `nsrsqlsv` command.

Table 17 Command options for `nsrsqlsv`

Command options	Descriptions
<code>-a "NSR_INCLUDE_AAG_DATABASE=None"</code>	Excludes Always On Availability Group databases from instance-level backups.
<code>-a "SKIP_SIMPLE_DATABASE={TRUE FALSE}"</code>	<p>Specifies whether to exclude simple recovery model databases from a logs only backup. The default value is FALSE.</p> <p>This option is useful when SIMPLE recovery model databases are part of a logs only backup. SIMPLE recovery model databases do not support logs only backups and the backups are promoted to level full by default.</p> <p>However, this option does not affect the backup of system databases, such as master and msdb, which are promoted to level full during a logs only backup and cannot be omitted from the backup.</p> <p>This option does not apply to full and cumulative incremental backups.</p>
<code>-a "NSR_SKIP_NON_BACKUPABLE_STATE_DB={TRUE FALSE}"</code>	<p>Specifies whether to check the status of the selected databases and ignore the databases that are not ready or available to be backed up. If the status of a database is online, the database is ready and available to be backed up. If the status of a database is offline, emergency (single user mode), suspect, restoring, recovering, or recovery pending, the database is not ready or available to be backed up.</p> <p>The default value is FALSE.</p>
<code>-a "NSR_COPY_ONLY={TRUE FALSE}"</code>	<p>Performs a copy-only backup from the SQL Server. Copy-only backups do not disturb the SQL Server backup chain and do not affect backup-level promotion and log truncation. This option works only if you specify either full or transaction log as the backup type in the <code>-l <Backup_Type></code> option.</p> <p>The default value is FALSE.</p> <p>You rarely need to perform copy-only transaction log backups because they are generally required only to perform online restores.</p>
<code>-a "SKIP_SYSTEM_DATABASE={TRUE FALSE}"</code>	<p>Specifies whether to skip the system databases during the instance-level transaction log backups.</p> <p>The default value is False.</p>
<code>-a "BACKUP_PROMOTION={ALL NONE NONE_WITH_WARNINGS SKIP_RECOVERY_MODEL}"</code>	<p>Specifies backup promotion options. The following values are valid:</p> <ul style="list-style-type: none"> ALL (Default): Enables backup promotion to occur in any applicable scenario. NONE: Disables backup promotion. Logs a warning when backup promotion would normally occur. NONE_WITH_WARNINGS: Disables backup promotion. Logs a warning when backup promotion would normally occur.

Table 17 Command options for `nsrsqlsv` (continued)

Command options	Descriptions
	<ul style="list-style-type: none"> SKIP_RECOVERY_MODEL: Disables database recovery model change detection. Backup promotion as a result of recovery model change will not occur, but backup promotion in other scenarios will still occur. <p>The setting specified with this option applies to every database in the SQL instance when this setting is specified along with the SQL instance level save set <code>MSSQL:</code> or <code>MSSQL\$<instance name></code>.</p> <p>Similarly, the setting specified with this option applies only to specified databases when this setting is specified with a database level save set <code>MSSQL: <database name></code> or <code>MSSQL\$<instance name>: <database name></code>.</p>
<code>-a "Deduplication backup=yes"</code>	Enables backup deduplication.
<code>-a "Device interface={data domain avamar}"</code>	Specifies whether data deduplication should be set with Data Domain or Avamar.
<code>-a "Deduplication node=node_name"</code>	Specifies the deduplication node to use for the backup.
<code>-a "Deduplication client name=client"</code>	Specifies the deduplication client to use for the backup.
<code>-A virtual_server</code>	Specifies the virtual server name when a SQL Server is configured to run in an MSCS cluster.
<code>-b pool_name</code>	Assigns a destination pool for a save set. Specifying <code>-b pool_name</code> overrides all other pool-selection criteria either that is either provided by the NetWorker client or specified in the Backup Options dialog box of the NetWorker User for SQL Server GUI. The pool must be created with a corresponding label template before a command that includes the <code>-b</code> option runs.
<code>-c client</code>	Specifies the NetWorker client name for which the SQL Server data is to be backed up.
<code>-C</code>	Specifies compression of the backup data before the NetWorker client sends the data to the NetWorker server.
<code>-D <debug_level></code>	Generates detailed logs that you can use to troubleshoot backup issues. The default value is 0 (zero). If you want to enable debug logs, specify a level between 1 and 9.
<code>-f aes</code>	Enables the NetWorker server to back up data using AES encryption.
<code>-g group</code>	Specifies the savegroup. The NetWorker server and the <code>savegrp</code> command use the group parameter to select the media pool.
<code>-G</code>	Specifies a NO_LOG transaction log backup before backing up the database. This command option is valid only for level full backups. This option is deprecated in SQL Server 2008 or later.
<code>-h <"db_name_with_optional_wildcard"></code>	<p>Excludes databases from the backup. You can use this option to specify exact database names or use wildcard characters. Two wildcard characters are supported:</p> <ul style="list-style-type: none"> Question mark (?): Matches any single character Asterisk (*): Matches zero to unlimited characters

Table 17 Command options for `nsrsqlsv` (continued)

Command options	Descriptions
	<p>When you use wildcards you must enclose the database name in square brackets, for example, <code>["DB?"]</code>.</p> <p>For example, consider the following scenarios:</p> <ul style="list-style-type: none"> To exclude only DB_1 and DB_2 from the backup, type the following command: <code>nsrsqlsv -s bv-customer.belred.emc.com -h "DB_1" -h "DB_2" MSSQL:</code> To exclude all databases that are named with the format of DB_x, such as DB_9 and DB_a, type the following command: <code>nsrsqlsv -s bv-customer.belred.emc.com -h ["DB_?"] MSSQL:</code> To exclude all databases with names ending in DB, type the following command: <code>nsrsqlsv -s bv-customer.belred.emc.com -h ["*DB"] MSSQL:</code>
<code>-H</code>	Specifies the NORECOVERY option when backing up transaction logs. This option leaves the database in Restoring state.
<code>-I text_file</code>	Specifies a text file that lists multiple SQL Server save sets. The <code>-I</code> option can also be specified with the <code>nsrsqlsv</code> command for the Backup Command field in the NetWorker Administration GUI.
<code>-k</code>	Performs checksum before writing to media.
<code>-j database_name</code>	<p>Performs a database consistency check (DBCC) before the backup is performed. For example, <code>-j testdb3,CopyOfDB010,test4</code> runs a consistency check for databases testdb3, CopyOfDB010, and test4.</p> <hr/> <p>Note</p> <p>When Perform DBCC consistency check is selected in the SSMS plugin, the script generated contains <code>-j ALL</code> even when specific databases are selected. This does not affect functionality.</p> <hr/>
<code>-l backup_level</code>	<p>Specifies the backup level. The following values are valid:</p> <ul style="list-style-type: none"> full diff txnlog <p>The "Combining data objects to create backup levels" section provides more information about which backup levels are supported for various SQL Server data objects.</p> <p>The <code>-l</code> option is valid only for manual backups that are initiated from a Windows command prompt on a client host. For scheduled backups, set the backup level in the Set Level dialog box of the schedule resource in the NetWorker Administration GUI. Do not use the <code>-l</code> option when you start a backup in the NetWorker Administration GUI.</p> <p>Sample outputs for each of the three different levels of backups follow:</p> <ul style="list-style-type: none"> Full backup of the database to a specified NetWorker server: <code>nsrsqlsv -s server -c client -b poolname -l full MSSQL:dbname</code>

Table 17 Command options for `nsrsqlsv` (continued)

Command options	Descriptions
	<p>nsrsqlsv -s swraj -c SQL2012RC1Named -b Sub9VDI -l full MSSQL:5</p> <pre> 43708:(pid 6004):Start time: Sat Jan 28 09:07:42 2012 43621:(pid 6004):Computer Name: SQL2012-NODE3 User Name: Administrator NSR_BACKUP_LEVEL: full; NSR_CLIENT: SQL2012RC1Named.joy.com; NSR_DATA_VOLUME_POOL: Sub9VDI; NSR_LOG_VOLUME_POOL: Sub9VDI; NSR_SAVESET_NAME: "MSSQL:5"; NSR_SERVER: Swraj.joy.com; 37994:(pid 6004):Backing up 5... 4690:(pid 6004):BACKUP database [5] TO virtual_device='EMC#4018d580-f511-4457-abc3-a62c4c3f0ff9' WITH name=N'EMCNWMSQL' 53085:(pid 6004):Backing up of 5 succeeded. nsrsqlsv: MSSQL: 5 level=full, 2261 KB 00:00:02 1 file(s) 43709:(pid 6004):Stop time: Sat Jan 28 09:07:49 2012 </pre> <ul style="list-style-type: none"> Cumulative incremental backup: nsrsqlsv -s NetWorker_server_name -l diff MSSQL:dbname Output similar to the following appears for a backup of a database named <code>my_database</code>: <pre> nsrsqlsv: Backing up my_database... nsrsqlsv: BACKUP database my_database TO virtual_device='BSMSQL' WITH differential, stats nsrsqlsv: my_database level=diff, 719 KB 00:00:05 1 file(s) </pre> Logs-only backup: nsrsqlsv -s NetWorker_server_name -l txnlog MSSQL:dbname Output similar to the following appears for a backup of a database named <code>my_database</code>: <pre> nsrsqlsv: Backing up my_database... nsrsqlsv: BACKUP log my_database TO virtual_device='BSMSQL' nsrsqlsv: my_database level=txnlog, 61 MB 00:00:05 1 file(s) </pre> <p>At least one SQL Server data item (file, file group, or database) must be specified, and the data items and list of data objects must follow all other command options and parameters at the command prompt.</p>
-m name	Specifies the cluster or availability group listener name.
-N name	<p>Specifies the symbolic name of the save set. By default, the most common prefix of the path arguments is used as the save set name.</p> <p>When performing a federated backup of an availability group, you must specify the save set name with -N.</p>
-O "option1, option2"	<p>Specifies advanced backup options.</p> <p>You can specify the following advanced backup options:</p> <ul style="list-style-type: none"> BUFFERCOUNT=number_of_IO_buffers This option specifies the total number of IO buffers that can be used during a backup operation. READ_WRITE_FILEGROUPS This option backs up only the read/write (active) filegroups within the database.

Table 17 Command options for `nsrsqlsv` (continued)

Command options	Descriptions
	<p>If you are specifying multiple options, separate each argument with a comma. For example:</p> <pre>nsrsqlsv -s NetWorker_server_name -O "BUFFERCOUNT=number_of_IO_buffers, READ_WRITE_FILEGROUPS" MSSQL:</pre>
<code>-P password</code>	<p>Specifies the SQL Server user password. When the <code>-U</code> user name command option is specified, the password command option must also be provided, as follows:</p> <pre>nsrsqlsv -s NetWorker_server_name -U username -P password MSSQL:</pre> <p>Use the SQL Server username and password to log in to SQL Server with SQL server authentication.</p>
<code>-q</code>	Displays <code>nsrsqlsv</code> messages in quiet mode; only summary information and error messages appear.
<code>-R</code>	Specifies the <code>NO_TRUNCATE</code> option when backing up transaction logs.
<code>-s NetWorker_server_name</code>	Specifies the name of the NetWorker server that will be used for the backup operation.
<code>-S number_of_stripes</code>	<p>Backs up the specified data items using stripes.</p> <p>To use backup and recovery striping successfully, see the striping configuration described in Striped Backup and Recovery.</p>
<code>-T</code>	<p>Performs a <code>TRUNCATE_ONLY</code> transaction log backup before backing up the database and is valid for only full backups.</p> <hr/> <p>Note</p> <p>This flag is deprecated in SQL Server 2008 or later.</p> <hr/>
<code>-u</code>	Continues the backup even in the event of a checksum error.
<code>-U username</code>	<p>Specifies the SQL Server username. When this command option is specified, the <code>-P</code> password command option must also be provided, as follows:</p> <pre>nsrsqlsv -s NetWorker_server_name -U username -P password MSSQL:</pre> <p>Use the SQL Server username and password to log in to SQL Server with SQL server authentication.</p>
<code>-v</code>	Displays the <code>nsrsqlsv</code> command messages in verbose mode, providing detailed information about the progress of the backup operation.
<code>-X</code>	Indicates that SQL Server internal backup compression is used.
<code>-y retention_period</code>	Specifies the retention period of the backup. The value must be typed in time and date formats accepted by the <code>nsr_getdate</code> program. The <i>NetWorker Command Reference Guide</i> or the UNIX man pages provide more information about the <code>nsr_getdate</code> program.
<code>-Z</code>	Applies to the backup of up databases for a SQL Server 2005, and is usually used in the online recovery scenario from the command prompt. The <code>-Z</code> option specifies

Table 17 Command options for `nsrsqlsv` (continued)

Command options	Descriptions
	that the logs-only (transaction log) backup after recovery is not promoted to full backup. Without the <code>-z</code> option, the backup is promoted to full.
<code>-z federated_backup_option</code>	<p>Specifies federated backup options. Use the following values with the <code>-z</code> option:</p> <ul style="list-style-type: none"> • <code>-z FEDERATED_SLAVE=true</code> This option denotes that the process is SQL Federated secondary process. • <code>-z FEDINDEX_NAME =SQL instance name under which the backups are browsable</code> For example, <code>MSSQL</code> or <code>MSSQL\$InstanceName</code> • <code>-z FEDCLIENT_NAME =Windows cluster client name</code> This option is useful for cluster federated backups. <p>When you use the <code>-z</code> option, you must also use the <code>-N</code> option to specify the save set name.</p> <p>For example, to perform an incremental federated backup, type the following command:</p> <pre>nsrsqlsv.exe -s mars.jupiter.com -m WIN-FED-3 -a device interface=device interface -a device interface data domain -a Data Domain interface=IP -b Default -l incr -q -z FEDERATED_SLAVE=true -z FEDINDEX_NAME=MSSQL -z FEDCLIENT_NAME=cluster_fed -N MSSQL#MARS:MARSDB MSSQL:MARSDB</pre>

Backup and recovery command syntax for SQL Server data

Use the additional command syntax that is shown in the following table to back up or restore SQL Server data with the `nsrsqlsv` and `nsrsqlrc` backup and recovery commands.

You can specify more than one data object and can combine different types of data. You must specify the SQL data objects with the syntax shown in the following table.

Table 18 Command syntax for SQL Server data

SQL Server data	Syntax for SQL Server data objects
All databases in the SQL Server storage hierarchy (optional)	<p>MSSQL:</p> <p>Typing MSSQL: yields an instance-level backup of all databases on the SQL Server host.</p>
Specified databases	<p>MSSQL: dbName</p> <p>or</p> <p>[MSSQL: dbName MSSQL: dbName2 ...]</p>
All file groups in specified databases	<p>MSSQL: dbName</p> <p>or</p>

Table 18 Command syntax for SQL Server data (continued)

SQL Server data	Syntax for SQL Server data objects
	<code>[MSSQL: <i>dbName</i>. MSSQL: <i>dbName2</i> ...]</code>
Specified file groups in specified database	<code>MSSQL: <i>dbName</i>. <i>fgName</i></code> or <code>[MSSQL: <i>dbName</i>. <i>fgName</i> MSSQL: <i>dbName</i>. <i>fgName2</i></code> <code>MSSQL: <i>dbName2</i>. <i>fgName</i> MSSQL: <i>dbName2</i>. <i>fgName2</i> ...]</code>
Specified files in specified file groups in specified databases	<code>MSSQL: <i>dbName</i>. <i>fgName</i>. <i>filename</i></code> or <code>[MSSQL: <i>dbName</i>. <i>fgName</i>. <i>filename</i> MSSQL: <i>dbName</i>. <i>fgName2</i>. <i>filename</i></code> <code>MSSQL: <i>dbName2</i>. <i>fgName</i>. <i>filename</i></code> <code>MSSQL: <i>dbName2</i>. <i>fgName2</i>. <i>filename</i> ...]</code>

Specifying `MSSQL` before each data object name is optional and does not affect the expression or the resulting operation. However, when `MSSQL` is specified, it must be followed by a colon (:).

For example, the following two commands are equivalent:

```
nsrsqlsv -s NetWorker_server_name dbName.fgName
nsrsqlsv -s NetWorker_server_name MSSQL:dbName.fgName
```

In a non-clustered, named instance configuration, `MSSQL$` is required, followed by the instance name and a colon. For example:

```
nsrsqlsv -s NetWorker_server_name MSSQL$instanceName:dbName.fgName
```

Syntax for a named instance configuration

When a configuration contains non-clustered named instances of the SQL Server, you must specify the name of the instance before the data.

```
MSSQL$instanceName:[dbName ...] [.fgName ...] [.fileName ...]
```

Note

The `nsrsqlsv` and `nsrsqlrc` commands support specification of only a single instance. If save sets for more than one instance are specified, the backup fails. The `nsrsqlrc` command supports mixing instances for a copy restore operation.

Example 3 Back up all databases in an instance

To back up all the databases for instanceOne, type the following:

```
nsrsqlsv -s NetWorker_server_name MSSQL$instanceOne:
```

Example 4 Recover several file groups in an instance

To recover several file groups for instanceTwo, type the following:

Example 4 Recover several file groups in an instance (continued)

```
nsrsqlrc -s NetWorker_server_name MSSQL$instanceTwo:dbName.fgName
MSSQL$instanceTwo:dbName.fgName2
```

Example 5 Using clustered instance SQL Server virtual server names with -A or -c

Instead of using clustered named instances in this syntax, use clustered instance SQL Server virtual server names with -A or -c options, as shown in the following command:

```
nsrsqlsv -s NetWorker_server_name -A SQL_virtual_server_DNS_name -c
SQL_virtual_server_DNS_name MSSQL:
```

where:

- *NetWorker_server_name* is the hostname of the NetWorker server.
- *SQL_virtual_server_DNS_name* is the Domain Name System (DNS) name for the SQL Server virtual server.

Create a client resource under the SQL Server virtual server DNS name.

For scheduled backups of a SQL Server virtual server client, you do not need to specify the -A or -c option with the SQL Server virtual server name. The `savegrp` process automatically specifies the virtual server name to the `nsrsqlsv` process by using the -m option.

Syntax for names containing a period, backslash, or colon

NMM provides command line syntax that enables you to back up and recover file names, file groups, and databases containing a period (.), backslash (\), or colon (:). By typing a backslash before the period or backslash, the `nsrsqlsv` and `nsrsqlrc` commands interpret the period or backslash as a literal character.

The tables in this section list the syntax for file names, file groups, and databases containing a period, backslash, colon, or any combination of the three.

The following notes apply to the information in the tables:

- The syntax that is shown in the right column applies to both the `nsrsqlsv` and `nsrsqlrc` commands.
- The notation `MSSQL:` is optional only for the `nsrsqlsv` command.
- A single period (.) continues to delimit SQL identifiers.
- The syntax also applies to named instances.
- The backslash period (\.) character sequence replaces each literal period in the SQL identifier.
- The double backslash (\\) character sequence replaces each literal backslash in the SQL identifier.

The following table lists the syntax for file names, file groups, and databases containing a period.

Table 19 Command syntax for names containing a period

Name visible from SQL utilities	Equivalent command-line syntax
SQL database named MyDatabase.COM	MSSQL:MyDatabase\ .COM MSSQL\$MyInstance:MyDatabase\ .COM
SQL file group named MyFileGroup.2 for the SQL database named MyDatabase.COM	MyDatabase\ .COM.MyFileGroup\ .2 MSSQL:MyDatabase\ .COM.MyFileGroup\ .2 MSSQL\$MyInstance:MyDatabase\ .COM.MyFileGroup\ .2
SQL file named MyFile.2, which is a member of the SQL file group named MyFileGroup.2 for the SQL database named MyDatabase.COM	MyDatabase\ .COM.MyFileGroup\ .2.MyFile\ .2 MSSQL:MyDatabase\ .COM.MyFileGroup\ .2.MyFile\ .2 MSSQL\$MyInstance:MyDatabase\ \COM.MyFileGroup\ \2.MyFile\ \2

The following table lists the syntax for file names, file groups, and databases containing a backslash.

Table 20 Command syntax for names containing a backslash

Name visible from SQL utilities	Equivalent command-line syntax
The SQL database named MyDatabase\COM	MyDatabase\ \COM MSSQL:MyDatabase\ \COM MSSQL\$MyInstance:MyDatabase\ \COM
The SQL file group named MyFileGroup\2 for the SQL database named MyDatabase\COM	MyDatabase\ \COM.MyFileGroup\ \2 MSSQL\$MyInstance:MyDatabase\ \COM.MyFileGroup\ \2

The following table lists the syntax for file names, file groups, and databases containing a colon.

Table 21 Command syntax for names containing a colon

Name visible from SQL utilities	Equivalent command-line syntax
SQL database named MyDatabase:COM	MyDatabase:COM MSSQL:MyDatabase:COM MSSQL\$MyInstance:MyDatabase:COM
SQL file group named MyFileGroup:2 for the SQL database named MyDatabase:COM	MyDatabase:COM.MyFileGroup:2 MSSQL:MyDatabase:COM.MyFileGroup:2 MSSQL\$MyInstance:MyDatabase:COM.MyFileGroup:2
SQL file named MyFile, which is a member of the SQL file group named MyFileGroup:2 for the SQL database named MyDatabase:COM	MSSQL\$MyInstance:Mydatabase:com.MyFileGroup:2.MyFile

Example backup command syntax

Back up a WSFC cluster

The Windows account that you use to back up the cluster must have WSFC administrator privileges. To determine which accounts have WSFC administrator privileges, refer to the WSFC online help. If the Windows account does not have WSFC administrator privileges, NMM cannot communicate with WSFC and the various WSFC cluster resources, including the SQL Server virtual servers.

To back up a SQL virtual server database, type the following:

```
nsrsqlsv -A SQL_virtual_server_name -c SQL_virtual_server_name -s
NetWorker_server_name MSSQL:dbName
```

where:

- *SQL_virtual_server_name* is the virtual server name when the SQL Server is configured to run in a WSFC cluster.
- *NetWorker_server_name* is the NetWorker server that is designated for the backup.
- *dbName* is the name of the SQL Server database that is to be backed up.

Specifying *-A SQL_virtual_server_name* starts the following actions:

- Contacts the SQL virtual server.
- Creates save set entries under *SQL_virtual_server_name* in the NetWorker client index.

Perform a full backup of an Always On Availability Group

```
nsrsqlsv.exe -s nwsrvr -m automation -a device interface="data domain"
-b Default "MSSQL$SQ12INST4#sql2012-aag3:"
```

where:

- *nwsrvr* is the NetWorker server
- *automation* is the Windows cluster name
- *MSSQL\$SQ12INST4#sql2012-aag3:* is the save set name:
 - *MSSQL* is a mandatory term.
 - *sq12INST4* is the SQL Server instance name.
If you use a named instance, ensure that you perform the backup by using the same instance name, even if the SQL Server has failed over to another node in the Always On Availability Group.
 - *#* indicates a federated backup.
 - *sql2012-aag3* is the Always On Availability Group name.
- *MSSQL\$SQ12INST4#sql2012-aag3:* is the backup object name, where all the databases of the *sql2012-aag3* Always On Availability Group are backed up.

Back up a specific database in an Always On availability group

To restrict the backup to specific databases, specify the database names. For example, to back up *database1* only, type "*MSSQL\$SQ12INST4#sql2012-aag3:database1*" "*MSSQL\$SQ12INST4#sql2012-aag3:database1*"

Back up multiple instances on multiple nodes

When you have multiple instances on multiple nodes in an Always On availability group, select one instance name to use as the save set name and target for the Always On

availability group backup. Use that same instance name when naming the backup object, regardless of which node in the availability group the backup is started on.

For example, consider the following scenario:

- There are two instances, SQ12INST4 and SQ12INST5, are present on Node 1 and Node 2 respectively.
- You want to use SQ12INST4 for the backup.
- Node 1 is down, and as the result, the SQ12INST4 instance is down.
- You must perform the backup on Node 2 because Node 1 is down.

In this scenario, to use SQ12INST4, you must specify the save set and the backup object as "MSSQL\$SQ12INST4#sql2012-aag3:" "MSSQL\$SQ12INST4#sql2012-aag3:" instead of "MSSQL\$SQ12INST5#sql2012-aag3:" "MSSQL\$SQ12INST5#sql2012-aag3:".

Performing manual backups from the NetWorker User for SQL Server GUI

The **Backup** window of the NetWorker User for SQL Server GUI displays data that is available for backup based on the SQL Server database settings. The SQL database items that cannot be backed up are not displayed in the **Backup** window.

Before you begin

Before starting a backup, ensure that a labeled media volume is mounted in the backup device because otherwise a backup cannot run. The volume should be labeled for the volume pool where the backup is to be directed. If there is no volume in the backup device when a backup is started, no messages appear in the **Backup Status** window and the backup waits for operator intervention.

The SQL database items include, but are not limited to, databases in the following states:

- Standby
- Recovering
- Suspect
- Offline
- Not recovered
- Loading
- Pre-recovery
- Restoring
- Recovery pending
- Single user with active user connections

You must complete the required steps to perform a manual backup.

Procedure

1. From the **Start** menu, select **Programs > NetWorker > NetWorker User for SQL Server**.

If multiple instances of SQL Server are active on the computer, the **Select SQL Instance** dialog box appears.

2. Select the SQL Server instance for NMM backup, and then click **OK**.

The NetWorker User for SQL Server GUI connects to the selected SQL Server instance and the main window appears.

3. To select a NetWorker server other than the server that was specified during the NMM installation, from the toolbar, click **Select NetWorker Server**.

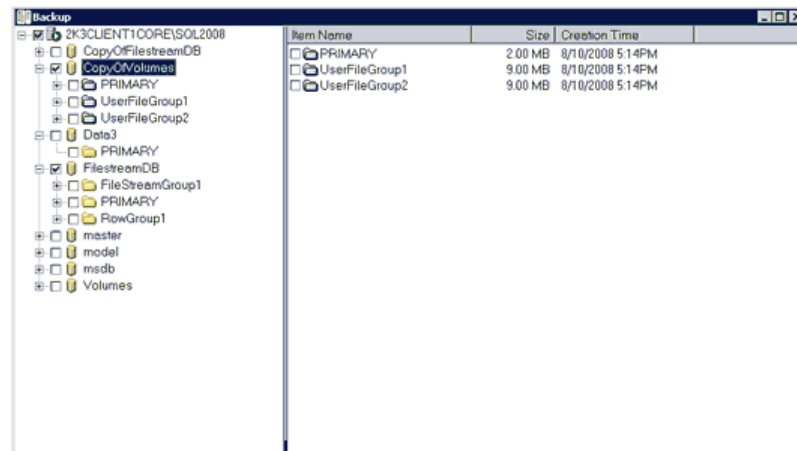
The **Change Server** dialog box appears.

4. Select a NetWorker server from the list, and then click **OK**.
5. To select the SQL Server data to back up, from the **Operation** menu, select **Backup**.

The **Backup** window displays a hierarchical list of SQL Server data objects that are available for backup. The NetWorker User for SQL Server GUI always performs a full backup for the SQL Server data.

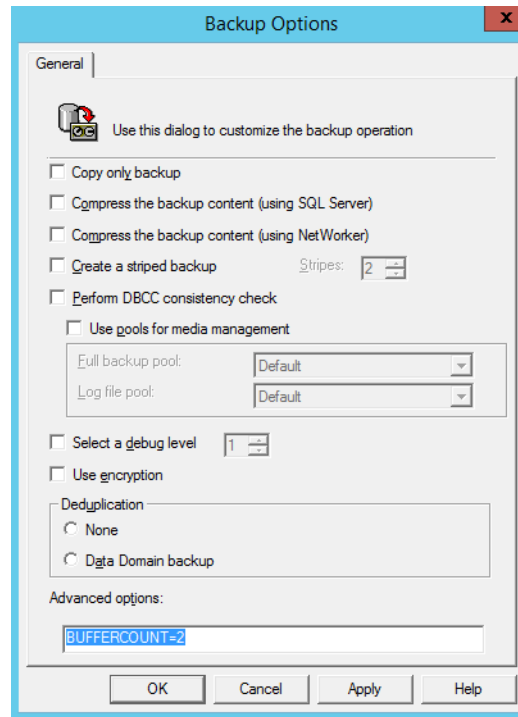
The following figure shows the **Backup** window with a hierarchical list of SQL Server data objects.

Figure 18 The Backup window with a hierarchical list of SQL Server data objects



6. In the left pane, expand an object. The descendants of the object are listed in the right pane.
7. Mark one or more objects for backup. Some SQL Server data items are visible in the **Backup** window, but they cannot be marked. If certain database options are set, the SQL Server prohibits a file or file group backup. These unavailable objects appear grayed out to distinguish them from objects that are available to back up.
8. In the **Backup Options** dialog box, select the required attributes before starting the backup.

The following figure shows the **Backup Options** dialog box.

Figure 19 The General tab in the Backup Options dialog box

The selected attributes are applied to each marked object for all backup operations that are performed until the NetWorker User for SQL Server GUI is closed. Selected pools, however, are stored in the registry and persist from one session to another. If a field is not selected, the default option is used.

9. (Optional) To perform a copy-only backup, select **Copy-only backup**.
A copy-only backup backs up the SQL server without interrupting the existing backup schedule or log chain.
10. (Optional) To apply XBSA compression to all marked databases before the backup data is written to the storage device, select **Compress the backup content (using NetWorker)**. In the same manual backup, you cannot back up some databases with compression and others without.

Compressing data for a backup generates less network traffic and uses less backup media space, but it consumes additional CPU resources. Most tape devices perform compression, which makes software compression unnecessary.
11. (Optional) To create a striped backup by using the SQL Striped feature, select **Create a striped backup**. If you select this option, the **Stripes** box is enabled.

In the **Stripes** box, type or select the number of stripes to use during the backup operation. The topic [Striped Backup and Recovery](#) provides more information about striping.

NMM supports a maximum of 32 stripes. However, the number of stripes cannot exceed the value set for NetWorker client parallelism.
12. (Optional) To run a database consistency check (DBCC) on a database, select **Perform DBCC consistency checks**.

If the DBCC check fails, the database backup is skipped and backup status of the protection group appears as failed. This process ensures that all backed-up databases are suitable for recovery.

13. (Optional) To view the list of media volume pools for storing the backup contents, select **Use pools for media management**.

This option applies only to manual backup operations performed from the NetWorker User for SQL Server program or from a command prompt. If this attribute is checked, the **Full Backup Pool** list and the **Log File Pool** list are enabled.

14. (Optional) To select volume pools, select **Use pools for media management**.
15. (Optional) Select a pool from the **Full backup pool** list to store full SQL Server backups, including databases and file groups.
16. (Optional) To store transaction log or differential (level 1) SQL Server backups, select a pool from the **Log file pool** list.

The pool names in the **Full backup pool** list or **Log file pool** list were created with the NMC and stored in the operating system registry.

17. (Optional) To define the level of troubleshoot information to be sent to the backup status window during the backup operation, select the **Select debug level**. Levels range from 1 to 9, with 1 representing the least amount of information.
18. (Optional) To specify that data is backed up with AES encryption, select **Use encryption**.

Data is encrypted with the default or current pass phrase which is provided by the NetWorker Server. If the NetWorker Server has a different pass phrase at recovery time, specify the pass phrase that was used at the time of backup.

The *NetWorker Administration Guide* provides detailed information about AES encryption and setting the pass phrase.

Note

It is recommended that no form of encryption or compression be used with Data Domain because these reduce the deduplication ratio.

19. (Optional) To specify deduplication, in the **Deduplication** area, select either of the following:
 - **None**—If data deduplication is not required.
 - **Data Domain backup**—To enable client-side Data Domain Boost deduplication backups.
20. (Optional) To add advanced SQL options for the backup, in the **Advanced options** field, type the required option. The only advanced backup option that is currently supported is the `BUFFERCOUNT` option. This option specifies the total number of IO buffers that can be used during a backup operation. The syntax is as follows:

```
BUFFERCOUNT=number_of_IO_buffers
```

21. Click **OK** to close the **Backup Options** window.

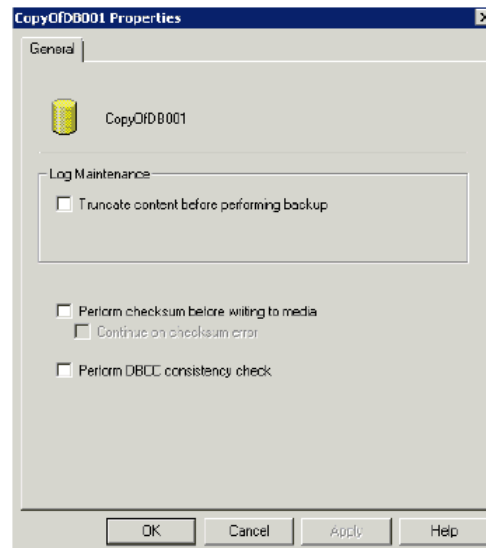
Before starting the backup operation, you must set the backup properties for each marked database. If a property is not specified, the default is used. After the backup operation completes and the **Backup** window closes, the property values revert to the default options.

22. In the **Backup** window, select one or more databases to set the backup properties, and then right-click the marked databases. Select **Properties** from the shortcut menu.

The **Properties** dialog box appears.

The following figure shows the **Properties** dialog box.

Figure 20 Properties dialog box



23. (Optional) In the **Log Maintenance** area, to truncate the transaction logs before backup, select **Truncate content before performing backup**.
24. (Optional) To perform a checksum operation with the backup and save the information to the backup media, select **Perform checksum before writing to media**. Another checksum is performed before a restore to ensure that the information matches the backup.

A checksum detects a fragmented backup or restore state. NMM verifies the checksum by calculating a local result and comparing it with the stored value. If the values do not match, you can choose to continue the backup or restore operation.

25. (Optional) Select **Continue on checksum error** if you want to continue with back up and restore operations when errors are detected.
26. (Optional) Select **Perform DBCC consistency checks**.
If the database consistency check (DBCC) fails, the database backup is omitted and the group backup status appears as failed. This feature ensures that all backed up databases are suitable for recovery.

27. Click **OK**.

28. From the **File** menu, select **Start Backup**.

Results

The backup operation starts.

From the NMC, in the **Backup Status** window, you can monitor the backup status messages.

After the backup is finished, a **Backup Completed** message appears. The amount of time that is required to back up a database depends on database size, network traffic, server load, and tape positioning.

From the **File** menu, click **End Backup** to cancel a backup operation.

CHAPTER 5

Scheduled Backups

This chapter includes the following sections:

- [Overview of scheduled backup](#).....88
- [Prerequisites](#).....88
- [Federated backup preferences for Availability Group databases](#).....88
- [Excluding incompatible databases in backups](#).....90
- [Configuring scheduled backups](#).....91
- [Monitoring scheduled backups](#)108

Overview of scheduled backup

The most reliable way to protect SQL Server data is to schedule backups of the SQL Server at regular intervals. Scheduled backups ensure that all SQL Server data is automatically saved.

Scheduled backups can be set for backup levels full, logs only, or cumulative incremental, and they can be configured to run at any time.

Review the *NetWorker Administration Guide* before reviewing the information in this chapter.

Prerequisites

Before you perform scheduled database backups with a SQL virtual server, you must meet several prerequisites.

Ensure that the following prerequisites are met before you perform scheduled backups:

- The NetWorker interface displays diagnostic attributes. In the **Administration** window, click **View > Diagnostic Mode**.
- When backing up filestream data, in the SQL Server Configuration Manager, on the **Properties** page of the SQL Server instance, select **Allow Remote Clients Access to Filestream Data**. This action is a requirement for backups to be successful. The SQL Server documentation provides details about the SQL Server Configuration Manager.

Prerequisites for cluster environments

You must meet additional prerequisites for cluster environments:

- Create a NetWorker client that is configured for database backups for each SQL virtual server that you want to protect.
- Configure the NetWorker client to either:
 - Back up all databases with `MSSQL:.`
 - Back up specific databases with `MSSQL:User_Database`.
Add more databases to an existing NetWorker client or create a new database client when additional databases are added to an existing NetWorker client or a new database client is created.
- Specify virtual servers by typing the following command at the command prompt:
`-A SQL_virtual_server`

Use the fully qualified domain name for the client configuration and the `-A` option to specify the SQL virtual server, as in the following example:

```
nsrsqlsv -sservername -ASQL_virtual_server
```

Federated backup preferences for Availability Group databases

The SQL Server backup preference for the Availability Group can be configured through the NMM SSMS plugin GUI or the Transact-SQL command.

NMM supports federated backups, during which NMM detects the SQL Server preferred backup setting for the Availability Group and performs the backup at the

preferred node. The database administrator can set the backup priority for the Availability Group or a database in the Availability Group and nominate a particular replica for the backup. This feature improves backup performance.

SQL Server 2012 or later provides the following options that determine where backups run for a specific AG replica:

- **AUTOMATED_BACKUP_PREFERENCE**—Specify any one of the following options:
 - **PRIMARY**—Specifies that the backups will always occur on the primary replica. This option is useful if you need backup features, such as creating differential backups, that are not supported when backup operations run on a secondary replica.
 - **SECONDARY_ONLY**—Specifies that backups will never be performed on the primary replica. If the primary replica is the only replica that is online, the backup will not occur.
 - **SECONDARY (Prefer Secondary)**—Specifies that backups should occur on the secondary replica. If the primary replica is the only replica online, the backup is performed using the primary replica. This is the default option.
 - **NONE (Any replica)**—Specifies that you prefer that backup jobs ignore the role of the availability replicas when choosing the replica on which to perform backups. Note backup jobs might evaluate other factors such as backup priority of each availability replica in combination with its operational state and connected state.
- **BACKUP_PRIORITY =n**—Specifies your priority for performing backups on this replica relative to the other replicas in the same availability group. The value is an integer in the range of 0 - 100. These values have the following meanings:
 - 0 indicates that this availability replica will never be chosen for performing backups.
This choice is useful for a remote availability replica for which you never want backups to fail over.
 - 1 indicates the lowest priority that an availability replica could be chosen for a backup operation.
If **BACKUP_PRIORITY = 1**, the availability replica will be chosen only if no higher priority availability replicas are available.
 - 100 indicates the highest priority that an availability replica could be chosen for a backup operation.

Follow the procedures provided in the article “Configure Backup on Availability Replicas (SQL Server)” on the Microsoft MSDN website <http://msdn.microsoft.com/en-us/library/hh710053.aspx> to decide which replica the backups will be run on.

Note

Federated backup operations fail if a replica has the **Readable Secondary** field set to **No**. For backups to succeed, in the **Availability Group Properties** window, set the **Readable Secondary** configuration option to either **Yes** or **Read-intent only** for each of the primary and secondary replicas.

Note

The following limitations apply to SQL Server backups of a secondary SQL Server replica in an Availability Group:

- For full backups of a secondary replica, SQL Server supports only copy-only backups.
 - Backup promotion is not supported for copy-only backups, including copy-only transaction log backups.
 - SQL Server does not support differential backups on secondary SQL Server replicas.
-

Excluding incompatible databases in backups

If you schedule a backup of a client with the `MSSQL: save set` or named instance `save set`, and if any of the databases in the protection group are incompatible, the scheduled backup fails.

Note

Incompatible databases are silently omitted during manual (non-scheduled) backups that are started from NMM on the client computer.

A database that is in any of the following states causes a scheduled backup to fail because the database is part of a previously configured backup:

- Standby
- Mirror copy
- Recovering
- Restoring
- Recovery Pending
- Suspect
- Offline
- Not recovered
- Loading
- Prerecovery

Definitive results about incompatible databases are available in the `daemon.log` file on the NetWorker server and in the `nsrsqlsv.log` file on the client computer.

Examples of error messages listed in the `daemon.log` file after a backup failure:

- Database 'Acme' cannot be opened because it is offline.
- Processing Acme failed, the item will be skipped.
- Database 'Acme' is in warm-standby state (set by executing `RESTORE WITH STANDBY`) and cannot be backed up until the entire load sequence is completed.
- Processing Acme failed, the item will be skipped.

To exclude databases that are not in a state that can be backed up, perform either of the following steps during client resource configuration:

- In the **Client Backup Configuration** wizard on the **Specify the Backups Options** page, select **Skip databases that are in an incompatible state**.
- In the NetWorker Administrator GUI, in the **Client Properties** dialog box, in the **Application Information** field, type `NSR_SKIP_NON_BACKUPABLE_STATE_DB=TRUE`.

Configuring scheduled backups

To configure a client resource, you must set up backup levels, configure a client resource, and then assign data protection policies.

All the procedures that are described in the following table must be performed on a NetWorker server through the NetWorker Administration GUI. Review the *NetWorker Administration Guide* for details.

Table 22 Tasks for configuring SQL Server VDI backups

Tasks	Considerations
Setting up backup levels	<p>The availability of a backup level depends on the type of data that are selected for backup and any SQL Server settings on those objects. You can specify any of the following backup levels:</p> <ul style="list-style-type: none"> • Full • Logs-only • Cumulative incremental
Configuring a client resource	<p>You can use either of the following methods to configure the client resource:</p> <ul style="list-style-type: none"> • Client Backup Configuration wizard • Client Properties dialog box
Setting data protection policies	<p>To set up the required data protection policies, complete the following tasks:</p> <ul style="list-style-type: none"> • Create a protection group—The type of group that you create depends on the actions that you plan to perform for the group. • Create a policy—When you create a policy, you specify the name and notification settings for the policy. • Within the policy, create a workflow—When you create a workflow, you specify the name of the workflow, the schedule for running the workflow, notification settings for the workflow, and the protection group to which the workflow applies. • Create one or more actions for the workflow.

Table 22 Tasks for configuring SQL Server VDI backups (continued)

Tasks	Considerations
	The <i>NetWorker Administration Guide</i> provides more information about data protection policies.

Setting up backup levels

You can specify backup levels other than database full, database cumulative incremental, and database logs only.

The availability of a backup level depends on the type of data that are selected for backup and any SQL Server settings on those objects, as listed in the following table.

Table 23 Backup levels for SQL Server data

SQL Server data objects	Supported SQL Server backup levels		
	Full	Cumulative incremental	Logs only
All databases of SQL default or named instances	Yes	Yes	Yes
Specified databases	Yes	Yes	Yes
All file groups in specified databases	Yes	Yes	N/A
Filestream data in specified databases	Yes	Yes	Yes
Specified file groups in specified database	Yes	Yes	N/A
Specified files in file groups in specified databases	Yes	Yes	N/A

For SQL Server data objects for which logs only backup can be performed, ensure that the SQL Server database options are correctly configured. The Microsoft SQL Server documentation provides more information. Individual items are subject to promotion.

Configuring a client resource

Use NMC to configure each SQL Server host that is to be backed up as a NetWorker client resource. Multiple SQL Server databases that exist on the same SQL Server host can be configured as separate NetWorker client resources.

By default, the system account that is the `nsrexecd` service account does not have the system administrator (sysadmin) role for the SQL server. Therefore, when you configure a client resource, you must set the username and password of a Windows account that can back up the SQL server. The required privileges are explained in the section [Access privileges for backup and recovery](#) on page 44.

You can configure a client resource for scheduled backups in NMC using either of the following tools:

- The Client Backup Configuration wizard
- The **Client Properties** dialog box

Client resources for SQL Server clusters

For federated backups using NMM to succeed, there must be a client resource for the Windows cluster name from which the backup is performed and dummy client resources for the other participating nodes in the cluster.

It is recommended that you use the Client Backup Configuration wizard to configure client resources for cluster environments. The Client Backup Configuration wizard simplifies configuring client resources for SQL Server 2012 or later scheduled backups in a Windows cluster environment by creating dummy client resources for the participating cluster nodes. To use the Client Backup Configuration wizard in a federated environment, ensure that the active node of the Windows cluster contains a SQL Server 2012 or later instance.


If you create the client resources manually using the **Client Properties** dialog box, you must manually create dummy client resources for each participating node in the cluster.

Note

NMM supports configuring federated backups of an AG that is configured with multiple SQL server instances, including server instances in a failover cluster. You must use the **Client Properties** dialog box to configure this type of backup.

Configuring a client resource with the Client Backup Configuration wizard

You must complete the required steps to configure a client resource with the Client

Backup Configuration wizard. Click the question mark icon  at the lower left of each page of the Client Backup Configuration wizard for details about each field on the page.

When you configure a client resource for a SQL Server cluster, remember the following points:

- Do not use a short name alias for a virtual server node that is not registered on the NetWorker server with a fully qualified domain client name.
- Select the SQL Server cluster instance or individual databases for backup. Do not select databases from two instances for backup. Databases that are in offline, restoring, or loading state are unavailable for selection.

Note

If you are using NetWorker server 8.2.3 or later and NMM 18.1:

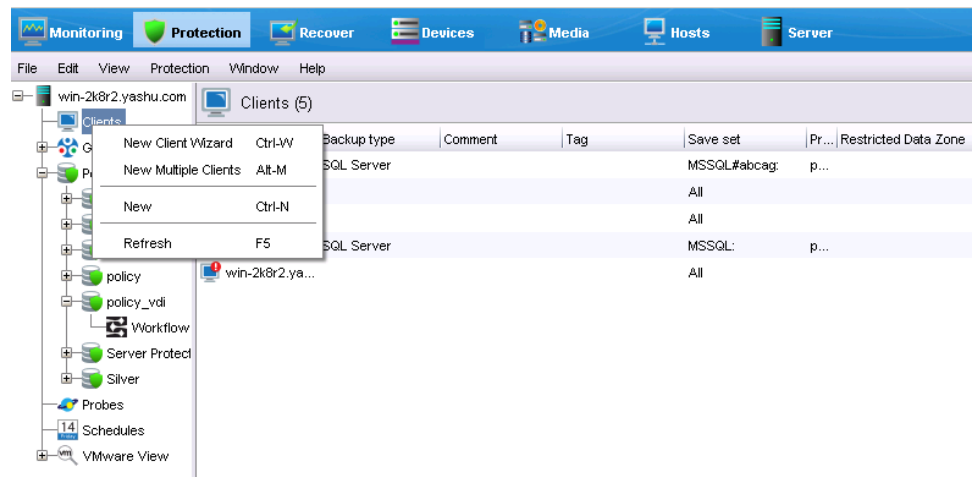
- Configure a regular NetWorker backup group instead of configuring a data protection policy. Do not enable the Snapshot option.
- For the Client Backup Configuration wizard to function correctly, ensure that JRE 8 is installed on the system where NMC is used. While the NMC for NetWorker 8.2.3 and later is compiled with JRE 7, the NMM 18.1 Java plugin for NMC is compiled with JRE 8.
- The procedure to create a client resource is different between NetWorker server 8.2.3 and NetWorker server 18.1. Follow the procedure provided in the *NetWorker Module for Microsoft for SQL VDI User Guide* version 8.2 SP1 when you are using NetWorker server 8.2.3 or later.

Procedure

1. In the **NetWorker Administration** window, click the **Protection** tab.
2. In the expanded left pane, right-click **Clients** and select **New Client Wizard**.

The following figure shows the NetWorker **Administration** window menu option to start the Client Configuration wizard.

Figure 21 Starting the Client Configuration wizard



The **Client Backup Configuration** window opens to the **Specify Client Information** page.

3. On the **Specify Client Information** page, specify the following information:
 - In the **Client Name** field, type the hostname or fully qualified domain name (FQDN) of the SQL Server, Windows cluster, or availability group listener.

Note

Do not type the IP address as the client name.

- In the **Comment** field, type a description for the client resource.
- In the **Tag** field, type one or more tags to identify this client resource for the creation of dynamic client groups for data protection policies. Dynamic client groups automatically generate a list of clients for a data protection policy that is based on the tags assigned to the client and group.

- In the **Group** field, select the group that was previously created for the required backup type.
- In the **Type** area, select **Traditional**.

The following figure shows the Client Configuration wizard **Specify Client Information** page.

Figure 22 Specify Client Information page

4. Click **Next**.
5. On the **Specify Backup Configuration Type** page, specify the following information:
 - In the **Client Operating System** field, the client operating system that is used during the configuration setup is automatically displayed.
 - In the **NetWorker Version** field, the NetWorker version that is used during the configuration setup is automatically displayed.
 - From the **Available Application** list, select **SQL Server**.
 - Do not select **Enable NetWorker Snapshot Management on the selected application**.

The following figure shows the Client Configuration wizard **Specify Backup Configuration Type** page.

Figure 23 Specify Backup Configuration Type page

Specify Backup Configuration Type

Select the backup configuration type that you want to configure for the host. The Available Applications table provides a summary of the available backup application types, based on the NetWorker or module software installed on the client host.

Client Details

Client Operating System: Windows Server 2016 Datacenter 0.0

NetWorker Version: 99.0.99.8819.Build.8819

Extended Client Installed: No

Available Applications

Available Applications	Support NetWorker Snapshot Management
Filesystem	No
SQL Server	No

☐ Enable NetWorker Snapshot Management on the selected application

< Back Next > Cancel

6. Click **Next**.
7. On the **Select the NetWorker Client Properties** page, specify the following information:
 - In the **Priority** field, select the priority level.
 - In the **Parallelism** field, select the level of parallelism.
 - In the **Remote Access** field, type the required attributes.
During the client resource configuration, you can control client recover access with the attributes in the **Remote Access** field. The **Remote Access** attribute displays a list of the users that can recover save sets for a client. Depending on the level of security that the files require, add or remove user names.
 - In the **Data Domain Interface** field, select the device type from the list.
 - In the **Block Based Backup** field, leave this option clear. This feature is not supported for SQL Server VDI.
 - The **Client Direct** field, selected by default, enables the client to bypass the storage node and send backup data directly to the storage device.

Note

To use the default NetWorker Client settings, do not update the options that are provided on the page.

8. Click **Next**.
9. On the **Select the Backup Objects** page, from the prepopulated list, select the SQL Server instance at root level or individual databases for backup.

Note

Do not select databases from two or separate instances for backup. Also, databases that are in offline, restoring, or loading state are unavailable for selection.

10. Click **Next**.
 11. On the **Specify the Virtual Server Credentials** page, in the **Remote User Name** and **Password** fields, type the login information for a Windows account that can backup the SQL server.
-

Note

Before performing this step, ensure that the User Account Control for Windows Server 2008 R2 is disabled for administrators. [Access privileges for backup and recovery](#) on page 44 provides details.

12. Click **Next**.
The **Specify the Backup Options** page appears.
13. (Optional) On the **Specify the Backup Options** page, select the following optional settings as required:
 - Select **NetWorker compression** to use NetWorker compression during the backup.
 - Select **256-bit AES software encryption** to use 256-bit AES software encryption.
 - Select **SQL server compression** if you are using SQL Server 2008 or later.
 - Select **Perform checksum before writing to media** to perform a checksum operation with the backup and save the information to the backup media, and then, if required, select **Continue to backup even on checksum error**.
 - Select **Do not truncate log contents when backing up transaction logs** when you are backing up transaction logs for a SQL server 2005.
 - Select **Skip simple databases during incremental backup** to omit user-created simple databases during a logs only backup.
The Microsoft SQL Server does not support logs only backups for simple recovery model databases. During a logs only backup that includes simple recovery model databases, the simple databases are backed up at level full. This option does not affect the backup of system databases, such as master and msdb, which are promoted to level full during a logs only backup and cannot be omitted from the backup.
This option does not apply to full and cumulative incremental backups.
 - Select **Turn off backup promotion** to disable backup promotion.
 - Select **Skip databases that are in an incompatible state**—Select to skip databases that are not compatible for backup.
 - In the **DBCC Consistency Check Options** area, select the appropriate options:
 - Select **Select all databases** to run database consistency checks (DBCC) on all of the databases. If the DBCC check fails, the database backup is skipped and backup status of the group appears as failed. This action ensures that all backed-up databases are suitable for recovery.

- In the **Databases** box, select specific databases for the DBCC check.
- In the **Striping Options** area, specify striping options:
 - a. To enable striping during a backup, select **Enable Striping**.
 - b. In the **Specify the number of stripes** field, type or select the number of stripes.

The following figure shows the Client Configuration wizard **Specify the Backup Options** page.

Figure 24 Specify the Backup Options page

14. Click **Next**.

The **Client Configuration Summary** page appears.

15. On the **Client Configuration Summary** page, check to ensure that the configuration choices are correct. Click **Back** to make changes, or click **Create** to configure the client resources.

The **Check Results** page appears.

16. On the **Check Results** page, review the messages to ensure that the client was successfully created. You should see the following message:

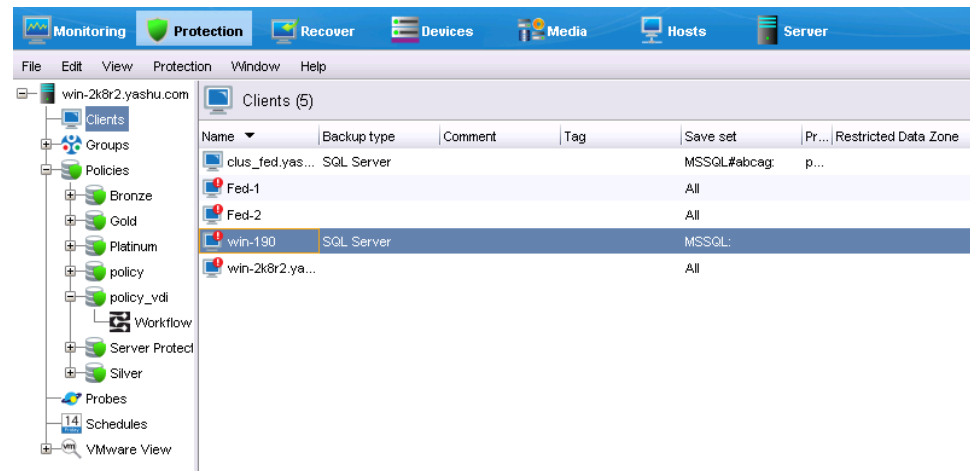
```
Successfully added new client "client_name"
```

17. Click **Finish** to exit the wizard.

After you finish

Open the NetWorker **Administration** window to the **Protection** tab. Click **Clients** in the expanded left pane and ensure that the newly created client is listed.

The following figure shows the NetWorker **Administration** window open to the **Protection** tab.

Figure 25 Viewing clients in the Protection tab**Note**

To make updates to an existing client resource, right-click the client resource and select **Client Backup Configuration > Modify Client Wizard**.

Configuring a client resource manually with the Client Properties dialog box

You must complete the required steps to create a SQL VDI client resource. The *NetWorker Administration Guide* and the *NetWorker Module for Microsoft Administration Guide* provide details about data protection policy. Review the information in these guides before manually creating the client resources.

When you create client resources manually for cluster configurations, ensure that a client resource is created for Windows cluster name and that dummy client resources are created for the other participating nodes in the cluster. If you do not create dummy clients for each of the nodes in the cluster, the backup fails. Do not assign a group for the dummy clients.

Note

If you are using NetWorker server 8.2.3 and later and NMM 18.1:

- Configure a regular NetWorker backup group instead of configuring a data protection policy. Do not enable the Snapshot option.
- The procedure to create a client resource is different between NetWorker server 8.2.x and NetWorker server 18.1. Follow the procedure provided in the *NetWorker Module for Microsoft for SQL VDI User Guide* version 8.2 SP1 when you are using NetWorker server 8.2.x.

Click the question mark icon  at the lower left of each page for details about each field in the page.

Procedure

1. In the **Administration** window, click **Protection**.
2. In the expanded left pane, select **Clients**.
3. From the **File** menu, select **New**.

4. On the **General** tab:

- In the **Name** field, type the fully qualified domain name (FQDN) of the SQL Server, Windows cluster, or availability group listener. If you create multiple client resources for the same SQL Server, use the same name for each.

Note

Do not type the IP address as the client name.

- In the **Comment** field, type a description. If you create multiple client resources for the same NetWorker client host computer, use this attribute to differentiate each resource's purpose.
- In the **Tag** field, type one or more tags to identify this client resource for dynamic client groups in data protection policies.
- In the **Retention Policy** field, select a retention policy from the list. The retention policy determines the time period during which the rolled-over data is available.
- Leave the **Block based backup** option clear. This feature is not supported for SQL Server VDI.
- From the **Directive** list, select an option. Directives are resources that contains special instructions that control how the NetWorker server processes files and directories during backup.
- Select **Scheduled Backups**. The **Backup type** field displays the SQL Server.
- In the **Save Set** field, specify the components to be backed up as listed in the following table.

Note

When nodes within an availability group (AG) reside on SQL server instances with different names, any instance name can be used. NMM will automatically detect the AG that resides on the instance name.

Table 24 Save sets

Save set	Description
All	The All save set appears in the Save Set field by default.
MSSQL:	Typing only MSSQL: yields a backup of all databases on the SQL Server host or SQL default instance.
MSSQL\$<SQL named instance>:	Specifying the named instance with MSSQL: yields a backup of the specified SQL named instance.
MSSQL\$<SQL named instance>:dbName	Specifying the database name along with the named instance along with MSSQL: yields a backup of the database of the specified SQL named instance.
MSSQL:dbName1 [MSSQL:dbName2 MSSQL:dbName3 ...]	Specifying the database names with MSSQL: yields a backup of the specified databases from SQL default instance.

- Select the appropriate option in the **Protection group list** field. If client resources for the same NMM host are added to different backup groups, ensure that the **Start Time** field for each backup group is spaced such that the backups for the host's client resources do not overlap.

- In the **Schedule** field, select a backup schedule.
5. On the **Apps & Modules** tab:
- In the **Remote user** and **Password** fields respectively:
 - For SQL Server 2012 or later, type the remote username and password in the **Remote User** and **Password** fields. To enable NMM to back up the SQL Server virtual server or a mirrored server, type the username for a Windows user account that has SQL Server administrator privileges.
 - For other SQL Server versions, leave the fields empty.
 - In the **Backup command** field, type the `nsrsqlsv` command and any necessary command options.
[Command options for nsrsqlsv](#) on page 70 provides the supported backup command options.
 - For **Application Information** field, type the required values. The following table describes the available application information variables.

Table 25 Application Information field values

Values	Default and valid values
ENABLE_GLR	<ul style="list-style-type: none"> • TRUE (Default): Enables granular-level recovery (GLR) capable backups. • FALSE: Prevents the backup from being GLR-capable. This setting is useful if you are using technology that is not supported for GLR (such as compression, encryption, or tape-type devices) and want to avoid error messages in the operation logs.
NSR_INCLUDE_AAG_DATABASE	<ul style="list-style-type: none"> • AAG_ANYREPLICA (Default): Includes Always On availability group databases in instance-level backup operations. • None: Omits Always On availability group databases from instance-level backup operations.
NSR_SKIP_SIMPLE_DB	<ul style="list-style-type: none"> • TRUE: Omits user-created simple databases during a logs only backup. This option does not affect the backup of system databases, such as master and msdb, which are promoted to level full during a logs only backup and cannot be omitted from the backup. • FALSE (Default): Includes simple databases in the backup. During a logs only backup that includes simple recovery model databases, the simple databases are backed up at level full. <p>This option does not apply to full and cumulative incremental backups.</p>
NSR_SKIP_NON_BACKUPABLE_STATE_DB	<ul style="list-style-type: none"> • TRUE: Omits databases that are in not in a state to be backed up from the backup operation. The savegroup reports success for the SQL instance level backups. • FALSE (Default): Includes databases that are in not in a state to be backed up from the backup operation. The backup of these databases fail and the savegroup reports failure of the SQL instance level backups.
NSR_BACKUP_PROMOTION	<ul style="list-style-type: none"> • ALL: Enables backup promotion to occur in any applicable scenario. • NONE: Disables backup promotion. Logs a warning when backup promotion would normally occur.

Table 25 Application Information field values (continued)

Values	Default and valid values
	<ul style="list-style-type: none"> NONE_WITH_WARNINGS: Disables backup promotion. Logs a warning when backup promotion would normally occur. SKIP_RECOVERY_MODEL: Disables database recovery model change detection. Backup promotion as a result of recovery model change will not occur, but backup promotion in other scenarios will still occur.
NSR_MIRROR_INSTANCE_PORT =<port number>	Specifies the port number when a SQL mirror database is configured with custom TCP port.

- Select **None** in **Proxy Backup**.
- For data deduplication using a Data Domain device, under **Deduplication**, select **Data Domain backup**.

- On the **Globals (1 of 2)** tab, complete fields, as required.
- For federated backups, type all known aliases for the SQL Server host as in the following example:

```
mars mars.emc.com
```

Include both the DNS short name and long name for the SQL Server host.

- On the **Globals (2 of 2)** tab, in the **Remote Access** field, type the user IDs or hostnames of other clients.

This field grants copy restore type permissions, which enables the named hosts to access the NetWorker server and receive directed recover data. If this field is empty, only administrators and users who are logged in to the SQL Server host have access.

For mirroring, this should be the same user account and password that you used to set up the mirroring relationship. If multiple accounts have been set up, only one must be specified. According to Microsoft SQL documentation, a domain account must be used to set up the mirroring relationship.

- Click **OK**.
- Run the backup from the savegroup.

Configuring cluster client resources with the Client Properties dialog box

When you create client resources manually for cluster configurations, ensure that a client resource is created for Windows cluster name and that dummy client resources are created for the other participating nodes in the cluster. If you do not create dummy clients for each of the nodes in the cluster, the backup fails. Do not assign a group for the dummy clients.

Creating client resources for each cluster node

When you create client resources, edit the client resource for each physical node of the cluster. In the NetWorker Administration GUI, in the left pane, select **Clients**, and then in the right pane, right-click to create a new resource.

On the **General** tab, in the **Name** field, type the fully qualified domain name for the cluster node name. For example:

```
wash-2k.belred.emc.com
```

Creating client resources for a virtual server

You must create client resources for each virtual server in the cluster.

Procedure

1. Open the NetWorker Management Console.
The **NMC Enterprise** window appears.
2. Right-click the server and select **Launch Application**.
The NetWorker Administration GUI appears.
3. In the NetWorker Administration GUI, click **Protection**.
4. In the left pane, select **Clients**.
5. From the **File** menu, select **New**.
6. Click the **General** tab.
7. In the **Name** field of the client resource, type the short name for the SQL virtual server without the domain specification. This field should contain the name of the virtual server, not a node name.
8. On the **Apps & Modules** tab (for NetWorker), in the **Backup command** field, type the `nsrsqlsv` command with the necessary options.

For virtual server backups, the `-A SQL_virtual_server_name` command option is required.

9. On the **Globals (2 of 2)** tab, in the **Remote Access** field, add entries similar to the following one to grant access to all physical nodes in the cluster.

RemoteUser@physicalnode_hostname

where:

- *RemoteUser* is the account under which the backup will run.
- *physicalnode_hostname* is the fully qualified domain name of the physical node.

The **Remote Access** field enables the NMM server to access the cluster node to authenticate the computer (on which the virtual server is running) as an NMM client before any backup or restore operation begins. Follow this step for each virtual server client resource in the cluster.

10. On the **Apps & Modules** tab, in the **User Name** and **Password** fields, add the username and password, respectively, for a Windows user account that has both SQL Server administrator privileges and Windows administrator privileges. The **User Name** and **Password** fields enable NMM to back up the SQL Server virtual server. Follow these steps for each virtual server Client resource in the cluster.
11. Use NMC to start the backup group manually, or wait for the next scheduled backup to occur.

Setting data protection policies

Note

NetWorker 8.2.x does not use data protection policies. If you are using NetWorker server 8.2.x and NMM 18.1, follow the steps provided in the *NetWorker Module for Microsoft for SQL VDI User Guide* version 8.2 SP1 to configure a NetWorker Group.

Complete the following tasks to set data protection policies for scheduled backups:

- Create a protection policy
- Create a workflow
- Create an action
- Create a protection group

Creating a policy

Policies provide an organizational container for the workflows, actions, and groups that support and define the backup, management, and system maintenance actions that you want to perform.

Procedure

1. Open the NetWorker Administration GUI. From the **Protection** tab, right-click **Policies** from the expanded left pane, and then click **New**.

The **Create Policy** window appears.

2. In the **Name** field, type a name for the policy.

The following figure shows the **Create Policy** window.

Figure 26 Create Policy window

3. Click **OK**.

Results

The policy is created and the NetWorker Administration GUI opens to the page for the newly created policy.

After you finish

Add workflows, actions, and protection groups to the policy.

Creating a workflow

Workflows define the start time for a series of actions, the frequency in which the actions run, the order of actions in a sequence, and the protection group to which the workflow applies.

Before you begin

Create a policy.

Procedure

1. Open the NetWorker Administration GUI.
2. From the **Protection** tab, click the policy that you must create a workflow for.
The **Policy** page opens.
3. Click **Create a new Workflow**.

The **New Workflow** window opens.

The following figure shows the **New Workflow** window.

Figure 27 Creating a workflow for the policy

New Workflow

Identity

Name: Workflow1

Comment:

Notifications

Send notification: Set at Policy Level

Running

Enabled: ☒

AutoStart Enabled: ☒

Start Time: 9:00 PM

Interval: 24 Hours

Restart Window: 24 Hours

Groups

Groups:

Actions:

Name	Comment	Type	Subtype	Enabled	Previous	Concurrent	Schedule Di...	Retention	Pool	St...
------	---------	------	---------	---------	----------	------------	----------------	-----------	------	-------

OK Reset Cancel

4. In the **Name** field, type a name for the workflow and either specify the workflow settings or accept the default settings.
5. Click **OK**.

Results

The workflow is created and the NetWorker Administration GUI opens to the page for the newly created workflow.

Creating an action within a workflow

Actions are the key resources in a workflow for a data protection policy and define a specific task, for example, a backup, clone, or snapshot.

Before you begin

Create a policy and a workflow within that policy.

Procedure

1. Open the NetWorker Administration GUI and click the **Protection** tab. Click the policy that you must create a workflow for.
2. Expand the policy in the left pane that contains the workflow you must create an action in, and click the workflow.

The **Workflow** page opens.

3. Click **Create a new Action**.

The Policy Action wizard opens.

4. On the **Specify the Action Information** page, in the **Name** field, type a name for the action.
5. In the **Comment** field, type a description for the action.
6. To ensure that the action runs when the policy or workflow that contains the action is started, select **Enabled**. To prevent the action from running when the policy or workflow that contains the action is started, clear this option.
7. From the **Action Type** list, select the action.
8. In the **Workflow** area, assign the action to a policy and workflow.
9. In the **Period** area, specify the backup schedule.

Note

NetWorker supports backup levels full, logs-only, and cumulative incremental with SQL Server VDI.

The following figure shows the Policy Action wizard.

Figure 28 Specifying action information in the Policy Action wizard

The screenshot shows the 'Policy Action Wizard' window with the 'Specify the Action Information' step selected in the left pane. The main area is divided into three sections: Identity, Workflow, and Period.

- Identity:**
 - Name: backup
 - Comment: (empty text area)
 - Enabled: ☒
 - Action Type: backup
 - Backup Subtype: traditional
- Workflow:**
 - Policy: policy_vdi
 - Workflow: Workflow1
 - Previous: (empty dropdown)
 - Concurrent: ☐
- Period:**
 - Weekly by day (For example Sunday, Friday, etc.): ☒
 - Monthly by day (For example 1st day, 5th day etc.): ☐
 - Click on the day to schedule the action: (calendar grid with Sun-Sat)
 - Buttons: Make All, Full

10. Click **Next**.
11. On the **Specify the Backup Options** page, specify backup properties for the action or accept the default properties, and click **Next**.
12. On the **Specify the Advanced Options** page, set the **Retries** field to 0, and then click **Next**.
13. On the **Action Configuration Summary** page, ensure that the correct selections are displayed.
 - If you must change any specifications, click **Back**.
 - If you are satisfied with the summary, click **Configure** to create the action.
14. Review the messages on the **Action Wizard Results** page and ensure that the action was successfully created.
15. Click **Finish** to exit the wizard.

Results

The action is created within the specified workflow and is displayed in the NetWorker Administration GUI, on the **Workflow** page.

Creating a protection group

Basic client groups define a static list of client resources for an action.

Before you begin

Create a policy, workflow, and action.

Procedure

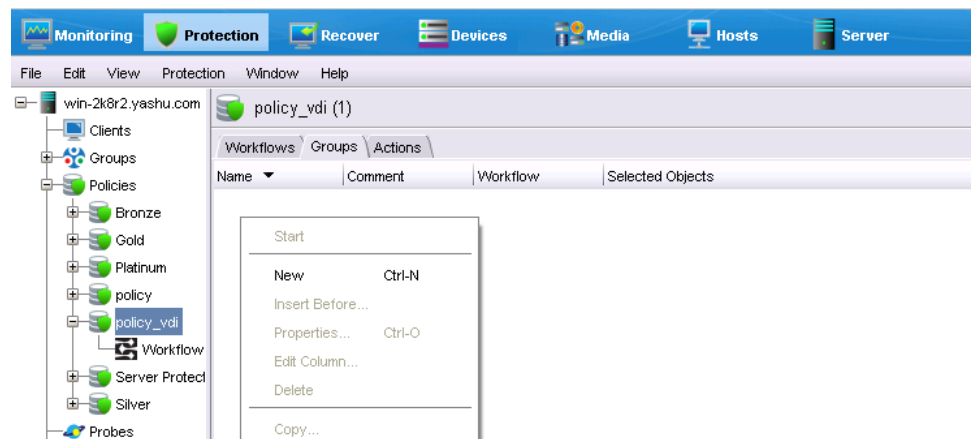
1. Open the NetWorker Administration GUI. From the **Protection** pane, click the policy that you must create a protection group within.

The **Policy** page opens.

2. Click the **Groups** tab. Right-click within the tab, and then select **New**.

The following figure illustrates adding a group to a policy from the **Protection** pane.

Figure 29 Adding a group to a policy from the Protection pane



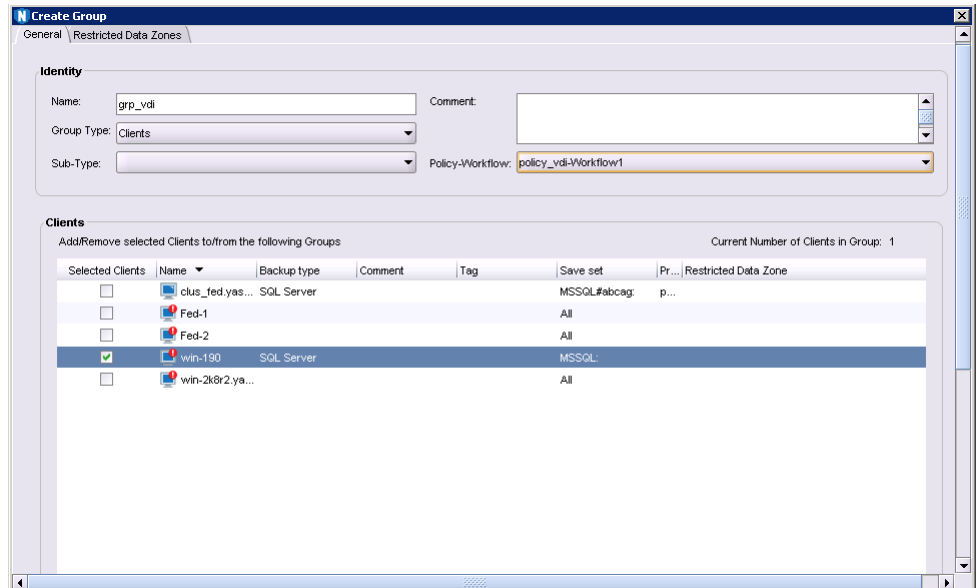
The **Create Group** dialog box opens.

3. In the **Name** field, type a name for the group.
4. From the **Group Type** list, leave the default selection of **Clients**.

5. In the **Comment** field, type a description of the group.
6. Select the workflow in which to assign the group from the **Policy-Workflow** list.

The following figure illustrates specifying a workflow and client in the **Create Group** window.

Figure 30 Specifying workflow and client in the Create Group window



7. Click **OK**.

Results

The group is created and appears in the **Group** tab of the policy in the NetWorker Administration GUI.

Monitoring scheduled backups

During a backup, you can check for status messages in the NetWorker Administration GUI.

The NetWorker server reports the successful and failed database backups, including the databases in an Always On Availability Group. To monitor the success of a backup, use the following procedure:


Procedure

1. From the NMC, open the NetWorker Administration GUI.
2. Click the **Monitoring** tab.
3. Right-click a backup in the **Policies** pane, and then select **Show Details**.


The **Show Messages** window appears.

Results

The **Show Messages** window displays successful and failed backup reports. The following figure shows as example of a successful backup message.

Figure 31 Successful backup messages

The following figure shows as example of a failed backup message.

Figure 32 Failed backup messages

CHAPTER 6

Data Recovery

This chapter includes the following sections:

- [Overview](#) 112
- [Prerequisites](#) 112
- [Performing a recovery from the NMM SSMS plugin GUI](#) 113
- [Performing a recovery from the command prompt](#) 120
- [Performing a recovery from the NetWorker User for SQL Server GUI](#) 135

Overview

You can recover SQL Server data through one of the following tools:

- NMM SQL Server Management Studio (SSMS) plugin
- The `nsrsqlrc` recovery command
- NetWorker User for SQL Server

Note

Consider the following limitations:

- You cannot recover SQL Server data with NMC.
 - You cannot recover data that was backed up by third-party vendors.
 - You cannot select SQL Server Always On Availability Group (AG) databases in the NetWorker User for SQL Server GUI or the NMM SSMS plugin GUI for recovery. All AG databases are unavailable.
-

Prerequisites

Review the prerequisites in this section before performing a recovery.

- Ensure that the user that is performing the recovery is assigned the required privileges, as defined in the section [Access privileges for backup and recovery](#) on page 44.
- Read the SQL Server product documentation to understand the limitations that are associated with recovery types on the various SQL Server versions.
- Ensure that the NetWorker server software is running on the host and the NetWorker Remote Exec Service is started on the SQL Server host.
- Ensure that no applications or services are waiting for or trying to log in to the SQL instance because restoring the SQL master database requires a restart of the SQL instance in single user mode followed by logging in to that instance. Single user mode allows only one administrator to log in.
- Wait for any database backups to finish. SQL Server will not restore a database while another database backup is in progress.
- Wait for any SQL Server startups to finish before starting a restore operation.
- Review the `ERRORLOG` file to determine if a database is being recovered, or search the `ERRORLOG` file for the `Recovery complete` string.
If the `nsrsqlrc` program starts while the SQL Server is recovering databases, the following error message appears:

```
Could not find database ID. Database may not be activated yet or may be in transition.
```
- Ensure that all database users are logged off the database. A restore fails if other users try to use the database during the restore operation.

Note

For the supported SQL Server versions, if the primary file group is not under restore, the online (piecemeal) restore functionality allows users to access a database while a backup or restore operation is in progress.

- Ensure that the SQL Server filestream feature is enabled on the recovery instance when you are recovering filestream data of SQL Server 2008 or later.
- Ensure that the primary file group is part of each piecemeal restore. SQL Server Enterprise Edition databases consisting of multiple file groups that can be restored in stages with piecemeal restore. The database remains offline during the restore process. Piecemeal restore is a multistage process that restores a database to itself or to another location. The initial stage must include the primary file group and optionally other file groups. After the primary file group is restored, you can bring the database online and continue restoring the remaining file groups in subsequent stages.

Performing a recovery from the NMM SSMS plugin GUI

It is recommended that you use the NMM SQL Server Management Studio (SSMS) plugin GUI to perform SQL Server VDI recovery operations.

In the **NetWorker** window, the **Script** view is available in each page. Click **Script** to generate the command prompt equivalent script. You can use the script to create a `.bat` file to perform scheduled backups, automation, and other tasks.

You cannot perform the following types of recovery from the **NetWorker** window in the NMM SSMS plugin GUI:

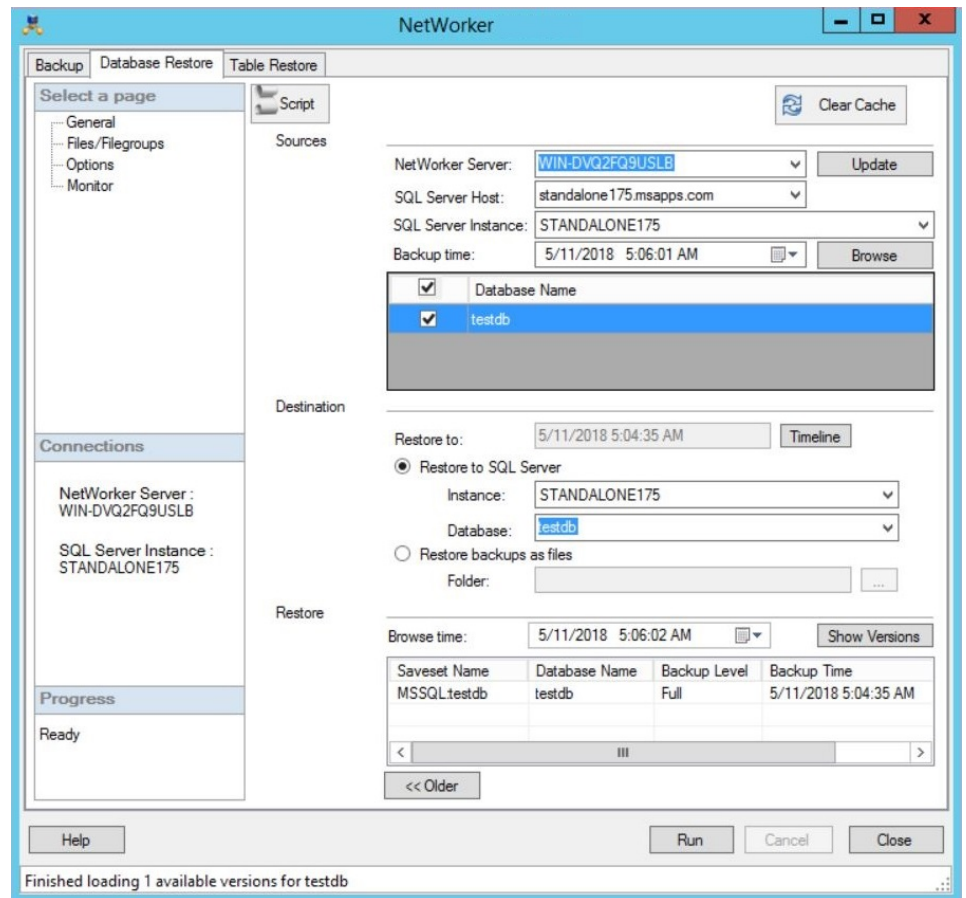
- Copy recovery
- Piecemeal recovery

Procedure

1. In the SQL Server Management Studio, select the SQL Server instance that you want to recover, and then open the **NetWorker** window.
2. In the **NetWorker** window, on the **Database Restore** tab, under **Select a page**, click **General**.

The **General** page appears.

The following figure shows the **NetWorker** window open to the **Database Restore** tab, on the **General** page.

Figure 33 NetWorker dialog box—Database Restore tab General page

Loading icons appear while the information from the last recovery is being loaded. While loading, you can specify other recovery settings. Once the information is loaded, NMM populates all of the fields under **Sources**. To clear the saved settings, click the **Clear Cache** button.

3. Under **Sources**, make the following selections:
 - a. From the **NetWorker Server** list, select the server that contains the backup.
Click **Update** to update the list of available NetWorker servers.

The **SQL Server Instance** list is refreshed after you select the NetWorker server.
 - b. From the **SQL Server Host** list, select the client that contains the backup.
If you used a cluster or availability group listener to take the backup, select the same cluster or listener name for the restore operation.
 - c. From the **SQL Server Instance** list, select the server instance that contains the backup.

Available backups are populated in the database list.
 - d. From the **Backup time** list, select or browse for the time when the database backup was taken.
 - e. From the database list, select one or more databases for recovery.
4. Under **Destination**, choose the location where the backup is restored to:

- To perform a recovery directly to the database, select **Restore to SQL Server**, and then specify the destination:
 - a. From the **Instance** list, select the SQL Server instance.

Note

If an availability group resides on multiple SQL Server instances, use the same SQL Server instance name that was specified when you configured the backup.

- b. From the **Database** list, select the database where the backup will be recovered to.
-

Note

Redirected restore operations for multiple databases is not supported.

- To perform a flat-file recovery, select **Restore backups as files**, and then in the **Folder** field, specify the destination for the files.
5. Under **Restore**, perform the following steps:
 - a. Specify the browse time through one of the following methods:
 - From the **Browse time** list, select a time.
 - Click **Show versions**, and then select a time.

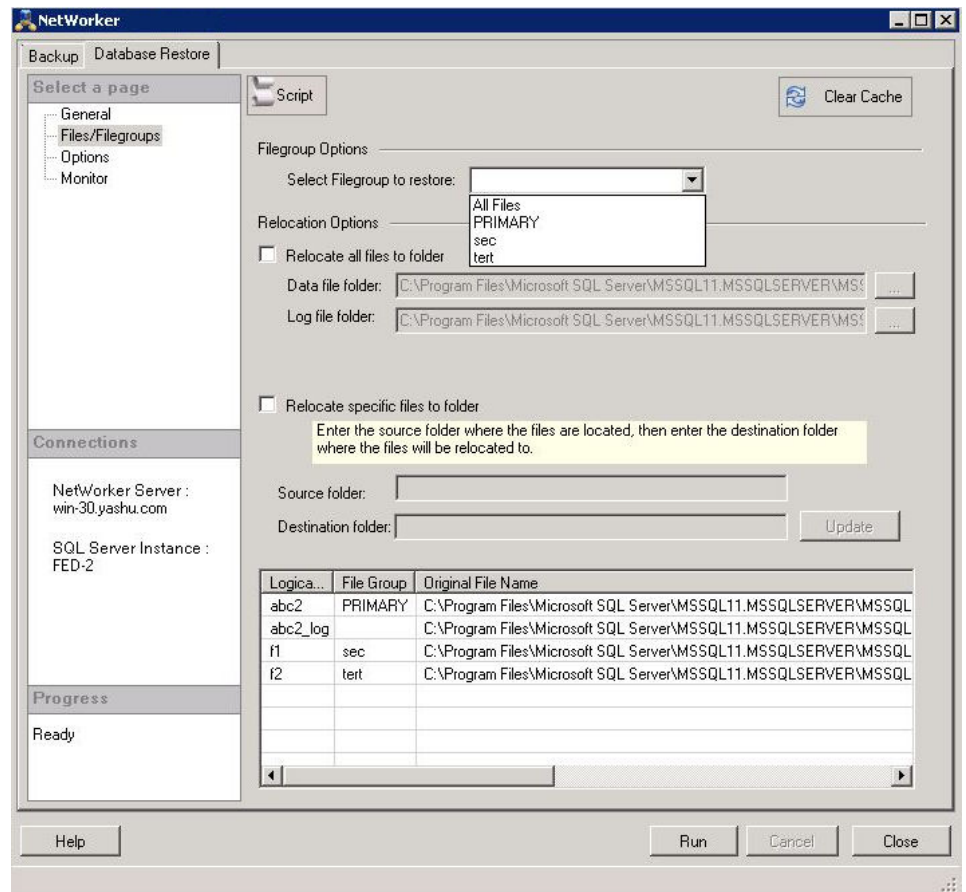
The table is populated with the save sets in the SQL database instance that are available for recovery. The number of available backup versions appears in the bottom bar.

- b. From the table, select the save set that you must recover.

6. To specify optional file settings, under **Select a page**, click **Files/Filegroups**.

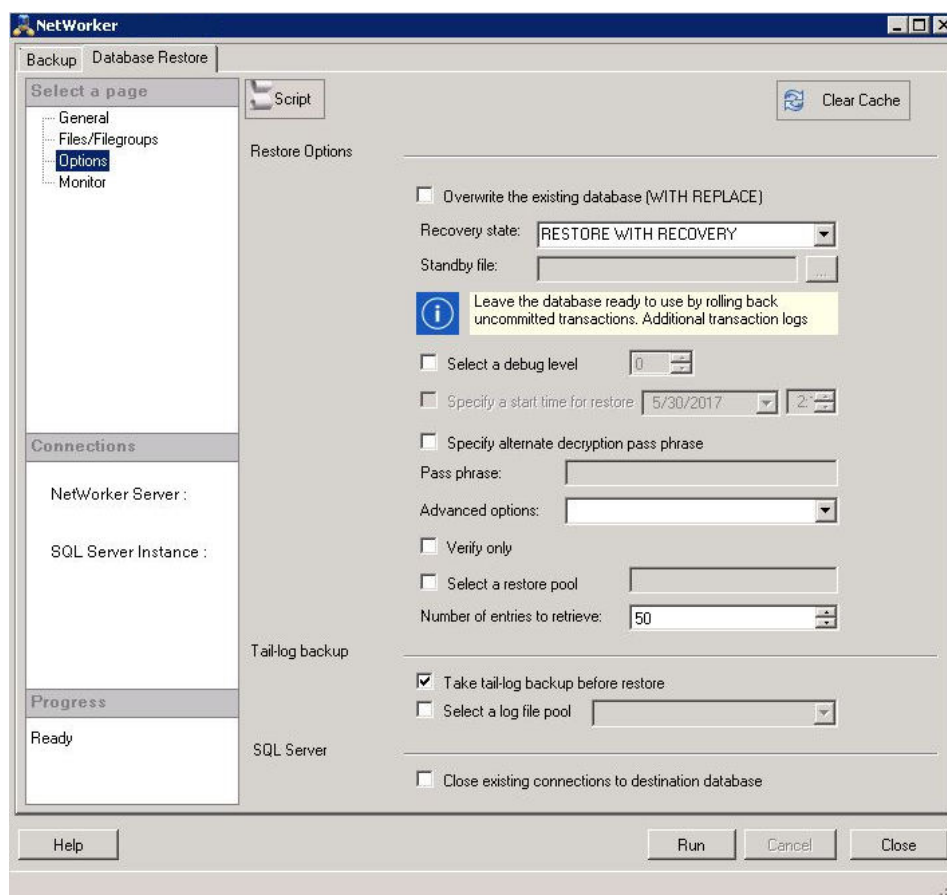
The **Files/Filegroups** page appears.

The following figure shows the **NetWorker** window open to the **Database Restore** tab, on the **Files/Filegroups** page.

Figure 34 NetWorker dialog box—Database Restore Files/Filegroups page

7. (Optional) To filter the list of available database files by file group, under **Filegroup options**, select a file group from the **Select Filegroup to restore** list.
To list all available files, select **All files** from the list.
8. (Optional) To recover files to a different location, select a relocation option:
 - To relocate all datafiles to one folder and the log file to another folder, select **Relocate all files to folder**.
Click the buttons on the right of **Data file folder** and **Log file folder** to specify the destination paths of the datafiles and the log file respectively.
 - To relocate individual file paths, select **Relocate specific files to folder**.
 - a. In the **Source folder** field, type the location path for the datafiles.
 - b. In the **Destination folder** field, type the destination path. The files are recovered to this location.
 - c. Click **Update**.
 Repeat these steps to relocate multiple file paths during a recovery.
9. To specify additional recovery options, under **Select a page**, click **Options**.
The **Options** page appears.
The following figure shows the **NetWorker** window open to the **Database Restore** tab, on the **Options** page.

Figure 35 NetWorker dialog box—Database Restore Options page



10. (Optional) Under **Restore Options**, select the following options as required:
 - To instruct the SQL Server to create the specified database and its related files, even if another database exists with the same name, select **Overwrite the existing database**.
If another database exists with the same name, that database is deleted.
 - To specify a recovery state, in the **Recovery state** field, select one of the following options:
 - To instruct the SQL Server to leave the database in operational state after the restore, select **Normal mode**.
Normal is the default mode.
 - To activate the SQL Server NORECOVERY database restore option for the last stage restored, select **No Recovery mode**.
This mode places the database in a state that cannot be loaded after the restore. However, the database can still process additional transaction log restore operations.
 - To enable the standby undo file feature for SQL Server to use when rolling back the transaction, select **Standby mode**.
 - If the recovery is in standby mode, from the **Standby File** field, browse and select an undo file.
 - To select the level of information recorded in the logs that you can use for troubleshooting recovery issues, select **Select a debug level**, and then select the troubleshooting level. Levels range 0-9, with 0 representing no information and 1 representing the least amount of information.

- To specify a timeframe for flat-file recovery, select **Specify a start time for restore**, and then select a date and time.
The date and time that is selected in this field is the start time for the restore period, and the end time is the time selected on the **General** page, from the **Browse time** list.
- To provide an alternate pass phrase, select **Specify alternate decryption pass phrase**.
If the NMC pass phrase is unavailable, then the alternate pass phrase is used.
- To specify advanced recovery options, in the **Advanced options** field, select or type advanced recovery options. Separate multiple entries with a comma, for example:

`BUFFERCOUNT=2, KEEP_CDC, KEEP_REPLICATION`

The following table outlines the advanced recovery options that NMM supports:

Table 26 Advanced recovery options

Recovery option	Description
<code>BUFFERCOUNT=number_of_IO_buffers</code>	This option specifies the total number of IO buffers that can be used during recovery.
<code>KEEP_CDC</code>	<p>This option enables change data capture (CDC) recovery. When restoring a database with CDC enabled, the recover operation works differently depending on the recovery destination.</p> <p>Use the to <code>KEEP_CDC</code> option to:</p> <ul style="list-style-type: none"> ■ Recover the CDC enabled database on the same SQL instance by overwriting an existing database. In this scenario, <code>KEEP_CDC</code> is optional. ■ Recover the CDC enabled database with a different name on the same SQL instance. In this scenario, <code>KEEP_CDC</code> is required. ■ Recover the CDC enabled database on a different SQL server instance. In this scenario, <code>KEEP_CDC</code> is required.
<code>KEEP_REPLICATION</code>	This option preserves the replication when recovering a published database. This option is required if a database was replicated when the backup was created.

- To verify that the backup is complete and readable, select **Verify only**.
- To specify the clone pool where the save set is restored, select **Select a restore pool**, and then type the name of the restore pool in the field.
- To change the number of entries that are retrieved from the server, in the **Number of entries to retrieve** field, select a number.
The value that is specified in this field is the page size for number of backup versions that are listed for recovery in the **Restore** table in the General view. This value also determines the number of entries that are retrieved when querying backup versions from the server.

The smaller the number that is specified, the faster the backup versions list is available. The minimum value is 1, the default value is 50, and the maximum value is 10000.

11. (Optional) Under **Tail-log backup**, select the following options as required:

- To perform a logs-only backup of the data before you perform a restore operation, select **Take tail-log backup before restore**.
A tail-log backup ensures that the operation backs up the data that has changed from the previous backup.

Note

NMM does not support tail-log backups when multiple databases are being restored.

- To specify a log file pool, select **Select the log file pool**, and then select the log file pool from the list of available pools.
12. (Optional) Under **Server connections**, select **Close existing connections to destination database** to ensure exclusive access to the database during the restore operation if multiple connections exist.
 13. To start the restore operation, click **Run**.

If you have not enabled the instant file initialization feature, a warning message appears before the restore operation starts. To continue with the restore, click **OK**. To abort the operation, click **Cancel**.

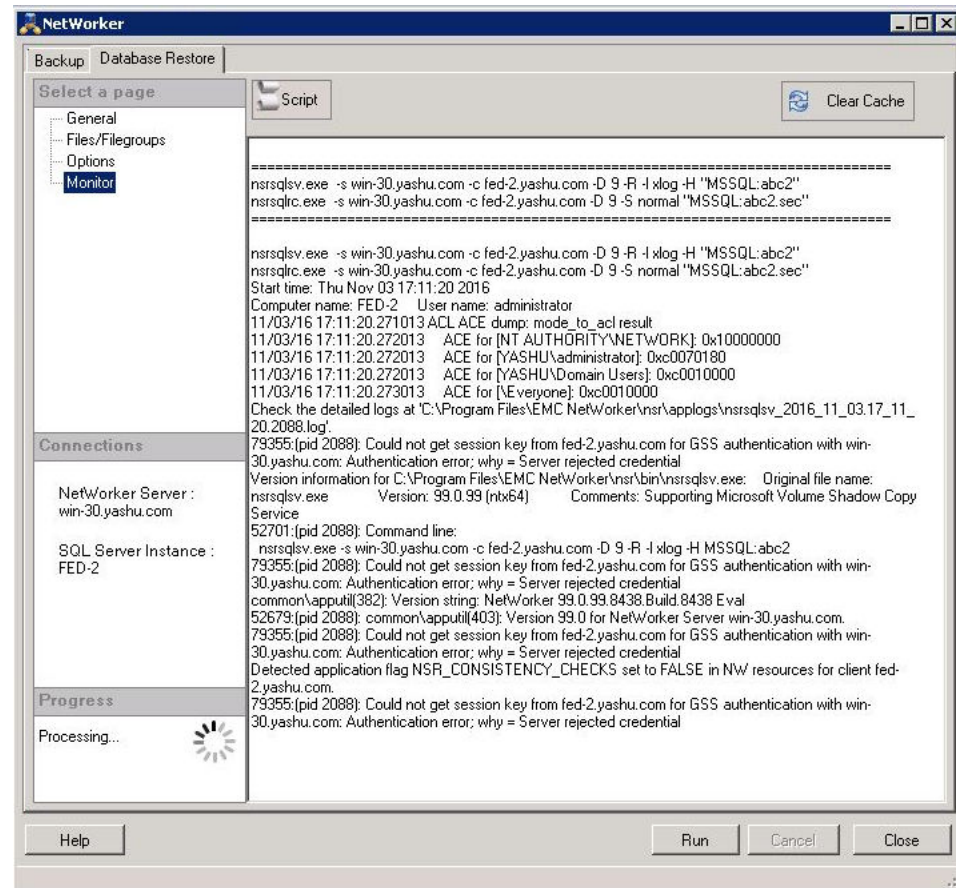
See Microsoft documentation for details on enabling the database instant file initialization feature on the SQL Server.

Results

The restore operation runs.

To view the status of the restore operation, open the **Monitor** page.

The following figure shows the **NetWorker** window open to the **Database Restore** tab, on the **Monitor** page.

Figure 36 NetWorker dialog box—Database Restore Monitor page**Note**

If you do not have the instant file initialization feature enabled, the following warning message will appear in the output:

```
Instant file initialization is not enabled, please enable this feature for improved restore performance.
```

However, the restore operation will continue.

Performing a recovery from the command prompt

The `nsrsqlrc` command recovers specified SQL Server data (files, file groups, and databases) from the NetWorker server.

To start a recovery operation, specify the `nsrsqlrc` command and any of its options at a Windows command prompt.

Before you use the restore command, review the following considerations:

- Ensure that each command option is either lowercase or uppercase and, frequently both the cases of a letter are in the set of command options. Case is very important when specifying command-line flags. For example, `-c` specifies the NetWorker client name, while `-C` copies the database that is being restored to either the same SQL Server or a different SQL Server.

- Ensure that you use correct spacing. Depending on the command option, the space separator between an option and its corresponding argument can be optional, required, or not allowed.

For example, both of the following expressions are valid:

```
-l backup_level
```

```
-lbackup_level
```

The following expression is invalid because a space is not allowed between the + argument and *log_mark_name*:

```
-M + log_mark_name
```

- Ensure that you use brackets ([]) to denote the optional portions of a command. When starting an actual backup or recovery operation, do not include brackets.
- Ensure that data items follow all other command options and parameters at the command prompt.
- If you do not have the instant file initialization feature enabled, the following warning message will appear in the output:
Please ensure that instant file initialization is enabled on the destination SQL Server instance for improved restore performance.

However, the restore operation will continue. See Microsoft documentation for details on enabling the database instant file initialization feature on the SQL Server.

Command syntax for nsrsqlrc

This section contains the command syntax for the `nsrsqlrc` command.

```
nsrsqlrc [-fjqVku]

[-$ instance_name]
[-a "option_key=value"]
[-A virtual_server_name]
[-b pool_name]
[-c client_name]
[-C file=path,file2=path2,...]
[-d MSSQL:destination_dbName]
[-e pass_phrase]
[-f]
[-H "'source_path';'destination_path'"]
[-j]
[-k]
[-M [+|-]log_mark_name]
[-O "Option1, Option2, Option3"]
[-P password]
[-q]
[-R fgName1,fgName2,...]
[-s NetWorker_server_name]
[-S normal|standby:undo_file|norecover]
[-t date]
[-u]
[-U username]
[-V MSSQL:dbName]
```

[-z]

[MSSQL: *dbname dbname.fgName dbName.fgName.filename ...*]

Command options for nsrsqlrc

The following table outlines the list of options available for the `nsrsqlrc` command.

Table 27 Command options for nsrsqlrc

Command options	Descriptions
<code>-\$ MSSQL\$ instance_name</code>	<p>Specifies a named SQL Server instance as the source of the copy recovery.</p> <p>The following example copies the Sales database from the SQL Server prod-sql, instance Venus, to the SQL Server test-sql, instance Mars:</p> <pre>nsrsqlrc -s nw_server.company.com -c prod-sql. company.com -d MSSQL\$MARS: -\$ MSSQL\$VENUS: -C " 'Sales '='D:\Program Files\Microsoft SQL Server \MSSQL.1\MSSQL\Data\Sales.mdf', 'Sales_log'='D:\Program Files\Microsoft SQL Server \MSSQL.1\MSSQL\Data\Sales_log.ldf' " -t "Fri Dec 01 08:01:19 2006" "MSSQL\$VENUS:Sales"</pre> <p>If the <code>-\$</code> switch is used and no instance is named, for example <code>-\$ MSSQL:</code>, or the <code>-\$</code> switch is omitted, the default instance is used.</p>
<code>-a "RESTORE_TO_SOURCE_PATH=TRUE"</code>	Ensures that the backup is restored to the original path by default during a redirected restore operation. The <code>-a "RESTORE_TO_SOURCE_PATH=TRUE"</code> option is overridden if the <code>-C</code> or <code>-H</code> relocation options are used.
<code>-a "CLOSE_SQL_CONNECTIONS=TRUE"</code>	Ensures exclusive access to the target database by changing the database to single user mode during the restore operation to ensure a successful recovery.
<code>-a "DDBOOST_COMPRESSED_RESTORE=TRUE"</code>	<p>Enables the Data Domain Boost compressed restore feature. A compressed restore uses less bandwidth by restoring the backed-up data in a compressed form from the Data Domain system to the SQL Server host.</p> <p>A compressed restore can be beneficial in a constrained bandwidth environment, but may impact the restore performance due to the usage of compression resources on the Data Domain system and SQL Server host.</p>
<code>-a "FLAT_FILE_RECOVERY=TRUE"</code>	Performs flat-file recovery, which recovers the backup as files instead of restoring directly to the databases.

Table 27 Command options for nsrsqlrc (continued)

Command options	Descriptions
<code>-a "FLAT_FILE_RECOVERY_DIR=<destination_directory>"</code>	Specifies the directory where the recovery files are saved for flat-file recovery.
<code>-a "RESTORE_START_TIME=start_date" -t "end_date"</code>	<p>Specifies a time range for flat-file recovery. Applicable backups within the time range are recovered. If the <code>RESTORE_START_TIME</code> variable is not defined when other flat-file recovery options are defined, required backups from the last level full backup until the <code>-t</code> time are recovered. If the last backup within the time range is a level differential backup, then all incremental backups in the time range are not recovered. In this scenario, only full and differential backups are recovered to file.</p> <hr/> <p>Note</p> <p>Follow the <code>nsr_getdate</code> command syntax guidelines when formatting the dates. See the <i>NetWorker Command Reference Guide</i> for details on the <code>nsr_getdate</code> command syntax.</p> <hr/>
<code>-A virtual_server_name</code>	Specifies the virtual server name when the SQL Server is configured to run as a clustered service in an MSCS cluster.
<code>-b pool_name</code>	Specifies the clone pool where the save set is restored from.
<code>-c client_name</code>	Specifies the NetWorker client name from which the SQL Server data is to be recovered.
<code>-C file=path,file2=path2,...</code>	<p>Copies the database that is being restored to either the same SQL Server or a different SQL Server. This option can be used for normal, copy, and partial restores. Use the relocation list to specify new locations for recovered database files. The relocation list is composed of pairs of logical database file names and fully qualified domain database file name relocation paths. Specify the relocation list only when you restore a database. Each file name and relocation path is joined by the equal sign (=), and pairs are delimited in the list by commas.</p> <p>The syntax is:</p> <pre>["] ['] file ['] = ['] path ['] , ['] file2 ['] = ['] path2 ['] , . . . ["]</pre> <p>For example, to copy the Project database from a client host that is named ClientHost1 to a different location on ClientHost1, specify the relocation list for the database files, but do not include the client hostname command option:</p> <pre>nsrsqlrc -s NetWorker_server_name -d MSSQL:CopyOfProject -C "'Project_Data'='C:\Relocation\Path\Project_Data.mdF', 'Project_Data2'='C:\Relocation\Path\Project_Data2.ndF',...</pre>

Table 27 Command options for nsrsqlrc (continued)

Command options	Descriptions
	<p><i>'Project_Log'='C:\Relocation\Path\Project_Log.ldf MSSQL:Project' "</i></p> <p>The relocation list must be enclosed by double quotation marks to allow for spaces in the relocation elements and path names. A logical file name or relocation path must be enclosed by single quotation marks to also allow for spaces. If a file name or path contains a single quotation mark, precede the single quote with another single quotation mark to prevent the NMM from parsing the single quotation mark as a delimiter, for example:</p> <pre>nsrsqlrc -s NetWorker_server_name -d MSSQL:CopyOfProject -C File='C:\Relocate Dir\Path\, ='C:\Relocate Dir\Path\','...', '='C:\Relocate Dir\Path\'' " MSSQL:Project</pre> <p>When no relocation list is specified, NMM reads the source database file names and location from the client index metadata or the backup. This information is used to generate a default list by relocating all files to the default SQL data path for the target SQL Server. The file names are guaranteed to be unique, but sufficient disk space is not ensured.</p> <p>To relocate an entire file path, see the table entry for the -H option. You cannot use the -C option with -H.</p>
-d MSSQL:destination_dbname MSSQL:source_dbname	<p>Performs a copy operation. The copy operation recovers SQL Server data from a client host to another database name on the same client host. The syntax is:</p> <pre>nsrsqlrc -s NetWorker_server_name -C client_name -d MSSQL:destination_dbname MSSQL:source_dbname</pre> <p>where:</p> <ul style="list-style-type: none"> <i>destination_dbname</i> is the name of the SQL database to which the source database is to be recovered. <i>source_dbname</i> is the name of the SQL database to restore. <p>When no relocation list is specified, NMM reads the source database file names and location from the client index metadata or the backup. This information is used to generate a default list by relocating all files to the default SQL data path for the target SQL Server. The file names are guaranteed to be unique, but it is not guaranteed that there will be sufficient disk space.</p> <p>When -C, -M, -R, or -d are used, the list of data items can include only one database. The list of data items must follow all other command options and parameters at the command prompt.</p>
-e pass_phrase	<p>Enables the use of an alternate pass phrase with AES encryption when recovering data.</p>

Table 27 Command options for nsrsqlrc (continued)

Command options	Descriptions
-f	Performs a recovery operation by replacing the target database with the source. This option recovers a source database to an existing, incompatible database of the same name on the target host. This option is also used to recover damaged files.
-H "'source_path';'destination_path'"	<p>Relocates file paths and all files contained in a path to a different location during a restore operation. The syntax for using this option is:</p> <pre>nsrsqlrc -s networker_server_name -c client_name -H " 'source_path';'destination_path'" -t "MM/DD/YY HH:MM:SS" -S restore_mode "MSSQL:source_database"</pre> <p>For example, if the testdb1 database has four datafiles that are in the C:\datapath source path and you want to relocate the database to the D:\datapath destination path, type the following:</p> <pre>nsrsqlrc -s networkerserver -c client -H " 'C:\datapath';'D:\datapath'" -t "11/11/2015 10:30:57" -S normal "MSSQL\$SQLSERVER2012:testdb1"</pre> <p>The -H command option can be used to relocate multiple file paths in the same command line. For example:</p> <pre>nsrsqlrc -s networkerserver -c client -H " 'C:\datapath1';'D:\datapath2';'E:\datapath3';'F:\datapath4';'G:\datapath5';'H:\datapath6'" -t "11/11/2015 10:30:57" -S normal "MSSQL\$SQLSERVER2012:testdb1"</pre> <p>The -H option is supported on standard and redirected restores, which includes the following:</p> <ul style="list-style-type: none"> • Normal restore—Same server and same instance. • Different instance restore—Same server and different instance. • Restore to different database file. • Different server restore. <p>To relocate individual files, see the table entry for the -C option. You cannot use the -C option with -H.</p>
-j	Performs a database consistency check between the backed-up SQL Server data and the recovered SQL Server data.
-k	Perform checksum before reading from media.
-m name	Specifies the cluster or availability group listener name.
-M [+ -]log_mark_name	<p>Performs a SQL Server data recovery of the named transaction that is specified in log_mark_name (for SQL Server 2005 only).</p> <p>The way the mark name is prefixed determines how the data is recovered:</p> <ul style="list-style-type: none"> • When the mark name is prefixed with a plus sign (+), the data is recovered to and includes the named transaction.

Table 27 Command options for nsrsqlrc (continued)

Command options	Descriptions
	<ul style="list-style-type: none"> When the mark name is prefixed with a minus sign (-), the data is recovered up to the named transaction. <p>The <code>log_mark_name</code> should immediately follow the plus (+) or minus (-) symbol. The use of a space separator is not allowed. The default is the plus sign.</p> <p>For example, to recover the SQL data and include the named transaction <code>transaction_mark</code>, type the following command:</p> <pre>nsrsqlrc -sNetWorker_server_name -M +transaction_mark MSSQL:dbName</pre> <p>To recover the SQL data only to the named transaction <code>transaction_mark</code>, type the following command:</p> <pre>nsrsqlrc -s NetWorker_server_name -M - transaction_mark MSSQL:dbName</pre> <p>Specify only one SQL Server database, followed by all the other command options and parameters at the command prompt.</p>
<p><code>-O "Option1, Option2, Option3"</code></p>	<p>This option specifies advanced recovery options. The syntax is as follows:</p> <pre>-O "Option1, Option2, Option3"</pre> <p>You can specify the following advanced recovery options:</p> <ul style="list-style-type: none"> BUFFERCOUNT=number_of_IO_buffers This option specifies the total number of IO buffers that can be used during a recovery operation. KEEP_CDC This option enables change data capture (CDC) recovery. KEEP_REPLICATION This option preserves the replication when you recover a published database. This option is required if a database was replicated when the backup was created. READ_WRITE_FILEGROUPS This option recovers only the read/write (active) filegroups from the backup. This option can only be used to recover backups taken with the <code>READ_WRITE_FILEGROUPS</code> option. You cannot specify filegroup or file level recover targets or use the <code>-R</code> recover option with the <code>READ_WRITE_FILEGROUPS</code> option. <p>After you restore a backup with the <code>READ_WRITE_FILEGROUPS</code> option, any read-only filegroups in the database will enter the recovery pending state and the rest of the filegroups will go online.</p> <p>There are two different scenarios to recover cumulative incremental <code>READ_WRITE_FILEGROUPS</code> backups:</p> <ul style="list-style-type: none"> ■ READ_WRITE_FILEGROUPS This option recovers only the read/write (active) filegroups from the backup. This option can only be used to recover

Table 27 Command options for nsrsqlrc (continued)

Command options	Descriptions
	<p>backups taken with the <code>READ_WRITE_FILEGROUPS</code> option. You cannot specify filegroup or file level recover targets with the <code>READ_WRITE_FILEGROUPS</code> option.</p> <p>If you are relocating multiple filegroups during the restore, you can use the <code>-H</code> option for global relocation.</p> <p>After you restore a backup with the <code>READ_WRITE_FILEGROUPS</code> option, any read-only filegroups in the database will enter the recovery pending state and the rest of the filegroups will go online.</p> <p>There are two different scenarios to recover cumulative incremental <code>READ_WRITE_FILEGROUPS</code> backups:</p> <ul style="list-style-type: none"> – If you have taken a full <code>READ_WRITE_FILEGROUPS</code> backup prior to the cumulative incremental backup, the Microsoft application agent will chain together the two backups in the correct order and complete the restore. – If you have not taken a full <code>READ_WRITE_FILEGROUPS</code> backup but want to use a normal full backup as the differential base for the cumulative incremental <code>READ_WRITE_FILEGROUPS</code> backup, you must first restore the full backup normally and leave the database in no recovery mode, and then apply the <code>READ_WRITE_FILEGROUPS</code> differential backup with the <code>-z</code> option. <p>If you are specifying multiple options, separate each argument with a comma. For example:</p> <pre>nsrsqlrc -sNetWorker_server_name -O "BUFFERCOUNT=2, KEEP_CDC, KEEP_REPLICATION, READ_WRITE_FILEGROUPS" MSSQL:</pre>
<code>-P password</code>	<p>Specifies the SQL Server user password. When the <code>-U</code> username command option is specified, the password command option must also be provided, as follows:</p> <pre>nsrsqlrc -s NetWorker_server_name -U username -P password MSSQL:</pre> <p>Use the SQL Server username and password to log in to the SQL Server with SQL server authentication.</p>
<code>-q</code>	<p>Displays <code>nsrsqlrc</code> messages in quiet mode, which provides minimal information about the progress of the recovery operation, including error messages.</p>
<code>-R fgName1, fgName2, . . .</code>	<p>Performs a partial database recovery or a piecemeal recovery of the specified file groups. The partial database recovery operation restores specific file groups from a single full SQL Server database backup. Add the file groups to the <code>-R</code> command option in a list, with items separated by commas. Use the <code>-C</code> parameter, and specify all files for the database. The <code>-d</code> option is also required:</p>

Table 27 Command options for nsrsqlrc (continued)

Command options	Descriptions
	<p>nsrsqlrc -s NetWorker_server_name</p> <p>-R ["']["']fgName[''],["']fgName2[''],["']...['']["']</p> <p>-C Project_Data=C:\Relocation\Path\Project_Data.mdF, Project_Data2=C:\Relocation\Path\Project_Data2.ndF,..., Project_Log=C:\Relocation\Path\Project_Log.ldF MSSQL:Project-d MSSQL:PartOfProject MSSQL:Project</p> <p>where:</p> <ul style="list-style-type: none"> <i>fgName</i>,... are the names of the file groups to restore. Consider the following points when you list file groups to restore: <ul style="list-style-type: none"> To allow spaces between file group names, enclose the list of file group names within double quotes. Within file group names, enclose each file group name within single quotes. If a file group name contains a single quote, precede the single quote with another single quote to prevent the NetWorker client from parsing the single quote as a delimiter. <p>For example, to accommodate for the space in file group A, the space after the first comma, and the single quote in file group A', use the following syntax:</p> <p>-R "'Filegroup A', 'Filegroup A' '"</p> <p>When an empty relocation list is supplied, use the following syntax:</p> <p>-R ""</p> <i>Project</i> is the name of the SQL database to restore. <i>PartOfProject</i> is the name of the SQL database to which the source database is to be recovered. <p>NMM recovers only the primary file group.</p> <p>When -C, -M, -R, or -d are used, the list of data objects can include only one database. The list of data objects must follow all other command options and parameters at the command prompt.</p>
-s NetWorker_server_name	Specifies the NetWorker server to use for the recovery operation.
-S normal standby:undo_file norecover	<p>Performs the recovery operation according to the specified database restore mode of normal, standby, or no recovery. The syntax is:</p> <p>nsrsqlrc -s NetWorker_server_name -ddestination_dbName -S normal "standby:undo_file" norecover MSSQL:source_dbName</p> <p>where:</p>

Table 27 Command options for nsrsqlrc (continued)

Command options	Descriptions
	<ul style="list-style-type: none"> <i>destination_dbName</i> is the name of the SQL database to which the source database is to be restored. <i>source_dbName</i> is the name of the SQL database to restore. <p>The recovery modes are as follows:</p> <ul style="list-style-type: none"> The normal recovery mode recovers the database in normal, operational mode. The standby recovery mode activates the SQL Server STANDBY option, which forces the database to be in a read-only state between transaction log recovery operations. The norecover recovery mode activates the SQL Server NORECOVER option, which places the database in an unloadable state after the recovery. The database can still process additional transaction log recovery operations. <p>For example, to recover the Project database in normal, operational mode to a new database named NewProjectName, type the following command:</p> <pre>nsrsqlrc -s NetWorker_server_name -S normal -d MSSQL: NewProjectName MSSQL:Project</pre> <p>To recover the database in standby mode, the standby parameter must be immediately followed by a colon, and the standby undo file location and file name must be specified. If a file name or location contains a space, enclose the file name or location within double quotes, for example:</p> <pre>nsrsqlrc -s NetWorker_server_name -S "standby:C:\temp\undo filename" -d MSSQL:NewProjectName MSSQL:Project</pre>
-t "MM/DD/YY HH:MM:SS"	<p>Recovers the SQL Server data on a specified date. If the date of a backup version is before or is equivalent to the specified date, the backup version is recovered. Follow the <i>nsr_getdate</i> command syntax guidelines when you format the date. See the <i>NetWorker Command Reference Guide</i> for details on the <i>nsr_getdate</i> command syntax.</p> <p>To avoid adversely affecting the database, do not use the -t option if you are restoring a file or file group.</p>
-u	Continues the recovery even if a checksum error occurs.
-U username	<p>Specifies the SQL Server username. When the username command option is specified, the -P password command option must also be provided, for example:</p> <pre>nsrsqlrc -s NetWorker_server_name -U username -P password MSSQL:</pre>

Table 27 Command options for nsrsqlrc (continued)

Command options	Descriptions
	Use the SQL Server username and password to log in to SQL Server with SQL server authentication.
<code>-V MSSQL: dbName</code>	<p>Verifies the SQL Server database that is selected for the recovery. The <code>-V</code> command option verifies only that the selected database backup is suitable for restoring, the backup is not recovered. The syntax is:</p> <pre>nsrsqlrc -s NetWorker_server_name -V MSSQL: dbName</pre>
<code>-z</code>	<p>Enables implementation of a recovery plan to restore individual backups from a timestamp.</p> <p>By default, NMM builds the recovery plan, ensuring that all the required backups are available and that they run in the proper order and with the proper options. The <code>-z</code> option removes NMM safety checks.</p> <p>The <code>-z</code> option disables building the recovery plan and recovers only the backups that correspond to the specified timestamp. The <code>-z</code> command is required for all except the first (full) recovery command.</p> <p>Use the <code>-S norecover</code> option with the <code>-z</code> option to prepare the database for more restores. This command is required for all except the last recovery command.</p> <p>The final command recovers the last logs-only backup and brings the database online.</p> <p>The <code>-z</code> option is used in more complex recoveries. The example commands in this section use the following backup history:</p> <pre>savetime 1:00 - full savetime 2:00 - txnlog savetime 3:00 - txnlog</pre> <p>Consider the following example uses of the <code>-z</code> option:</p> <ul style="list-style-type: none"> To restore a database with a single command, specify the most recent backup timestamp or <code>now</code>. The following list provides examples: <ul style="list-style-type: none"> The following example uses the most recent timestamp: <pre>nsrsqlrc ... -t "savetime 3:00" ...</pre> The following example uses <code>now</code>: <pre>nsrsqlrc ... -t "now" ...</pre> <p>These commands restore the entire backup chain from the last backup, regardless of the backup level.</p> To restore a database in a series of independent commands, type a separate command for each backup. You must recover the backup chain in order. For the final command, specify the most recent backup timestamp or <code>now</code>. The following list provides examples: <ul style="list-style-type: none"> The following example uses the most recent timestamp for the final command:

Table 27 Command options for nsrsqlrc (continued)

Command options	Descriptions
	<pre>nsrsqlrc ... -z -S norecover -t "savetime 1:00" ...</pre> <pre>nsrsqlrc ... -z -S norecover -t "savetime 2:00" ...</pre> <pre>nsrsqlrc ... -z -S normal -t "savetime 3:00" ...</pre> <hr/> <p>Note</p> <p>You can specify a point-in-time restore within the last logs-only backup by replacing the save time with the point-in-time. For instance, in the previous example, you can replace 3:00 with 2:45.</p> <hr/> <ul style="list-style-type: none"> ■ The following example uses now for the final command: <pre>nsrsqlrc ... -z -S norecover -t "savetime 1:00" ...</pre> <pre>nsrsqlrc ... -z -S norecover -t "savetime 2:00" ...</pre> <pre>nsrsqlrc ... -z -S normal -t "now" ...</pre> <p>The database is not available for general use until after the final recovery completes. Any missing, incorrect, or out-of-order save times result in SQL Server reporting errors.</p>

Backup and recovery command syntax for SQL Server data

Use the additional command syntax that is shown in the following table to back up or restore SQL Server data with the `nsrsqlsv` and `nsrsqlrc` backup and recovery commands.

You can specify more than one data object and can combine different types of data. You must specify the SQL data objects with the syntax shown in the following table.

Table 28 Command syntax for SQL Server data

SQL Server data	Syntax for SQL Server data objects
All databases in the SQL Server storage hierarchy (optional)	MSSQL : Typing MSSQL : yields an instance-level backup of all databases on the SQL Server host.
Specified databases	MSSQL : dbName or [MSSQL : dbName MSSQL : dbName2 ...]
All file groups in specified databases	MSSQL : dbName or

Table 28 Command syntax for SQL Server data (continued)

SQL Server data	Syntax for SQL Server data objects
	[MSSQL: <i>dbName</i> . MSSQL: <i>dbName2</i> ...]
Specified file groups in specified database	MSSQL: <i>dbName</i> . <i>fgName</i> or [MSSQL: <i>dbName</i> . <i>fgName</i> MSSQL: <i>dbName</i> . <i>fgName2</i> MSSQL: <i>dbName2</i> . <i>fgName</i> MSSQL: <i>dbName2</i> . <i>fgName2</i> ...]
Specified files in specified file groups in specified databases	MSSQL: <i>dbName</i> . <i>fgName</i> . <i>filename</i> or [MSSQL: <i>dbName</i> . <i>fgName</i> . <i>filename</i> MSSQL: <i>dbName</i> . <i>fgName2</i> . <i>filename</i> MSSQL: <i>dbName2</i> . <i>fgName</i> . <i>filename</i> MSSQL: <i>dbName2</i> . <i>fgName2</i> . <i>filename</i> ...]

Specifying MSSQL before each data object name is optional and does not affect the expression or the resulting operation. However, when MSSQL is specified, it must be followed by a colon (:).

For example, the following two commands are equivalent:

```
nsrsqlsv -s NetWorker_server_name dbName.fgName
nsrsqlsv -s NetWorker_server_name MSSQL:dbName.fgName
```

In a non-clustered, named instance configuration, MSSQL\$ is required, followed by the instance name and a colon. For example:

```
nsrsqlsv -s NetWorker_server_name MSSQL$instanceName:dbName.fgName
```

Syntax for a named instance configuration

When a configuration contains non-clustered named instances of the SQL Server, you must specify the name of the instance before the data.

```
MSSQL$instanceName:[dbName ...] [.fgName ...] [.fileName ...]
```

Note

The `nsrsqlsv` and `nsrsqlrc` commands support specification of only a single instance. If save sets for more than one instance are specified, the backup fails. The `nsrsqlrc` command supports mixing instances for a copy restore operation.

Example 6 Back up all databases in an instance

To back up all the databases for instanceOne, type the following:

```
nsrsqlsv -s NetWorker_server_name MSSQL$instanceOne:
```

Example 7 Recover several file groups in an instance

To recover several file groups for instanceTwo, type the following:

Example 7 Recover several file groups in an instance (continued)

```
nsrsqlrc -s NetWorker_server_name MSSQL$instanceTwo:dbName.fgName
MSSQL$instanceTwo:dbName.fgName2
```

Example 8 Using clustered instance SQL Server virtual server names with -A or -c

Instead of using clustered named instances in this syntax, use clustered instance SQL Server virtual server names with -A or -c options, as shown in the following command:

```
nsrsqlsv -s NetWorker_server_name -A SQL_virtual_server_DNS_name -c
SQL_virtual_server_DNS_name MSSQL:
```

where:

- *NetWorker_server_name* is the hostname of the NetWorker server.
- *SQL_virtual_server_DNS_name* is the Domain Name System (DNS) name for the SQL Server virtual server.

Create a client resource under the SQL Server virtual server DNS name.

For scheduled backups of a SQL Server virtual server client, you do not need to specify the -A or -c option with the SQL Server virtual server name. The `savegrp` process automatically specifies the virtual server name to the `nsrsqlsv` process by using the -m option.

Syntax for names containing a period, backslash, or colon

NMM provides command line syntax that enables you to back up and recover file names, file groups, and databases containing a period (.), backslash (\), or colon (:). By typing a backslash before the period or backslash, the `nsrsqlsv` and `nsrsqlrc` commands interpret the period or backslash as a literal character.

The tables in this section list the syntax for file names, file groups, and databases containing a period, backslash, colon, or any combination of the three.

The following notes apply to the information in the tables:

- The syntax that is shown in the right column applies to both the `nsrsqlsv` and `nsrsqlrc` commands.
- The notation `MSSQL:` is optional only for the `nsrsqlsv` command.
- A single period (.) continues to delimit SQL identifiers.
- The syntax also applies to named instances.
- The backslash period (\.) character sequence replaces each literal period in the SQL identifier.
- The double backslash (\\) character sequence replaces each literal backslash in the SQL identifier.

The following table lists the syntax for file names, file groups, and databases containing a period.

Table 29 Command syntax for names containing a period

Name visible from SQL utilities	Equivalent command-line syntax
SQL database named MyDatabase.COM	MSSQL:MyDatabase\ .COM MSSQL\$MyInstance:MyDatabase\ .COM
SQL file group named MyFileGroup.2 for the SQL database named MyDatabase.COM	MyDatabase\ .COM.MyFileGroup\ .2 MSSQL:MyDatabase\ .COM.MyFileGroup\ .2 MSSQL\$MyInstance:MyDatabase\ .COM.MyFileGroup\ .2
SQL file named MyFile.2, which is a member of the SQL file group named MyFileGroup.2 for the SQL database named MyDatabase.COM	MyDatabase\ .COM.MyFileGroup\ .2.MyFile\ .2 MSSQL:MyDatabase\ .COM.MyFileGroup\ .2.MyFile\ .2 MSSQL\$MyInstance:MyDatabase\ \COM.MyFileGroup\ \2.MyFile\ \2

The following table lists the syntax for file names, file groups, and databases containing a backslash.

Table 30 Command syntax for names containing a backslash

Name visible from SQL utilities	Equivalent command-line syntax
The SQL database named MyDatabase\COM	MyDatabase\ \COM MSSQL:MyDatabase\ \COM MSSQL\$MyInstance:MyDatabase\ \COM
The SQL file group named MyFileGroup\2 for the SQL database named MyDatabase\COM	MyDatabase\ \COM.MyFileGroup\ \2 MSSQL\$MyInstance:MyDatabase\ \COM.MyFileGroup\ \2

The following table lists the syntax for file names, file groups, and databases containing a colon.

Table 31 Command syntax for names containing a colon

Name visible from SQL utilities	Equivalent command-line syntax
SQL database named MyDatabase:COM	MyDatabase:COM MSSQL:MyDatabase:COM MSSQL\$MyInstance:MyDatabase:COM
SQL file group named MyFileGroup:2 for the SQL database named MyDatabase:COM	MyDatabase:COM.MyFileGroup:2 MSSQL:MyDatabase:COM.MyFileGroup:2 MSSQL\$MyInstance:MyDatabase:COM.MyFileGroup:2
SQL file named MyFile, which is a member of the SQL file group named MyFileGroup:2 for the SQL database named MyDatabase:COM	MSSQL\$MyInstance:Mydatabase:com.MyFileGroup:2.MyFile

Example recovery command syntax

Example syntax to restore federated backups of a SQL Always On Availability Group

In the following example commands, the source backup object name, SQ12INST4 is the SQL server instance name used during the backup.

- The following command restores the federated backup to the same database:

```
nsrsqlrc.exe -s nwsrvr -c sql2012clus3.brsvlab.local -f -S normal -
$ "MSSQL$SQ12INST4:" -d "MSSQL$SQ12INST4:database1" "MSSQL
$SQ12INST4:database1"
```
- The following command restores the federated backup to a different instance and database (redirected restore)

```
nsrsqlrc.exe -s nwsrvr -c sql2012clus3.brsvlab.local -C"
'testdb1_Data'='E:\sql2012_data\database1.mdf', 'testdb1_Log'='F:
\sql2012_log\database1_log.LDF'" -f -S normal -$ "MSSQL$SQ12INST4:"
-d "MSSQL$SQ12INST5:copy-database1" "MSSQL$SQ12INST4:database1"
```
- The following command restores multiple databases within an Always On Availability Group:

```
nsrsqlrc.exe -s nwsrvr -c sql2012clus3.brsvlab.local -f -S normal -
$ "MSSQL $SQ12INST4:database1" "MSSQL$SQ12INST4:database2"
```

Example syntax to perform a piecemeal recovery in an active mirror session

In an active mirror session, the NetWorker User for SQL Server GUI prevents a piecemeal recovery of the principal database to a different location. However, you can perform a piecemeal recovery from the command prompt.

The following lines provide sample commands for a piecemeal recovery of the primary file group (mdf & ldf) and file group "a" (ndf) of AcmeBank to the new database AcmeOnline:

```
nsrsqlrc -s "bv-v-cgd2.belred.legato.com" -c "bv-v-
cgd2.belred.legato.com" -$ "MSSQL$THREE:" -R "'PRIMARY','a'" -d "MSSQL
$THREE:AcmeOnline" -C "'AcmeBank'='E:\Data\AcmeOnline.mdf',
'AcmeBank_log'='E:\Data\AcmeOnline_log.ldf', 'AcmeBank1'='E:\Data
\AcmeOnline1.ndf'" -t "Wed Sep 14 13:31:46 2005" "MSSQL
$THREE:AcmeBank"
```

Note

The AcmeOnline database name and file locations are different from AcmeBank.

Performing a recovery from the NetWorker User for SQL Server GUI

The **Restore** window displays the database objects that are available for recovery.

Based on the selected recovery type, the **Restore** window restricts the marking of database objects:

- For the normal restore type, the marking of database objects is not restricted.
- For the copy restore type, only database objects can be marked.

The amount of time that a database takes to recover depends on the following variables:

- Amount of data
- Network traffic
- Server load
- Backup volume availability
- Tape positioning

After recovery is complete, the completion time appears in the **Restore Status** window.

Note the following considerations before you start a recovery operation from the NetWorker User for SQL Server GUI:

- If the backup volume with the databases is loaded at a storage node (backup device) local to the NetWorker server, the restore continues. If the restore does not begin, it is possible that either the wrong volume or no volume is mounted in the backup device.
- In the **Files** tab, you must select **Overwrite the Existing Database** in the following scenarios:
 - When you restore an incompatible database by using the name of an existing database.
 - When you restore from a media failure where one or more database files were lost.

Setting the browse time

You can specify the browse time through the NetWorker User for SQL Server **Restore** window or the command prompt.

Specifying the browse time through the NetWorker User for SQL Server Restore window

In the NetWorker User for SQL Server Restore window, you can browse the online index and mark a database backup version that you want to restore.

The browse time controls which backup data is viewable in the **Restore** window. You can modify the browse time to display backup versions for a different restore time by selecting **Change Browse Time** from the **View** menu. However, the browse time cannot change to a point:

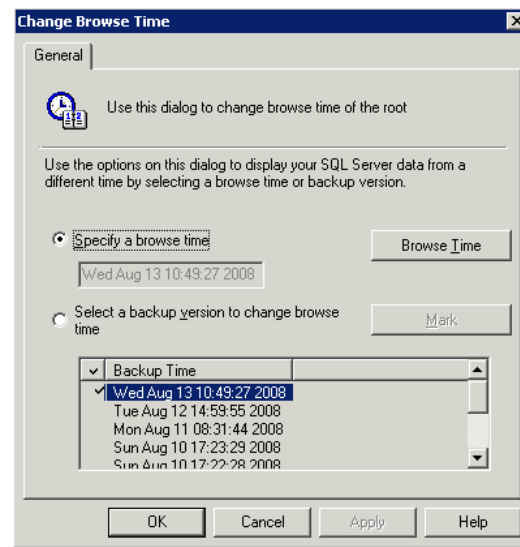
- Earlier than the first backup.
- Later than the most recent backup.

An invalid time entry results in an error message.

Procedure

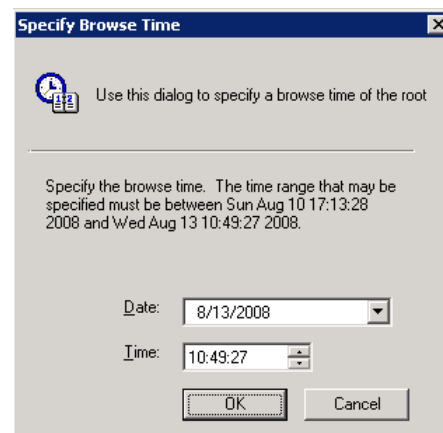
1. From the **View** menu, select **Change Browse Time**.

The **Change Browse Time** dialog box appears, as shown in the following figure.

Figure 37 Change Browse Time dialog box

2. In the **Change Browse Time** dialog box, by default, the **Specify a browse time**, or end time, is set to the current time and the start time is set to 14 days prior. Select this field to enable the **Browse Time** button.
3. To change the end time of the backups that are listed for browsing, click **Browse Time**.

The **Specify Browse Time** dialog box appears, as shown in the following figure.

Figure 38 Specify Browse Time dialog box

4. Select a date and time from the **Date** and **Time** lists.
5. Click **OK**.

The dialog box closes and the **Specify a browse time** field in the **Change Browse Time** dialog box appears with the new date and time. The backup version list is also updated. The start time is set to 14 days before the browse time that you select.

6. By default, the **Select a backup version to change browse time** field marks the current browse time in the backup version list. Select this field to enable the backup version list and **Mark** button. The backup version list displays the known backup version times, which correspond to save times of the root browse directory. The root browse directory is created whenever a SQL Server backup is performed.

7. To select a new browse time, do one of the following:
 - Double-click a backup time in the list, and then click **OK**.
 - Select a backup time in the list, click **Mark**, and then click **OK**.
The **Specify a Browse Time** field is updated to reflect the new setting.
When you change the browse time, previous selections in the **Restore** window are discarded. All file, file group, or database property settings are lost.

Specifying the browse time through the command prompt

After you set the end time with the **Select Time** button, you can modify the range of browsing backup versions with the `nwmssql.exe -b <n_days>` command from the command prompt, where the start time is *n* days prior to the end time set with the **Select Time** button.

For example, if `nwmssql.exe -b 30` is used and the end time in the GUI is set to "6/10/2015 16:50:00", all the backups up to 30 days from the end time are displayed on the **Change Browse Time** dialog box.

Setting up the recovery

Complete the following steps to set up the recovery.

Note

While you set up a recovery in the NetWorker User for SQL Server GUI, you can restrict the number of backups that are displayed on the **Change Browse Time** page, **Database Properties** page, and the **Filegroup Properties** page. If you apply this restriction, the selected browsing list loads faster.

Procedure

1. Open the NetWorker User for SQL Server GUI.
If multiple SQL Server instances are active on the computer, the **Select SQL Instance** dialog box opens before the main window.
2. Select the SQL Server for recovery, then click **OK**.
3. Click the **Select NetWorker Server** button on the toolbar to select a NetWorker server other than the default server.
The **Change Server** dialog box opens.
4. Select a NetWorker server from the list, then click **OK**.
5. In the main window of the NetWorker User for SQL Server interface, perform one of the following steps to open the **Restore Operation** dialog box:
 - From the toolbar, click **Restore**.
 - From the **Operation** menu, click **Restore**, and then click **Normal** or **Copy** (if running SQL Server Enterprise Edition).
 The **Restore Operation** dialog box appears.

Configuring a normal recovery


Procedure

1. From the **Restore Operation** dialog box, select **Normal**, and then click **Continue**.

The **Restore** window appears and lists the databases that can be restored.

The following window shows the **Restore** window.

Figure 39 Restore window

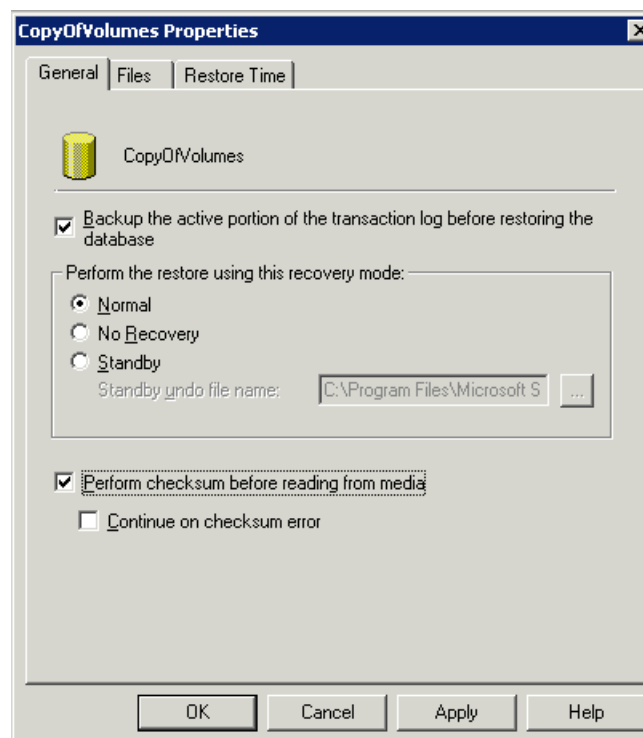


2. On the **Restore** window, select one or more databases to recover.
3. Right-click each marked database and select **Properties** to configure additional settings.

The **Properties** dialog box appears. Available options and actions depend on the version of SQL Server that is to be restored and the selected recovery type.

The following window shows the **Properties** dialog box.

Figure 40 Properties dialog box—General tab



4. Select **Backup the active portion of the transaction log before restoring the database**.

This option is selected by default for a Normal restore, and NMM tries the transaction log backup by using the NO_TRUNCATE and NORECOVERY SQL keywords. The recovery operation continues regardless of whether the transaction log backup succeeds or fails.

There is no log backup in Simple mode. If the option to back up the log is selected, a full backup of the log is performed on the database that is to be recovered.

The most common reason for recovering databases is to restore from operator errors, which are recorded in the transaction log. If you recover the database

without applying the transaction log, you lose the information since the last backup. If you apply the entire transaction log you recorrup the database. Point-in-time recovery data can be recovered to the time of the error minus approximately 1 second, assuming that the database is functional enough to complete the final transaction log backup. If **Specify a Restore Time** on the **Restore Time** tab is selected, the backup continues, but the latest transactions that are captured in the active transaction log backup are not restored.

5. In the **Perform the restore using this recovery mode** area, select one of the following for recovery:
 - **Normal mode**—Instruct the SQL Server to leave the database in operational state after the restore. This is the default mode.
 - **No Recovery mode**—Activates the SQL Server NORECOVERY database restore option for the last stage restored. This mode places the database in a state that cannot be loaded after the restore. However, the database can still process additional transaction log restore operations.
 - **Standby mode**—Enables the Standby Undo File Name, which specifies an undo file for SQL Server to use when rolling back the transactions.

By default, the **Standby mode** field displays a default file name and path:

```
%DriveLetter:\default_path\default_dbNameundo.ldf
```

where:

default_path is the default SQL Server backup path obtained from the SQL Server registry.

default_dbName is the name of the database backup that is selected for the restore.

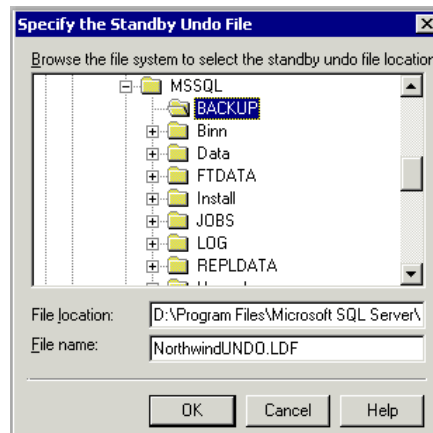
Specify another name and path for the standby undo file through the following steps:

- a. Type a valid name and path in the **Standby mode** field, or click the ellipses button.

The **Specify the Standby Undo File** dialog box appears.

The following window shows the **Specify the Standby Undo File** dialog box.

Figure 41 Specify the Standby Undo File dialog box

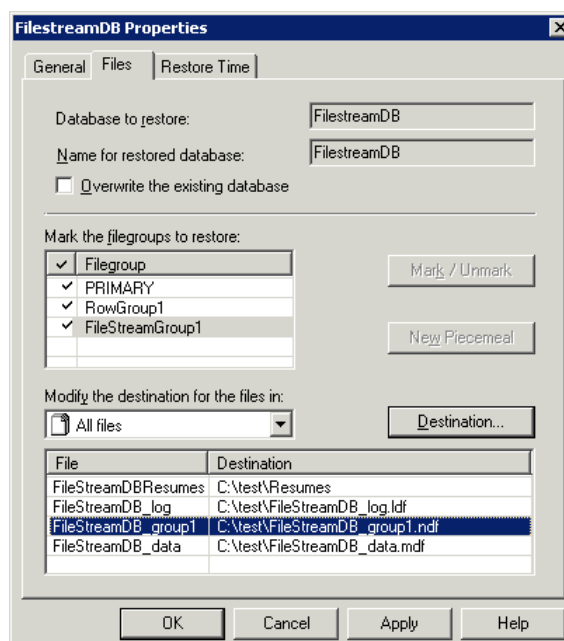


- b. In the **File Location** text box, type a path or browse the file system tree and highlight a file.

- c. In the **File Name** text box, type a file name or browse the file system tree and highlight an existing file.
 - d. Click **OK**.
6. (Optional) In the **Properties** dialog box, select **Perform checksum before reading from media** option to perform a checksum operation before a restore to ensure that it matches the backup.
- NMM verifies the checksum by calculating a local result and comparing it with the stored value. If the values do not match, you can choose to continue the restore operation by selecting **Continue on checksum error**.
7. Select the **Files** tab.

The following window shows the **Properties** dialog box open to the **Files** tab.

Figure 42 Properties dialog box—Files tab



Consider the following:

- The **Database to restore** field displays the database that is selected for the restore operation and cannot be modified.
 - The **Name for restored database** field specifies the name for the restored database and cannot be modified.
8. (Optional) Select **Overwrite the existing database**. The SQL Server creates the specified database and its related files, even if another database with the same name exists. In such a case, the existing database is deleted.
- The **Overwrite the existing database** operation includes the WITH REPLACE SQL keyword in the restore sequence, which restores files over existing files of the same name and location.
9. (Optional) If you are performing a normal, copy, or verify restore operation, the **Mark the filegroups to restore** area displays the file groups of the selected database. These file groups cannot be changed. If you are performing any other type of restore operation, select the file groups and files that you want to restore.

The selected set of file groups is copied to the **Modify the destination for the files in** list, which displays the options for restoring the database files. You can

filter the files that are visible in the list. The following table describes the options in the **Mark the filegroups to restore** list:

The **Modify the destination for the files in** table displays the SQL Server files, which are associated with the database that is to be restored, with their locations. When you perform a normal restore, this table displays the current name and destination based on the SQL Server physical file name and logical location for the restored file. Filestream data is displayed as a folder with no subordinate objects.

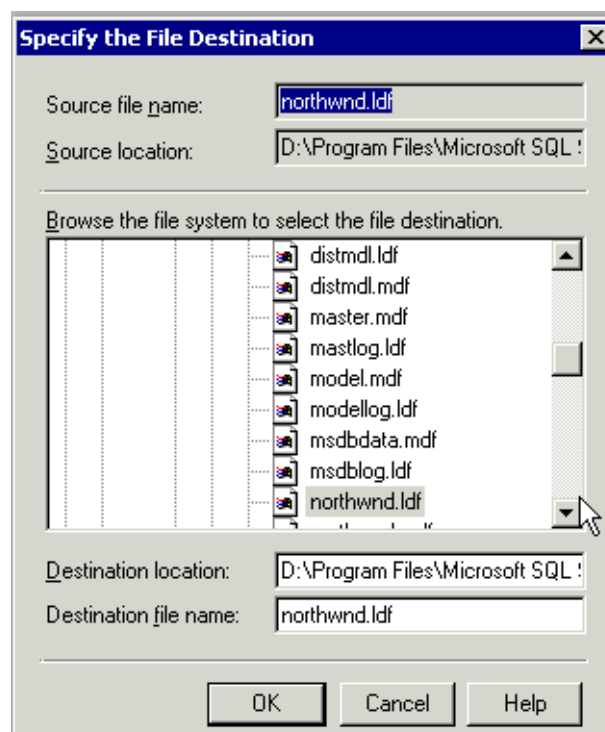
Table 32 Modify the destination for the files in list options

Views	Description
All files	All the files for the database, including transaction log files
All log files	Only the transaction log files
All datafiles	Only datafiles
File group name	Only datafiles for a specific file group
Drive letter	All files that are on a particular drive at the time the backup occurred, even if those files have since been relocated to a different drive

10. (Optional) To modify the destination, perform one of the following actions:

- Double-click a file in the **Modify the destination for the files in** table to display the **Specify the file destination** dialog box.
- Click a file in the list, and then click **Destination**.
The **Specify the File Destination** dialog box appears, as shown in the following figure.

Figure 43 Specify the File Destination dialog box



Consider the following:

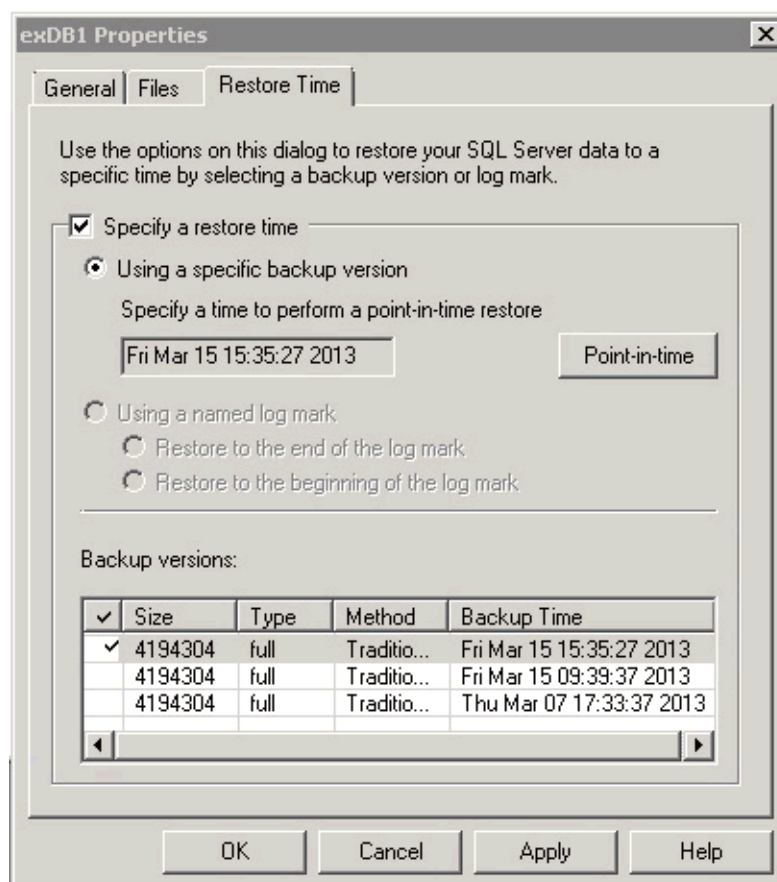
- In the **Source file name** field, the file that is selected in the **File and Destination** lists appears. This text cannot be modified. When multiple files are selected, this text box is empty.
 - In the **Source location** field, the location and the file that is selected in the **File and destination** list appears. This information cannot be modified. When multiple files are selected, the location of the first selected file in the list is displayed.
11. (Optional) To change the destination location, in the **Destination location** field, type a path name or browse the file system tree and highlight a directory or file.

The **Destination location** field displays the file system location for the restored file. When multiple files are selected, the default SQL data path is opened, but it is not selected. Type a pathname, or browse the file system tree and highlight a directory or file to change the location.
 12. (Optional) In the **Destination file name** field, to change the file name, type a new name or browse the file system tree and highlight a file.

The **Destination file name** field displays the name of the file that is selected in the **File and Destination** table. When multiple files are selected, the field is empty
 13. To return to the **Files** tab, click **OK**.
 14. Click the **Restore Time** tab.

The following window shows the **Properties** dialog box open to the **Restore Time** tab.

Figure 44 Properties dialog box—Restore Time tab



The **Restore Time** tab enables you to select a backup version and modify the restore date and time. The default selection for the restore is listed in the **Backup versions** table. When a point-in-time restore is performed, the restore procedure reinstates only transactions from the backup version that occurred before the specified restore date and time.

You can change the backup version or transaction time. The point-in-time for a restore can be changed only for logs-only backups. Unless the last backup time is specifically selected by the user, the active log backup option is set by default.

15. Select **Specify a restore time** to schedule the recovery.

On the **General** tab, if you selected **Backup the active portion of the transaction log before restoring the database** and select **Specify a restore time**, but do not specify the point-in-time in the transaction log, the latest transactions that are captured in the active transaction log backup are not restored.

The latest transactions that are captured in the active transaction log backup are restored to the specified point-in-time only if it is specified in the transaction log.

16. (Optional) To specify a time to perform a point-in-time restore, click **Point-in-Time**.

The restore time indicates the data from which the marked backup version is reinstated during the restore and when to stop restoring transactions. The **Point-in-Time** button is available only when the logs-only backup, or the latest backup (regardless of whether it is full, level logs-only or level cumulative incremental) is marked in the **Backup Versions** table. If a full or differential backup is marked but is not the latest backup, an error message appears.

Specify a Time to perform a point-in-time restore displays the date and time for the backup that is marked. Only one backup version can be selected and marked.

17. To perform the restore, select **By using a named log mark**. This field is enabled only when log marks exist for the selected database backup.

When you select **By using a named log mark**, **Restore to the End of the Log Mark** and **Restore to the Beginning of the Log Mark** are enabled.

18. Select one of the following options:

- Select **Restore to the End of the Log Mark** to restore the backup and stop immediately after the named log mark. This type of restore includes the named transaction in the restore.
- Select **Restore to the Beginning of the Log Mark** option to restore the backup and stop immediately before the named log mark. This type of restore excludes the named transaction.

19. Use the **Log Mark** table field to specify a log mark to use for the restore. Double-click a log mark or select a log mark in the list and click the **Mark** button.

20. When you are finished making selections, click **OK**.

21. To define additional restore options, from the **Options** menu, select **Restore Options**.

The **Restore Options** dialog box appears.

22. On the **General** tab, select the following options as required.

The following table describes the options that are available in the **Restore Options** dialog box

Table 33 Fields on the General tab of the Restore Options dialog box

Field name	Describe
Automatically overwrite any database marked for restore	Select this field to overwrite the marked databases. This action restores the database by using the WITH_OVERWRITE SQL keyword.
Detect available tape sessions prior to restoring a striped backup	The support for this field is deprecated and this field will be removed from NMM in future releases. In NMM 9.1 and later, this field should always be selected if any backups are being restored from tape type devices. Selecting this field does not affect the performance of restore operations from disk type devices.
Select a debug level	<p>Selecting a troubleshooting level issues the command line -D n option, where n is an integer value between 1 and 9. When the troubleshooting level is selected, the troubleshooting information appears in the Backup or Restore status window, and also in the log files. This information can be useful in diagnosing problems.</p> <p>If you specify troubleshooting level logging, watch the size of the NMM log files. The files are at:</p> <pre>installpath\nsr\applogs</pre> <p>where <i>installpath</i> is typically C:\Program Files\EMC NetWorker and can grow very large. Log files with troubleshoot output can be deleted (or archived).</p>
Disable log mark display	<p>Select this option to disable the potentially time-consuming retrieval of log mark information. This option can be selected when log marks will not be used for most restore operations, particularly if the client file index is large. This setting persists across multiple invocations of the user interface. When log mark display is disabled, the log mark in the backup will not be displayed in the Restore Time tab.</p> <p>When this option is selected, the log mark information is not displayed in the Restore Time tab of the Restore Properties dialog box. When it is not selected, the information is displayed.</p>
Transaction log backup before restore	<p>Select this option to perform a transaction log backup before the restore operation starts. With supported SQL Servers, transaction log backups are required before restoring databases. Transaction log backups ensure that the database can be restored to the most recent point-in-time possible. This option is selected by default.</p> <p>When the Transaction log backup before restore option is not selected, the Back Up the Active Portion of the Transaction Log Before Restoring the Database option (selected by default) is enabled in the Properties dialog box.</p> <p>You can choose not to back up the transaction log by selecting No for Restore with overwrite in the precautionary message box. This option replaces the existing database by using the T-SQL command with the WITHREPLACE flag. Selecting No for Restore with overwrite improves performance in cases where the database is too damaged to perform the transaction log backup or it is known that the transaction log is not of interest.</p>
Alternate decryption pass phrase	Select this option to specify a pass phrase for AES encryption other than the default or current phrase used by the NetWorker server. If data was backed up using AES encryption and an alternate pass phrase, (and the original pass phrase is no longer current), that phrase must be specified to decrypt data during the restore. If multiple

Table 33 Fields on the General tab of the Restore Options dialog box (continued)

Field name	Describe
	pass phrases are required, they must be specified at the command line. The <i>NetWorker Administration Guide</i> provides complete information about AES encryption, and setting the pass phrase.
Advanced options	<p>In the Advanced options field, type the required advanced recovery options. The following options are available:</p> <ul style="list-style-type: none"> • BUFFERCOUNT=number_of_IO_buffers This option specifies the total number of IO buffers that can used during a recovery operation. • KEEP_CDC This option enables change data capture (CDC) recovery. • KEEP_REPLICATION This option preserves the replication when it recovers a published database. This option is required if a database was replicated when the backup was created. <p>If you are specifying multiple advanced options each argument must be separated by a comma.</p>

23. Click **OK**.

24. In the toolbar of the **Restore** window, click the play button to start the restore.

After you finish

Monitor the status of the restore in the **Restore Status** window.


Configuring a copy restore

When you perform a copy restore of databases in an SQL instance to a folder in a new location, ensure that you add **Everyone** to the folder security settings, and provide all permissions to the folder. You must complete these actions for all users who have logged in with or without Domain Admin credentials when they perform copy restore operations.

Procedure

1. In the **Restore Operation** dialog box, select **Copy**, and then click **Continue**.

The **Select the SQL Server** dialog box appears, as shown in the following figure.


Figure 45 Select the SQL Server dialog box

2. Select the appropriate source for the backup:
 - To use the database backup from the SQL Server that is currently connected to the NMM client, select **The current SQL Server Instance** . The name of the current NetWorker client host appears.
 - To choose a different SQL Server host from the existing database backups, select **A NetWorker SQL Server client**.
3. Click **Continue**.
4. Specify the destination host and an appropriate user account or group in the list of NetWorker administrators on the NetWorker server (or remote access) to select other client hosts as the source of a copy restore.
 If either of the following situations apply to the selected NetWorker client, a new dialog box appears:
 - A backup exists.
 - A named instance directory in the index.
 This dialog box enables you to select the source from default and named instances.
5. Select either the SQL Server default instance or an SQL Server named instance, and then click **Continue**. After choosing a SQL instance, all database backups are listed.

A copyof prefix is added in front of the database (and database file) names only if a database by the same name exists on the target system. This is not the case for copy restore between systems or between instances of SQL Server (or if the original database has been removed). In these cases, the default database and file names are the same as the original names.

If backups for the SQL default or other named instance were not performed, or the backups are no longer available, an error message appears.


The following figure shows an example of the error message that appears.

Figure 46 NetWorker User for SQL Server dialog box

6. Select a SQL instance.

The **Restore** window lists the databases that can be restored.

The following figure shows the **Restore** window.

Figure 47 Restore window

7. From the left pane, select one database to recover.

For copy restore, you can mark only one database object. Files and file groups of the selected database are not visible in the browse window because they require a database backup. The files and file groups can be viewed and configured in the **Files** tab of the **Properties** dialog box.

8. Right-click the marked database and select **Properties**.

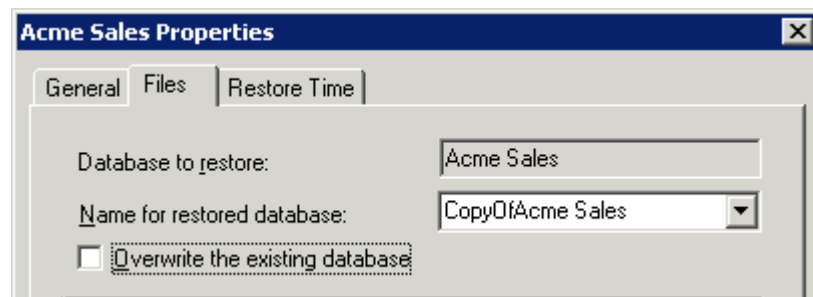
Note

For copy restore, specify the restore properties before starting the restore. Otherwise, the restore fails to start.

The **Properties** dialog box with the **General** tab appears.

9. In the **Properties** dialog box, click the **Files** tab.

The following figure shows the **Properties** dialog box open to the **Files** tab.

Figure 48 Properties dialog box

When you restore a copy of a database, you can overwrite an existing database or create one:

- To create a database, in the **Name for restored database** field, type a name. The default name is *CopyOfSelectedDB*.
- If you want to overwrite an existing database, from the **Name for restored database** field, select a name, and then select **Overwrite the existing database**.

Check the generated file names. After changing the name of the database from *CopyOfAcme Sales* back to *Acme Sales* by using overwrite, remove the *CopyOf* prefix from the associated file names. Default names are generated when the dialog box is first displayed. Any database name edits are not propagated to the database file names. Make these changes manually.

10. In the **Properties** dialog box, click the **Restore Time** tab.

11. After configuring the properties, click **OK**.

These settings are maintained until the restore operation starts.

12. In the toolbar of the **Restore** dialog box, click the play button to start the restore.

After you finish

On the **Restore Status** page, monitor the status of the restore operation.

Configuring a named instance recovery


If a default instance and a named instance are available on a client computer, after a backup is performed with NMM, a copy restore is available for all two instances.

If a default instance is uninstalled, a copy restore from the default instance is available because the instance is still listed in the client file index for the most recently named instance directory. If a more recent backup is performed of any of the instances after the default instance is uninstalled, a recovery will not be possible. A new named instance directory is created that excludes the default instance.

The solution is to reinstall the removed named instance and run another backup. This action restores the directory information and provides complete functionality for the default instance.

The following figure shows the **Restore** window.

Figure 49 Restore window



Perform the following steps to view and select restore items:

Procedure

1. In the left pane of the **Restore** window, expand any item to view a list of available data items.

The descendants of the item are displayed in the right pane.

2. In the left pane, select the items that you want to restore.

If the copy restore type is selected, only one database object can be marked. Files and file groups of the selected database are not visible in the browse window, and require a database backup. The files and file groups are available in the **Properties** dialog box.

3. After the restore is set up, start the operation. [Starting the recovery](#) on page 163 provides more information.

Configuring a verify-only restore

The verify-only option verifies that a database can be restored and that the media that holds the data is available.

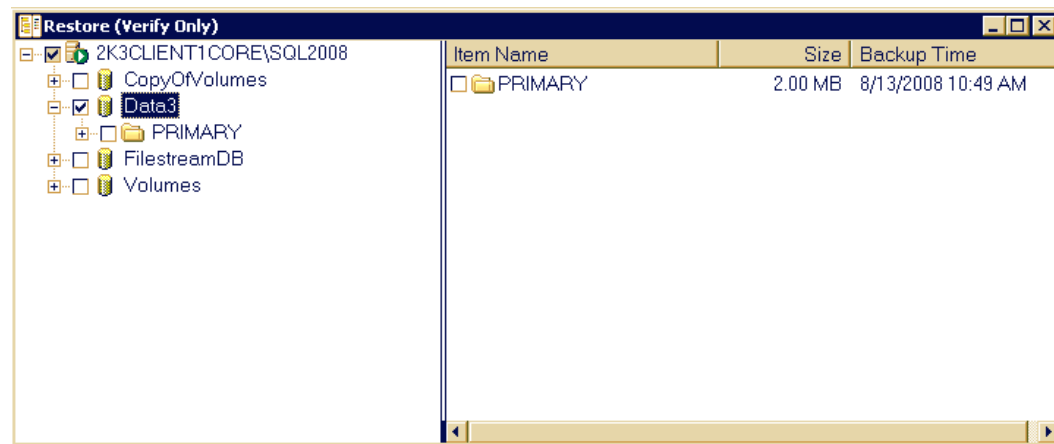
Procedure

1. From the **Restore Operation** dialog box, select **Verify Only**, and then click **Continue**.

The **Restore** window lists the databases that can be restored.

The following figure shows the **Restore (Verify Only)** window.

Figure 50 Restore (Verify Only) window



2. In the left pane, select one or more databases from the list.
3. (Optional) To specify additional options in the **Restore Options** dialog box, right-click the parent server.
4. Click **OK**.
5. In the toolbar of the **Restore** dialog box, click the play button to start the restore operation.

You can monitor the status of the restore in the **Restore Status** window.

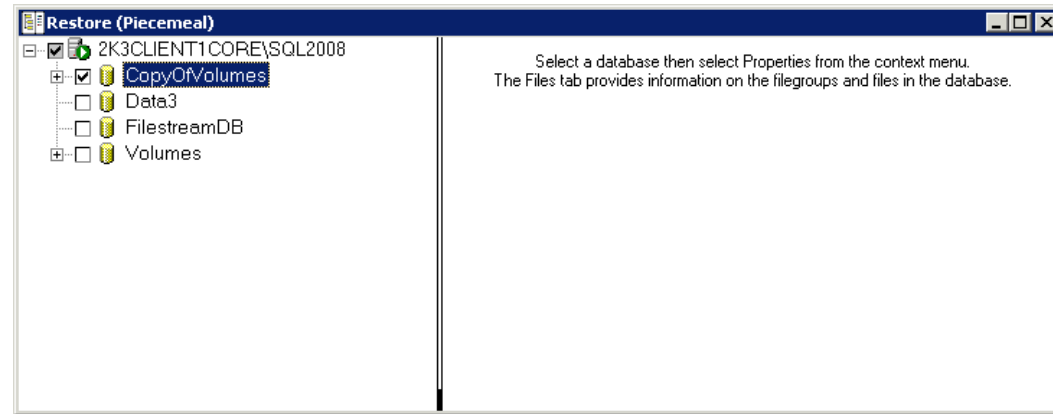
Configuring a piecemeal restore

Select the piecemeal restore option to restore SQL Server databases (Enterprise Edition). [Types of supported recovery](#) provides detailed descriptions.

Procedure

1. In the **Restore Operation** dialog box, select **Piecemeal**, and then click **Continue**.

The **Restore (Piecemeal)** page appears, as shown in the following figure.

Figure 51 Restore (Piecemeal) window

2. In the left pane, select a database from the list.

Only one database object can be marked for a piecemeal restore operation. Files and file groups of the selected database are not visible in the browse window. These operations require a database backup. The files and file groups can be viewed and configured in the **Properties** dialog box, in the **Files** tab.

3. (Optional) To specify additional options in the **Restore Options** dialog box, right-click the parent server.
4. Click **OK**.
5. Right-click the database and select **Properties**.

In piecemeal restore operations, you must specify the restore properties before starting the restore. Otherwise, the restore fails to start.

The **Properties** dialog box appears.

6. On the **General** tab, select a name from the **Name for restored database** field, and then select **Overwrite the existing database**.

Piecemeal restore operations overwrite a database when they restore to the same location.

7. In the **Properties** dialog box, click the **Restore Time** tab.
8. When you are finished configuring properties, click **OK**.

For piecemeal restore, if you redisplay the **Properties** dialog box before starting the restore operation, the following settings in the **Properties** dialog box are reset:

- Marked files and file groups of the selected database
- **Name for restored database** option (on the **Files** tab)
- **Backup the active portion of the transaction log before restoring the database** checkbox (on the **General** tab)

If you return to the **Properties** dialog box, a message appears.

The following figure shows an example of the message that appears.

Figure 52 Restore Options dialog box

9. To start the restore operation, in the toolbar of the **Restore** dialog box, click the play button.

Results

Piecemeal restore is iterative. You can continue to restore additional file groups in subsequent operations. Previously restored file groups are not available for selection unless you specify **New Piecemeal**.

Viewing the required volumes

After you select the databases to restore, you can view the backup volumes that contain the databases.

To check the backup volumes that contain the databases, do one of the following:

- Click **View > Required Volumes**
- From the toolbar, select **Required Volumes**.

Ensure that the necessary backup volumes are mounted in the NetWorker server's storage device.

If a required backup volume is not mounted, do one of the following:

- Contact the NetWorker administrator and request that the volume be mounted.
- Wait for the NetWorker software to prompt the administrator to mount the volume.

Setting the restore properties

You must complete the required steps to set the restore properties for a file, file group, or database.

Procedure

1. Open the NetWorker User for SQL Server GUI.
2. Right-click the database item, and select **Properties**.
3. Complete the **Properties** dialog box.
4. Click **OK**.

Results

All properties are context-sensitive, based on the version of SQL Server that is running and the selected restore type.

Preferences that are specified for these properties are used for the current restore operation. When a preference is not set, the default is used. After the restore completes, and the **Restore** window closes, the property settings revert to the NetWorker User for SQL Server defaults.

Note

If you select the copy restore type, you must specify the restore properties before starting the restore. Otherwise, the restore fails to start.

Setting database restore properties

The information in this section applies to normal and copy restore types.

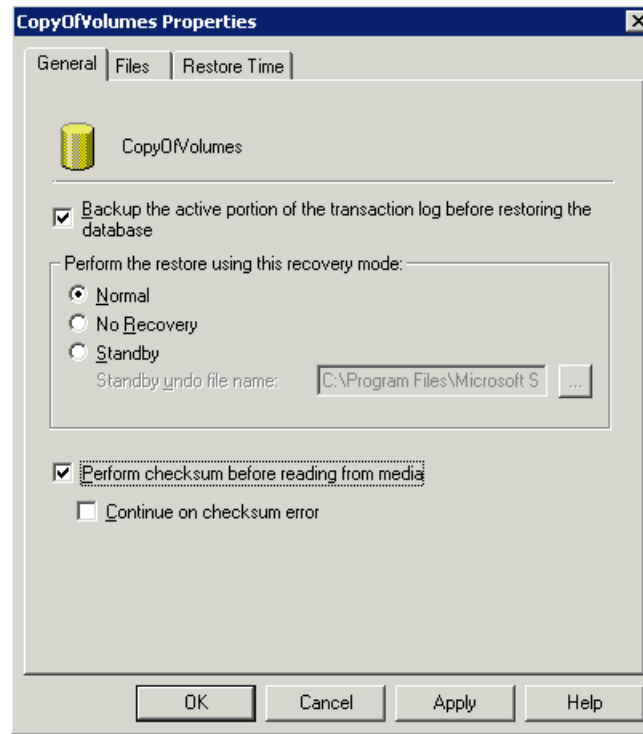
You must complete the required steps to set database restore properties.

Procedure

1. Open the NetWorker User for SQL Server GUI.
2. In the **Restore** window, select a database.
3. Right-click the database and select **Properties**.

The **Properties** dialog box appears, as shown in the following figure.

Figure 53 Properties dialog box



4. On the **General** tab, indicate whether to back up the active portion of the transaction.

When **Backup the active portion of the transaction log before restoring the database** is selected, the transaction log backup starts with the `NORECOVERY` keyword. The restore operation continues regardless of whether the transaction log backup succeeds or fails.

Specify a point-in-time restore within the transaction log backup. This option applies to all SQL Server restores and addresses database failure if one occurs after the last scheduled logs-only backup.

By backing up the active portion of the log before the restore, you can specify a restore for *time of failure*.

5. To specify a point-in-time restore, select the **Restore Time** tab, and make entries as appropriate.

Point-in-time restore can use a full, logs only, or cumulative incremental backup, if the selected backup is a logs backup and/or the latest backup.

6. Specify a recovery mode to use for the restore:

- **Normal mode**—Instructs the SQL Server to leave the database in operational state after the restore. This is the default mode.
- **No Recovery mode**—Activates the SQL Server `NORECOVERY` database restore option for the last stage restored. This mode places the database in a state that cannot be loaded after the restore. However, the database can still process additional transaction log restore operations.
- **Standby mode**—Enables the **Standby Undo File Name** option, which specifies an undo file for SQL Server to use when rolling back the transactions. By default, this field displays a default file name and path as follows:

```
%DriveLetter:\default_path\default_dbNameundo.ldf
```

where *default_path* is the default SQL Server backup path that is obtained from the SQL Server registry, and *default_dbName* is the name of the database backup selected for the restore.

7. Select **Checksum** and **Continue with error**.

If checksums are not calculated for the backup, specifying them for the restore causes the restore operation to fail.

Specifying an undo file name

You must complete the required steps to specify the undo file name.

Procedure

1. In the **Standby Undo File Name** field, type a name
2. Click the ellipses (. . .) button.
3. In the dialog box that appears, select a file from the file system tree.
4. Specify the required values for the following fields:
 - The **File Location** field—In the **File Location** field, type a path or browse the file system tree and highlight a file name to specify the path for the standby undo file.
 - The **The File Name** field—In the **File Name** field, type a file name or browse the file system tree and highlight a file name to specify the name for the standby undo file.
5. Click **OK**.

View file or file group restore properties

The file and file group restore properties in the **Properties** dialog box are provided for informational purposes only. SQL Server dictates the settings for these properties. They cannot be configured.

The properties in the **Properties** dialog box apply to one file or a file group for a normal restore type operation.

Note

These file and file group restore properties are not available for copy restore type operation.

You must complete the required steps to view the restore properties for a file or file group.

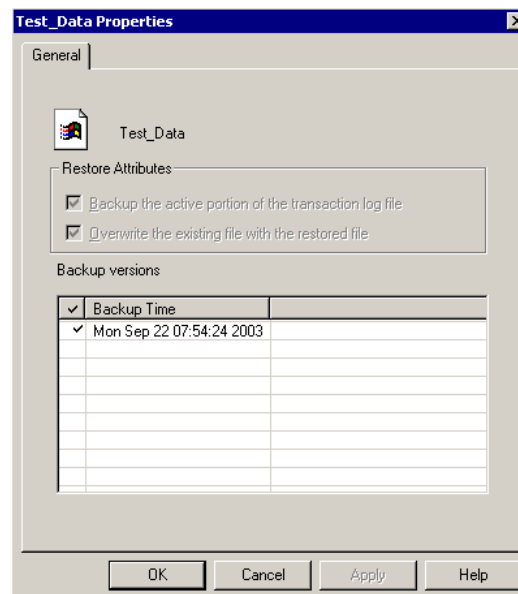
Procedure

1. From either pane of the **Restore** window, select a file or file group, and then right-click.
2. From the list, select **Properties**.

The **Properties** dialog box appears. Properties differ depending on the version of SQL Server.

The following figure shows the **Properties** dialog box.

Figure 54 Properties dialog box



3. Select **Backup the active portion of the transaction log file** to back up the active portion of the transaction log before the restore operation is performed.

The log is applied to the file group or file to ensure it is consistent with the rest of the database. The SQL Server requires the transaction log when it restores damaged or lost datafiles. The NetWorker User for SQL Server tries to perform a transaction log backup for SQL Server non-Enterprise Editions or Primary file group. The backup uses the `NO_TRUNCATE` and `NO_RECOVERY` SQL keywords.

For files that belong to secondary file groups and secondary file groups restore for SQL Server Enterprise Editions, the restore workflow requires that you first restore the file group and then take a backup of the active portion of the transaction log.

If a file or file group is restored with the NetWorker User for SQL Server program, the transaction log backups occur automatically. It is recommended that you use the NetWorker User for SQL Server for this type of restore operation.

4. Select **Overwrite the existing file with the restored file** to force the SQL Server to ignore errors that are caused by nonexistent files, which are a result of media failure. When you select this option, the recovery operation uses the `WITH REPLACE` SQL keyword in the restore sequence. The file or file group is restored to the exact location (drive and pathname) as the location on the source host from which the data was backed up.
5. The **Backup versions** table lists the date and time of all the backups that are available for the restore operation.

Selecting file groups to restore

Use the **Properties** dialog box to select a file group to restore. For normal and copy restore, the tab is labeled **Files**.

You must complete the required steps to select file groups to restore.

Procedure

1. In the **Properties** dialog box, click **Files**.

If the fields in this dialog box are empty, review the file configuration information. [Specifying Read File Configuration properties](#) provides details.

The **Database to restore** field displays the name of the database (on secondary storage) that is selected for the restore operation. This field is informational only.

2. The **Name for restored database** field specifies the name for the restored database:

- If you are performing a normal restore, this field displays the name of the selected database is disabled.
- If you are performing a copy restore, `CopyOf` is appended to the source database name and to all associated datafiles and log files.

In the **Name for restored database** field, type a new name in the field or select a name from the list to specify a different name. The name must comply with SQL Server naming conventions.

If you specify a different name, the data and log files retain the default name, as shown in the **File and Destination** area. For example, if you select copy restore when you want to restore a database named *Project* to a database named *Test*, the data and log file names retain the values of `CopyOfProject_Data.MDF` or `CopyOfProject_Log.LDF`. You must change the data and log file names.

When the **Name for restored database** field is set to the name of an existing database and you click **Apply** or **OK**, the **Overwrite the existing database** field is enabled. For example, when the two databases are incompatible, the name of the existing database is used for the restored one.

3. Select the **Overwrite the existing database** field for the SQL Server to create the specified database and its related files, even if another database exists with the same name. In such a case, the existing database is deleted.

This field causes the `WITH REPLACE` SQL keyword to be included in the restore sequence. The `WITH REPLACE` keyword restores files over existing files of the same name and location.

4. Select or clear the file groups to restore. If you are performing a normal or copy restore, this field displays the file groups of the selected database. The set of

file groups that is marked in this field is copied to the **Modify the destination for the files in list**:

- **Modify the destination for the files in** lists the different views for the database files that will be restored, and it enables the files that are visible in the **File and Destination** table to be filtered.
- The **File** column of the **File and Destination** table lists SQL Server logical file names. The **Destination** column lists physical file name and locations. The files that are listed in this table are associated to the database that is marked for the restore operation.
- If you are performing a normal restore, the **File and Destination** table displays the current name and destination based on the SQL Server physical file name and logical location for the restored file.
- If you are performing a copy restore, the **File and Destination** table displays a default name and destination based on the SQL Server physical file name and logical location for the restored file.
The default location for the datafiles and log files is in the data path of the default SQL Server installation directory. If this directory is on the system drive, provide enough disk space for the database files, or specify another location that has sufficient space.

You cannot edit the **File and Destination** table. You can, however, modify the destination location.

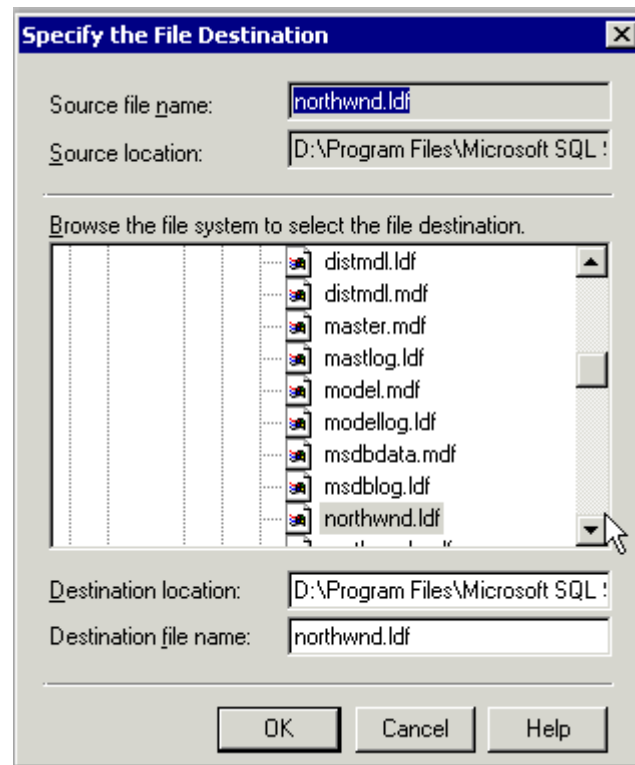
5. To modify the destination, perform do one of the following actions:

- Double-click a file to display the **Specify the file destination** dialog box. Then follow the instructions in the next section.
- Click a file, and then click **Destination** to display the **Specify the file destination** dialog box. Then follow the instructions in the next section.

Specify the restored file's destination and file name

Specify the destination locations for the restored files in the **Specify the File Destination** dialog box.

The following figure shows the **Specify the File Destination** dialog box.

Figure 55 Specify the File Destination dialog box

Note that the following fields provide information and cannot be modified:

- The **Source file name** field displays the file that is currently selected in the File and Destination lists. When multiple files are selected, this field is empty.
- The **Source location** field displays the file system location and the file that is currently selected in the File and Destination lists. When multiple files are selected, this field contains the file system location of the first selected file in the **File and Destination** table.

Modify the following fields as required:

- The **Destination location** field displays the file system location for the restored file. When multiple files are selected, the default SQL data path is open, but it is not selected.
To modify this field, type a pathname or browse the file system tree and highlight a directory or file. When a directory is highlighted, its path appears in the **Destination Location** field. If a file is highlighted, the directory for the highlighted file appears.
- The **Destination file name** field, by default, lists the name of the file that is currently selected in the **File and Destination** table. When multiple files are selected, the field is empty.
In the **Destination file name** field, type a new name or browse the file system tree and highlight a file to modify this field. When a file is highlighted, the file name appears in the **Destination file name** field.

Note

Default file names are generated when the dialog box first appears. Verify that the file names are correct. Verification is particularly important after you change the database name.

Specifying Read File Configuration properties

Some of the data that populates the fields on the **Files** tab of the **Properties** dialog box is obtained from new file-configuration metadata objects that are created in the client file index.

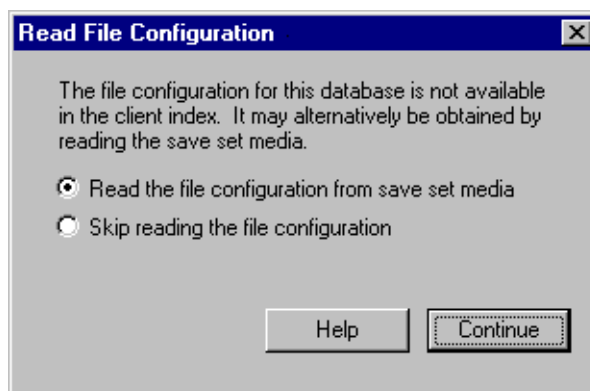
You must complete the required steps to specify Read File Configuration properties.

Procedure

1. Open the **Properties** dialog box for a marked database item that has no file-configuration metadata in the client file index. Load this information from the save set media.

The **Read File Configuration** dialog box appears, as shown in the following figure.

Figure 56 Read File Configuration dialog box



2. Select one of the following buttons:
 - To read the file configuration from the save set media in the background, select **Read the file configuration from save set media**. A message appears. When this process finishes, the **Properties** dialog box appears, and the **File and Destination** table display valid data. To cancel the reading process, click **Cancel**. If you cancel from the **Reading** dialog box, the **Properties** dialog box appears, but the **File and Destination** table is empty.
 - To skip the process of reading the file configuration, select **Skip reading the file configuration**. If you select this button, the **Properties** dialog box appears, but the **File and Destination** table is empty.

Setting Restore Time properties

The Restore Time properties support selecting of a backup version and changing the restore date and time. When a point-in-time restore is performed, the restore procedure reinstates transactions only from the backup version that occurred before the specified restore date and time.

Complete the following steps to set restore time properties:

Procedure

1. In the **Properties** dialog box, select **Restore Time**.
2. On the **General** tab, select **Back up the active portion of the transaction log**, and then **Specify a Restore Time**.

Not the transactions in the transaction log backup exist in the restored database. Only data up to the point-in-time restore operation appear.

You can specify a point-in-time restore for between the last scheduled full, logs-only, or cumulative incremental backup and the current time. The transaction log is backed up as the initial portion of the restore operation and, if the transaction log backup is successful, the log is also restored.

3. Select **Specify a Restore Time** to enable the restore time controls, and then Specify a Restore Time to enable the restore time controls.
4. Select **Using a specific backup version** to enable the **Specify a time to perform a point-in-time restore** field.

The **Specify a time to perform a point-in-time restore** field indicates what data from the marked backup version is reinstated during the restore, and what to stop restoring transactions. If a backup is the most recent or a logs-only backup, this field can be modified for a backup that is selected in the **Backup Versions** table.

5. Click **Point-in-Time** to modify the content in the **Specify a time to perform a point-in-time restore** field.

The **Point-in-Time** button is enabled if:

- A logs-only backup is selected in the **Backup Versions** table.
- The **Backup the active portion of the transaction log before restoring the database** option is selected on the **General** tab for the latest backup (full, cumulative incremental, or logs-only). If the marked backup is not the most recent backup, clicking this button displays an error message.

The **Backup Versions** table displays the following information:

- **Size**—The size of the backup.
- **Type**—The backup type, including “full” for full backups, “inc” for logs-only backups, and “1” for cumulative incremental backups.
- **Method**—The traditional recovery method creates the backup that is to be restored.
- **Backup Time**—The date and time in seconds, when the backup was created.

6. Select **By using a named log mark** to perform the restore. If log marks are not used, you can disable the display of this information.

This field has the following characteristics:

- Can use a maximum of 1024-log marks.
- Is enabled only when log marks exist for the selected database backup. When this field is selected, the **Restore to the End of the Log Mark** and the **Restore to the Beginning of the Log Mark** buttons are enabled.

7. Specify which type of named log mark restores to perform by selecting one of the following:

- To restore the backup and stop it immediately after the named log mark, select **Restore to the End of the Log Mark**. This type of restore includes the named transaction in the restore.
 - To restore the backup and stop it immediately before the named log mark, select **Restore to the Beginning of the Log Mark**. This type of restore excludes the named transaction.
8. Use the **Log Mark table** field to specify the log mark to use for the restore. Only one log mark is selected. The following columns appear in the **Log Mark** table:
- **Log mark**—The name of the transaction log mark.
 - **Date Modified**—The date and time, in milliseconds, on which the named transaction log mark was created.
 - **Description**—Any information about the log mark that was typed when the transaction was performed.
9. To select a log mark, perform one of the following actions:
- Double-click any log mark in the table.
 - Click any log mark in the list, and then click **Mark**.
A check mark appears next to the log mark name. You can mark only one log mark name at a time.

Setting point-in-time restore properties

Specify a date and time for the SQL point-in-time restore in the **Point-in-Time Restore** dialog box.

Ensure that the date and time are within a time frame spanned by the transaction log backup that was marked in the **Backup Versions** table. The create time of the transaction log backup defines the upper time frame limit.

For a normal restore, if a transaction log backup has been specified as part of the restore, the upper limit is the current time. The lower time frame limit is the create time of the first full backup of the database.

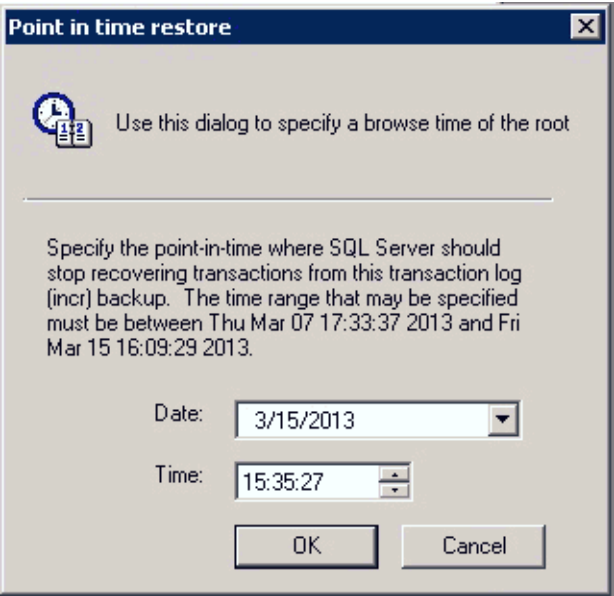
You must complete the required steps to set point-in-time restore properties.

Procedure

1. Click **Restore Time**.
2. Select **Specify a restore time**.
3. Click **Using a specific backup version**.
4. Click **Point-in-time**.

The **Point in time restore** dialog box appears, as shown in the following figure.

Figure 57 Point in time restore dialog box



5. In the **Date** field, specify the restore time date for the marked backup version. The syntax for this text box is *mm/dd/yyyy*.

To modify the date, type a date by using the syntax or click the arrow to display the **Calendar**. In the **Calendar**, click a date. Use the **Previous Month** and **Next Month** buttons to change from the current month.

The restore time date must be within the range of transaction dates in the selected backup version. If a transaction log backup has been specified as part of the normal restore operation then the upper limit is today's date.

6. In the **Time** field, specify the restore time for the marked backup version. The syntax for this text box is *hh:mm:ss*.

To modify the time, type or select a time.

The restore time must be within the range of transaction times in the marked backup version. If a transaction log backup has been specified as part of the restore operation, the upper limit is the current time.

Note

A point-in-time restore is successful only when the times on the NetWorker server and NMM client are synchronized.

Setting advanced restore options

NMM supports advanced restore options in the NetWorker User for SQL Server GUI.

The following table describes the advanced recovery options that NMM supports.

Table 34 Advanced recovery options

Recovery option	Description
<code>BUFFERCOUNT=number_of_IO_buffers</code>	This option specifies the total number of IO buffers that can be used during recovery.

Table 34 Advanced recovery options (continued)

Recovery option	Description
KEEP_CDC	<p>This option enables change data capture (CDC) recovery.</p> <p>When restoring a database with CDC enabled, the recover operation works differently depending on the recovery destination.</p> <p>Use the to KEEP_CDC option to:</p> <ul style="list-style-type: none"> Recover the CDC enabled database on the same SQL instance by overwriting an existing database. In this scenario, KEEP_CDC is optional. Recover the CDC enabled database with a different name on the same SQL instance. In this scenario, KEEP_CDC is required. Recover the CDC enabled database on a different SQL server instance. In this scenario, KEEP_CDC is required.
KEEP_REPLICATION	<p>This option preserves the replication when recovering a published database. This option is required if a database was replicated when the backup was created.</p>

Use the following procedure to set advanced restore options in the NetWorker User for SQL Server GUI.

Procedure

1. On the **Restore** page, select the database that requires advanced recovery options.
2. From the **Options** menu, select **Restore Options**.
The **Restore Options** window appears.
3. In the **Advanced options** field, type the required advanced recovery options.
Follow the syntax described in the Advanced restore options table. Separate multiple entries with a comma. For example:
`BUFFERCOUNT=2, KEEP_CDC, KEEP_REPLICATION`
4. Click **OK**.

Results

The recovery options that are specified in the **Advanced options** field are applied to the marked database.

Starting the recovery

During the restore operation, the **Restore Status** window displays progress messages for each restored database.

If the backup volume containing the databases is loaded at a storage node (backup device) local to the NetWorker server, the restore proceeds. If the restore does not begin, it is possible that either the wrong volume or no volume is mounted in the backup device.

When you restore an incompatible database by using the name of an existing database or when you restore from a media failure where one or more database files were lost, on the **Files** tab, select **Overwrite the Existing Database**.

To display the **Files** tab, right-click the database item and select **Properties**.

After the restore is finished, the restore completion time is listed in the **Restore Status** window.

Note

If you do not have the instant file initialization feature enabled, the following warning message will appear in the output:

```
Please ensure that instant file initialization is enabled on the destination
SQL Server instance for improved restore performance.
```

However, the restore operation will continue. See Microsoft documentation for details on enabling the database instant file initialization feature on the SQL Server.

Performing recovery of federated backups with NetWorker User for SQL Server

Recovering federated backups

Procedure

1. On the primary replica server, break the Always On replication for the Always On database by typing the following command in NMM SQL Server Management Studio (SSMS) plugin GUI query window:

```
ALTER AVAILABILITY GROUP group_name REMOVE DATABASE
database_name
```

2. In the NetWorker User for SQL Server GUI, select **Operation > Select NetWorker SQL Server Client**.

The **Select the SQL Server** dialog box appears.

3. In the **Select the SQL Server** dialog box, select the Windows cluster name from the list, and then click **Continue**.

The **Select the SQL Server Instance** dialog box appears.

4. In the **Select the SQL Server Instance** dialog box, select the instance (default or named) where the availability group databases are created, and then click **Continue**.

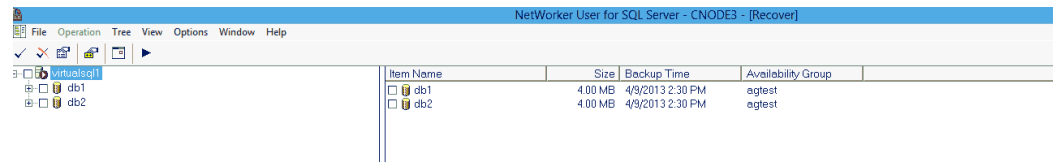
Note

If an availability group resides on multiple SQL Server instances, use the same SQL Server instance name that was specified when you configured the backup.

5. Browse the backed-up availability group databases.

The following figure shows the backed-up availability group databases displayed in the **NetWorker User for SQL Server** page.

Figure 58 Backed-up availability group databases displayed in the NetWorker User for SQL Server page



Item Name	Size	Backup Time	Availability Group
db1	4.00 MB	4/9/2013 2:30 PM	egiest
db2	4.00 MB	4/9/2013 2:30 PM	egiest

6. Select the availability group database and perform a recovery.

In the dialog box that appears, click **Yes** to run a full backup.

Note

A full backup is performed as part of the recovery process when SQL VDI federated recovery is started.

7. On the secondary replica server, to rejoin the Always On database to the Availability Group:
 - a. In the NMM SSMS plugin GUI, locate and expand **Availability Group** to view the Always On database.
 - b. Right-click Always On database and select **Join Availability Group**.
8. Ensure that both primary and secondary copies of Always On database are online.

NOTICE

Do not create a client resource using the Always On Availability Group virtual name when you perform a backup because the backup will fail.

Performing SQL Server 2012 or later recovery of an Always On Availability Group

Perform the steps in this section for SQL Server 2012 or later instance-level and database-level recovery with the NetWorker User for SQL Server GUI.

Note

You can also use the NMM SQL Server Management Studio (SSMS) plugin GUI to perform federated recovery of SQL Server 2012 and later. Refer to the section [Performing a recovery from the NMM SSMS plugin GUI](#) on page 113 for instructions.

For SQL Server 2012 or later instance-level and database-level recovery, ensure to:

1. Perform a full backup of the instance or Always On database.
2. On the primary replica server, break the Always On replication for the Always On database by typing the following command in NMM SSMS plugin GUI T-SQL query window:
`ALTER AVAILABILITY GROUP group_name REMOVE DATABASE database_name`
3. On the secondary replica server, delete the Always On database, which is in restoring mode.

Complete the following steps to perform instance-level recovery.

1. On the primary replica server:

Note

For database-level restores of Always On databases, do not complete step a to recover all the system databases.

- a. Recover all the system databases.
- b. Recover all the user databases or just Always On databases, as required.
2. On the secondary replica server, recover the same Always On database using the **No Recovery** option in Normal Restore under database properties.
3. On the primary replica server, to rejoin the Always On database to the Availability Group.
 - a. Right-click the Availability Group and select **Add Database**.
 - b. In the dialog box that appears, select the database and click **Next**.
 - c. In the dialog box that appears, specify the shared location and click **Next**.
4. Connect to the instance, and then click **Finish**.

Performing SQL Server 2012 or later recovery when an Always On Availability Group is configured on a SQL Server instance or database

This section provides instructions on how to restore an Always On Availability Group that includes both primary and secondary database replicas.

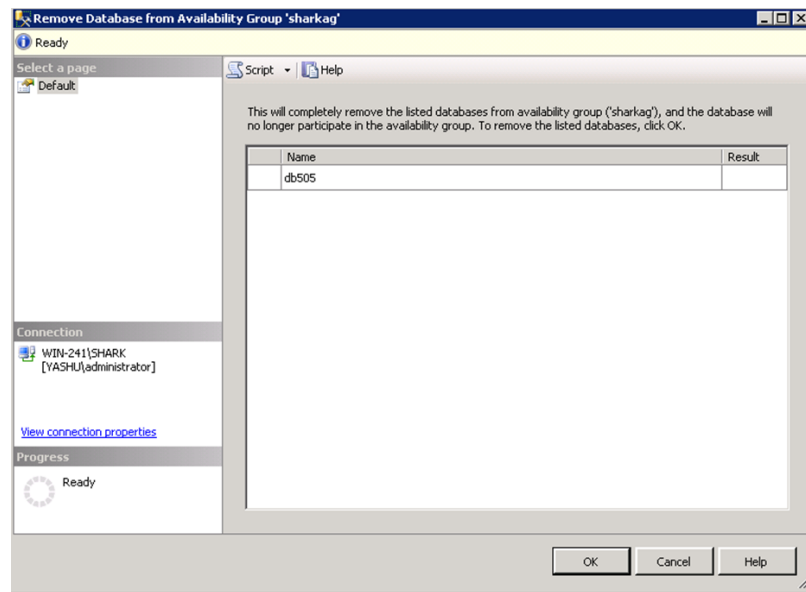
Note

Ensure that you have full and incremental backups of the SQL Server instance or Always On database.

Procedure

1. On the primary replica server, break the Always On replication for the Always On database using the following steps:
 - a. From the SQL Server Management Studio (SSMS) GUI, from the **Object Explorer** pane, expand the primary replica server.
 - b. Select the availability group, and then expand **Availability Databases**.
 - c. From the **Object Explorer** pane, select the Always On database that you want to recover.
 - d. Right-click the selected database, and then click **Remove Database from Availability Group**.

The **Remove Database from Availability Group** dialog box appears, as shown in the following figure.

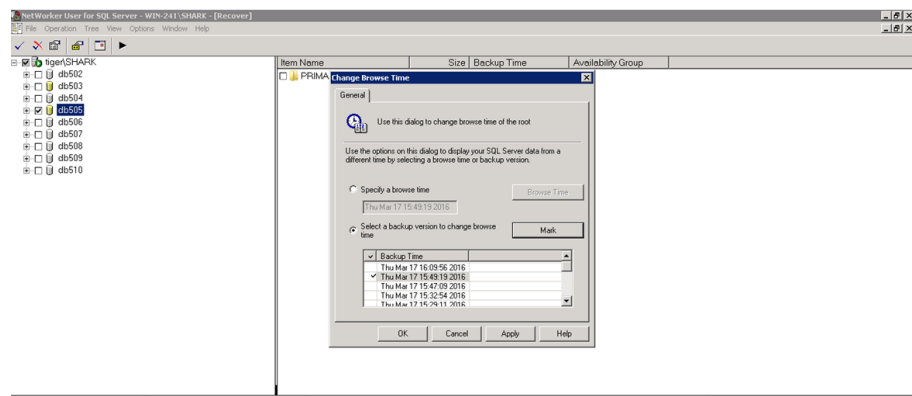
Figure 59 Remove Database from Availability Group dialog box

- e. In the **Remove Database from Availability Group** dialog box, verify that the database you want to recover is listed, and then click **OK** to break the Always On replication.
2. Open the NetWorker User for SQL Server GUI, and then from the **Normal Restore** page, recover the Always On database with default settings.
3. Select the database backup that you want to restore and which browse time you must change.

Note

Ensure that you have selected an incremental backup. This procedure does not work with full backups.

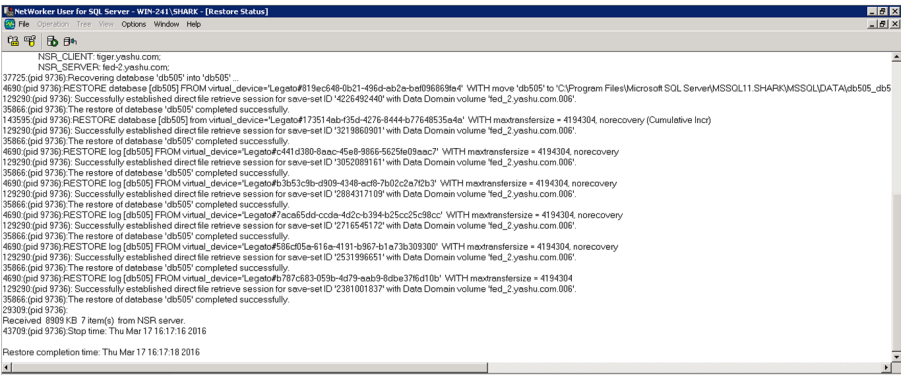
The following figure shows the **Change Browse Time** window.

Figure 60 Change Browse Time window

4. Click **OK**.
The **Restore Option** dialog box appears.
5. Click **Yes**.
The restore operation is completed on the primary server.

The following figure shows the output for a successful restore operation.

Figure 61 Restore Status window



6. Configure recovery of the same Always On database on the secondary replica server. From the **Normal Restore** page, under database properties, ensure that you select **NoRecovery** and **Overwrite the existing database**.

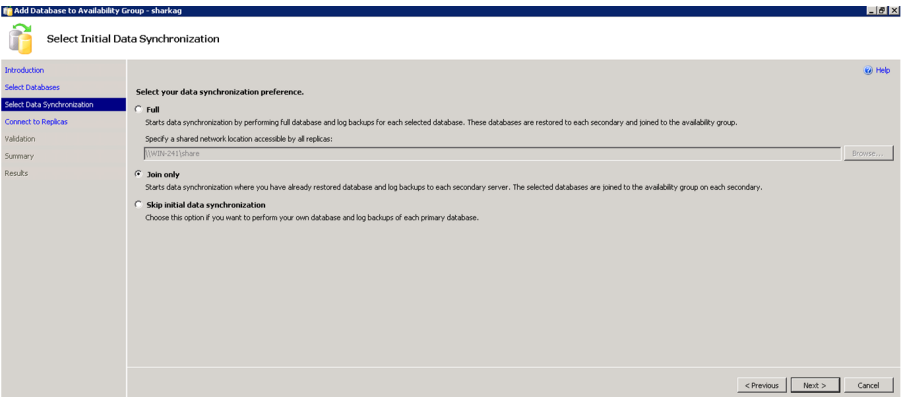
Note

The browse time must be the same as the one you selected for the restore operation on the primary node.

7. From the NMM SSMS plugin GUI on the primary replica server, in the **Add Database to Availability Group** window, select **Join only** to rejoin the Always On database to the Availability Group.

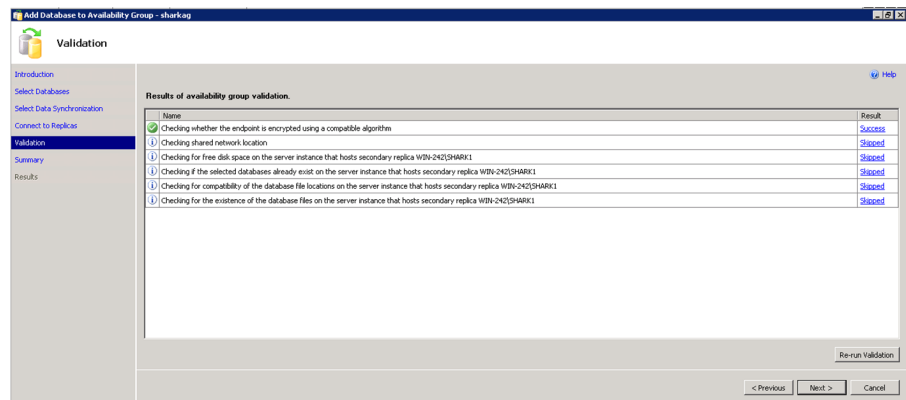
The following figure shows the **Add Database to Availability Group** window.

Figure 62 Add Database to Availability Group window - Select Initial Data Synchronization view



8. Click **Next**.

The **Validation** page appears, as shown in the following figure.

Figure 63 Add Database to Availability Group window - Validation view

9. Click **Next**.
The **Summary** page appears.
10. Click **Finish**.
The **Results** page appears.
11. To join the Always On database to the availability group, from the SSMS GUI on the secondary replica server, in the **Object Explorer** pane, right-click the Always On database, and then click **Join to Availability Group**.

CHAPTER 7

Granular Level Recovery

This chapter includes the following sections:

• Overview	172
• Considerations	172
• Performing granular-level recovery	177
• Dismounting backups after performing granular-level recovery	180

Overview

NMM supports granular-level recovery (GLR) for backups of SQL Server created with NMM and SQL VDI.

GLR enables you to recover specific items from a single backup without having to restore the entire backup. By using GLR with SQL Server, you can recover individual tables to the production database. This feature reduces the space requirements on local system storage during a restore operation. Depending on the size of the content database, GLR may also reduce recovery time.

NMM SQL GLR requires GLR enabled backups. NMM SQL GLR enabled backups are enabled by default for NMM 9.1 and later. GLR enabled backups perform additional inline processing during a backup operation, and store additional metadata on the save set backup records. GLR enabled backups are not supported for direct backups to tape devices.

NMM SQL GLR is managed from the Table Restore tab in the NMM SSMS plug-in GUI. The Table Restore tab allows you to select the database and backup version for GLR. The Table Restore tab constructs a list of save sets that are contained in the selected backup version. NMM uses a mount service and the NetWorker Virtual File System (NWFS) to mount all save sets in the selected SQL backup version. The NWFS exposes the save sets as SQL native backup files on a virtual file system on the NMM client.

After the Table Restore tab mounts all save sets as a virtual file system, ItemPoint™ for Microsoft SQL Server launches. ItemPoint reads the NWFS and presents the database and tables that are contained in the backup, enabling you to restore SQL tables to the production database. Once GLR is performed through ItemPoint and the program is exited, the NMM SSMS plugin GUI unmounts the backup.

The NMM SSMS plugin GUI mounts SQL backups as NWFS by using a mount service. The mount service enables lifecycle management of the mount. While a backup is mounted, the mount service system tray icon appears. You can use the tray icon to inspect the mount path, extend the mount time, and unmount the backup.

Considerations

ItemPoint environmental requirements

Before using ItemPoint for Microsoft SQL Server, consider the following environmental requirements:

- ItemPoint supports the 64-bit versions of the following operating systems, both physical and virtual:
 - Windows Server 2008
 - Windows Server 2008 R2
 - Windows Server 2012
 - Windows Server 2012 R2
 - Windows Server 2016
- ItemPoint requires Microsoft .NET Framework version 4.5
- ItemPoint supports the following SQL Servers as both source and target servers:
 - SQL 2008

- SQL 2008 R2
- SQL 2012
- SQL 2014
- SQL 2016
- ItemPoint supports the following Microsoft Azure source data:
 - Local backups of Microsoft Azure
 - Downloaded copies of database files
 - Downloaded copies of backup files

Note

SQL Server data that is hosted in the Microsoft Azure cloud is not available as source data. Move or copy the data to a local folder before using ItemPoint for Microsoft SQL Server.

- ItemPoint supports the following Microsoft Azure target servers:
 - SQL Server data files in Microsoft Azure
 - SQL Server virtual machines in Microsoft Azure
 - SQL database in Microsoft Azure
-

Note

For virtual environments, virtual operation of tape devices may have restrictions that are imposed by virtual operating systems.

The *ItemPoint for Microsoft SQL Server User Guide* provides the details about using ItemPoint to perform granular-level recovery of SQL backups.

Supported data types

ItemPoint supports the following data types:

- Table level restoring of data
 - Snapshots (if presented as a drive letter or volume of files)
 - Column (data) types (up to 2 GB):
 - BigInt, Int, SmallInt, TinyInt, Bit
 - Binary, VarBinary
 - Char, VarChar, NChar, NVarChar
 - Computed Columns
-

Note

Computed Columns appear in the Source List as <COMPUTED>. Computed Columns referencing an object that does not exist in the target database cannot be copied.

- DateTime, SmallDateTime
- Float, Real
- Image, VarBinary(max)
- Money, SmallMoney
- Numeric, Decimal

- Sql_variant

Note

If a sql_variant column is copied and:

- The sql_variant is a Money or SmallMoney type, it is converted to a decimal on the target.
 - The value is a string type, the collation is not preserved.
-

- Sparse columns without COLUMN_SET
- Text, NText, VarChar(max), NVarChar(max)
- Uniqueidentifier
- XML

Limitations

ItemPoint does not support the following data types:

- Databases backed up using encryption and compression
 - Backup data encryption
 - SQL Server database objects other than tables, rows, and columns (such as system tables, stored procedures, triggers, primary and foreign keys, and CLR) are not supported and do not appear in the ItemPoint for Microsoft SQL Server interface.
 - Files and filegroups
-

Note

Memory-Optimized tables cannot be accessed.

- Columnstore indexes
-

Note

Tables with a columnstore index cannot be accessed.

- ItemPoint does not fully support the following column (data) types:
-

Note

These column types appear in the Source List as either <UNSUPPORTED> or binary and are copied as binary.

- CLR UDT (User Defined Type)
- Cursor
- Date, DateTimeOffset, DateTime2, Time
- Geography, Geometry
- HierarchyId
- Timestamp/RowVersion
- ItemPoint does not support the following column (data) types:
 - COLUMN_SET column

Note

While sparse columns that are defined by COLUMN_SET column appear in the Source List, tables containing this column cannot be copied.

- Filestream

Note

Tables containing Filestream columns cannot be accessed.

Installation requirements

To use the SQL granular-level recovery feature, you must select the following options from the NMM installation wizard:

- The SQL GLR option
This option installs the following required components for granular-level recovery:
 - The Eldos CBFS driver
 - The NMM mount service
 - The NetWorker virtual file system (NWFS) plugin
 - ItemPoint for SQL Server
- The SQL SSMS plugin option
This option is required because SQL granular-level recovery is managed and initiated through the Table Restore tab in the NMM Microsoft SQL Server Management Studio (SSMS) plugin GUI.

The *NetWorker Module for Microsoft Installation Guide* provides more information.

Backup considerations

SQL granular-level recovery requires backups that are created with NMM 9.1 or later. No additional configuration is required for granular-level recovery compatible backups as granular-level recovery is enabled by default for backups that are taken with NMM 9.1 or later. Backups that are taken with older versions of NMM are not granular-level recovery enabled. Backups that use the NetWorker AES encryption or NetWorker compression technology are not supported for granular-level recovery.

Granular-level recovery is supported only for SQL Server backups that are taken with an AFTD or Data Domain device. Client Direct must be enabled for the volumes containing the save sets to be restored in order for the mount operation to succeed. If a Client Direct session cannot be established, the mount operation fails.

If the backup is on a tape device, the backup must be cloned to an AFTD or Data Domain device before you can perform a granular-level recovery operation. In this scenario, it is also required to set the recovery pool in the registry. Contact support for details.

The default creation of SQL granular-level recovery compatible backups does not affect backup workflow. For example, if you are performing a type of backup that is not supported for granular-level recovery, the backup completes successfully. However, if a backup is not granular-level recovery compatible, NMM writes warning messages to the operational logs. If these warning messages impact group completion statuses, disable granular-level recovery backups by typing `ENABLE_GLR=NO` in the application information field of the client resource.

The following backup types are not supported for SQL table level recovery:

- Columnstore indexes

Note

Tables with a columnstore index cannot be accessed.

- File level and file group backups
-

Note

Memory-Optimized tables cannot be accessed.

NMM SSMS plugin considerations

ItemPoint for SQL Server requires administrative privileges, and as a result the NMM SSMS plugin must be started with a user account that has administrative privileges.

The Table Restore tab may display backup versions that are not granular-level recovery enabled. The mount operation fails if any of the save sets in the selected backup version are not granular-level recovery capable. The VFS log file contains details of the cause of failure.

NMM allows only one backup mount to be active at a time. This restriction is built-in to the NMM SMSS plugin GUI workflow. While a mount is active, the Table Restore tab is blocked for use until the ItemPoint for SQL Server GUI is closed.

Mount considerations

The maximum number of recovery sessions the NWFS process (nwfs.exe) uses while mounting a backup is 31 sessions. If NetWorker server parallelism is set and the value is less than 31, the NWFS process uses the parallelism value.

If the number of open recover sessions causes any operational problems with the NetWorker Server, it is possible to reduce the number of recover sessions by modifying the client parallelism registry setting. Contact support for details.

Large database considerations

SQL table level restore can be inappropriate for large databases for scenarios where a large amount of data has been committed to the database before to the backup.

The committed data is represented as outstanding transactions when there are SQL Server full and incremental backups on the database. When you perform table restores, ItemPoint must load all the outstanding transactions that are present in a backup. To load large databases, ItemPoint requires a few hours to read the data because of suboptimal Virtual File System performance.

This problem occurs while ItemPoint is loading the database, and ItemPoint may appear to stop responding while reading the backup.

If you observe this problem, an alternative to using SQL table level recovery is to restore the backup as flat files to a server where you have enough space, and the use ItemPoint with the files on the local disk.

Recovery considerations

Consider the following ItemPoint for Microsoft SQL Server recovery limitations:

- Referential integrity, indexes, and permissions are not restored.
- When restoring tables with TimeStamp columns, the TimeStamp column is not preserved. The value is the time of the restoration for that row. This is the normal behavior for TimeStamp columns.
- Restoring tables with identity columns does not preserve the value of the identity column. The restored table follows the rules of the identity column on the target table.

For example, if the source table has three rows with identity values of 1, 3, 5, once restored, the target has values of 1, 2, 3 because this behavior is the normal behavior of identity columns.

Performing granular-level recovery

In NMM SSMS plugin **NetWorker** window, from the **Table Restore** tab, you can perform granular-level recovery.

Note

The **Table Restore** tab is only visible in the NMM SSMS plugin GUI if the SQL GLR option is selected during NMM installation.

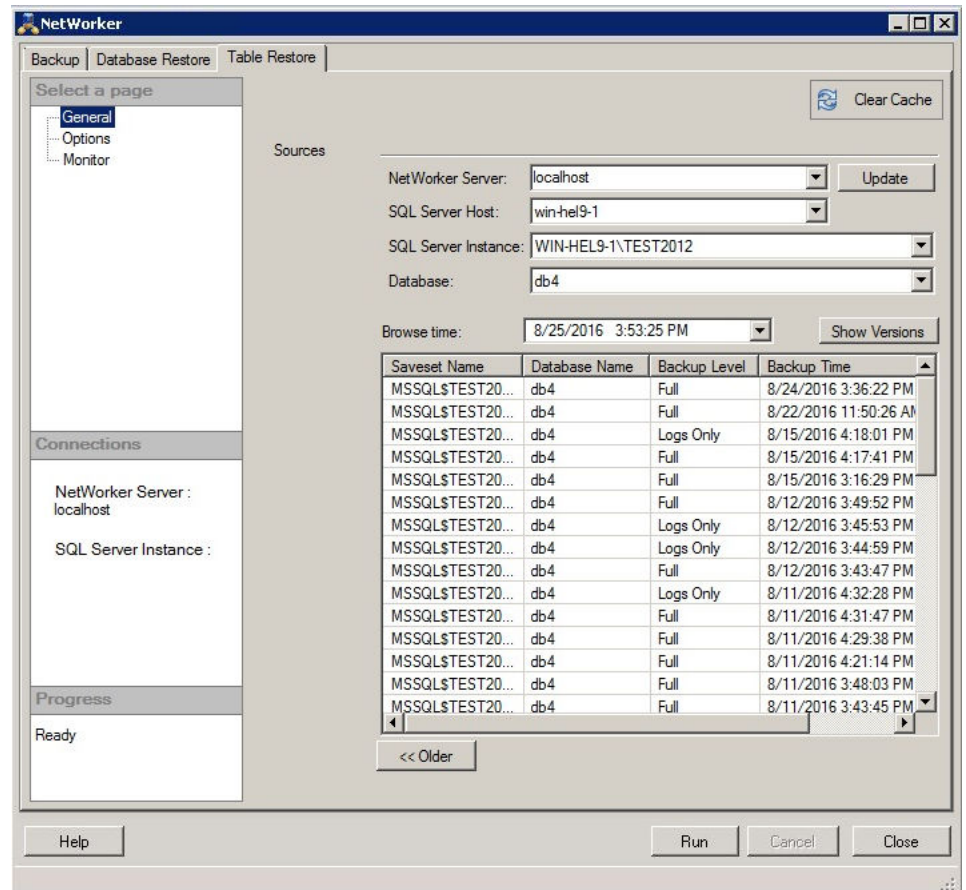
The **Table Restore** tab is used to mount the backup in the NetWorker virtual file system so that it is available for granular-level recovery using the ItemPoint GUI.

Procedure

1. In the SQL Server Management Studio, select the SQL Server instance that you want to recover, and then open the **NetWorker** window.
2. In the **NetWorker** window, on the **Table Restore** tab, under **Select a page**, click **General**.

The **General** page appears.

The following figure shows the **NetWorker** window open to the **Table Restore** tab, on the **General** page.

Figure 64 NetWorker dialog box—Table Restore General page

3. Select the NetWorker server where the backup is located from the **NetWorker Server** list.

Use the **Update** button to update the list of available NetWorker servers.

The **SQL Server Instance** list is refreshed after you select the NetWorker server.

4. From the **SQL Server Host** list, select the client that contains the backup.
5. From the **SQL Server Instance** list, select the server instance that contains the backup.
6. From the **Database** list, select the database.
7. Specify the browse time by using one of the following options:

- Use the **Browse time** list.
- Click **Show versions** and choose the browse date and time.

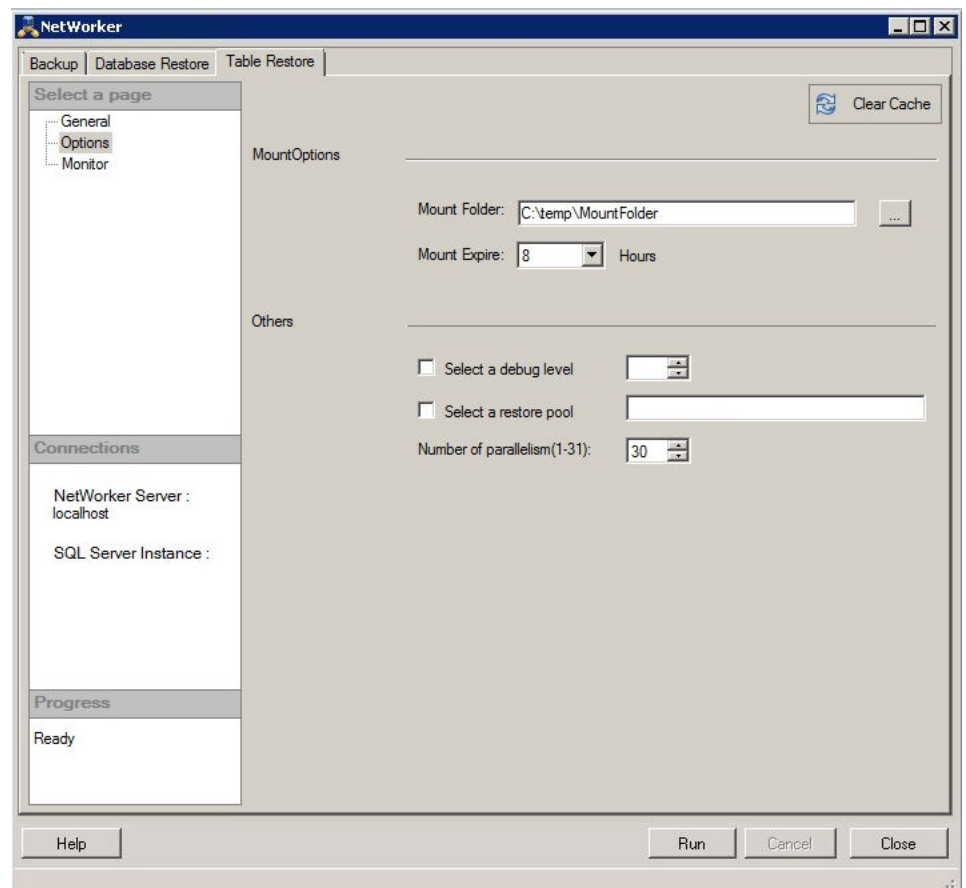
The table is populated with the save sets in the SQL database instance that are available for recovery. The number of available backup versions appears in the bottom bar.

8. From the table, select the save set that contains the items for granular-level recovery.
9. To specify additional granular-level recovery options, under **Select a page**, click **Options**.

The **Options** page appears.

The following figure shows the **NetWorker** window open to the **Table Restore** tab, on the **Options** page.

Figure 65 NetWorker dialog box—Table Restore Options page



10. (Optional) Under **Mount Options**, modify the following settings as required:
 - Edit the **Mount Folder** field to change the location where database is mounted as files. The default mount path is the registry key folder that is specified during installation.
 - Edit the **Mount Expires** field to change the amount of time that the backup remains mounted. The default setting is 8 hours, and the maximum setting is 24 hours.
Valid mount expiration values are 2, 4, 8, 16, and 24 hours.
11. (Optional) Under **Others**, modify the following settings as required:
 - To select the level of information recorded in the logs which can be used in troubleshooting recovery issues, select **Select a debug level**. Levels range 0-9, with 0 representing no information and 9 representing the most amount of information.
 - To specify the clone pool where the save set is restored, select **Select a restore pool**, and then type the name of the restore pool in the field.
 - To change the restore parallelism setting, type or select a value in the **Number of parallelism (1-31)** field.
Parallelism controls the number of recovery sessions the NWFS process uses while mounting a backup. The default and maximum value is 31.

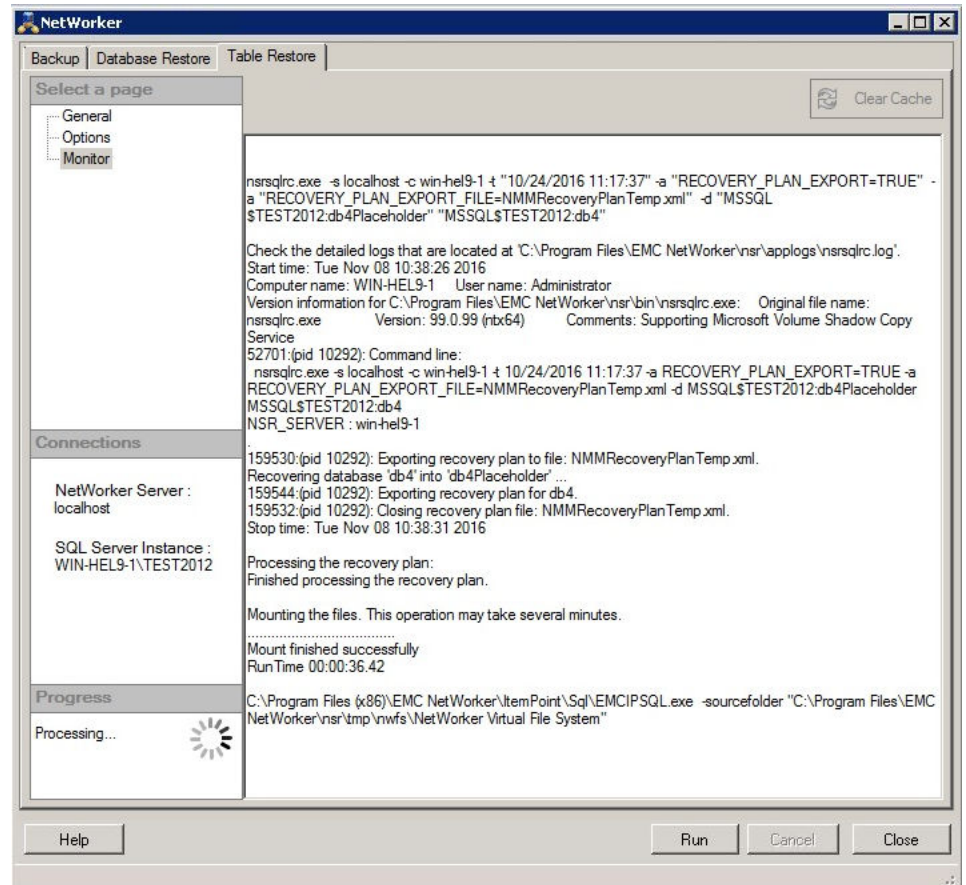
12. To mount the selected backup, click **Run**.

After clicking **Run**, ItemPoint for Microsoft SQL Server opens and the NMM SSMS plugin GUI switches to the **Monitor** page.

The **Monitor** page provides messages about the status of the mount operation.

The following figure shows the **NetWorker** window open to the **Table Restore** tab, on the **Monitor** page.

Figure 66 NetWorker dialog box—Table Restore Monitor page



13. Perform granular-level recovery of the mounted database through the ItemPoint for Microsoft SQL Server software.

The *ItemPoint for Microsoft SQL Server User Guide* provides instructions about performing granular-level recovery. The NMM SSMS plugin GUI is disabled until GLR is complete with ItemPoint.

Dismounting backups after performing granular-level recovery

After performing GLR using ItemPoint for Microsoft SQL Server, the backup must be dismounted from the virtual drive.

For standard restore operations, the backup is dismounted in the NMM SSMS plugin GUI after the restore is completed and ItemPoint is closed.

If the NMM SSMS plugin GUI cannot dismount the backup:

- You can manually dismount the backup by right-clicking the Mount system tray icon and clicking **Dismount Backups**.

- The backup is dismounted based on the timeout setting specified in the **Options** view. The default timeout period for a mounted backup is 8 hours.

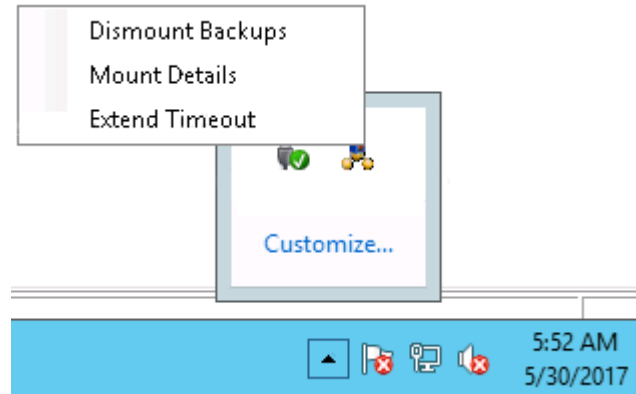
Using the Mount system tray icon

The Mount system tray icon appears when a backup is mounted.

Click the Mount system tray icon to choose from a menu of available actions.

The following figure shows the Mount system tray menu options.

Figure 67 Mount system tray menu

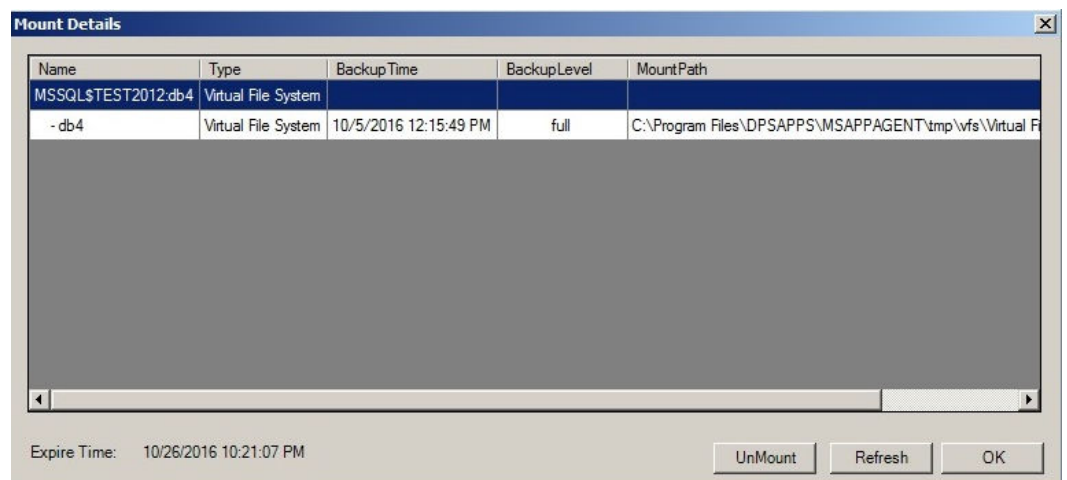


The following actions are available through the Mount system tray icon:

- **Dismount Backups**—Dismount a backup.
- **Mount Details**—View mount details.
- **Extend Timeout**—Extend the period for which the backup is mounted. The mount timeout can be extended by 2, 4, 8, 16, or 24 hours.

The following figure shows the **Mount Details** window.

Figure 68 Mount Details window



CHAPTER 8

Striped Backup and Recovery

This chapter includes the following sections:

- [Overview of striped backups](#).....184
- [Performance considerations](#)..... 184
- [Performing striped backups](#)..... 185
- [Performing striped recovery](#)..... 186

Overview of striped backups

NMM supports the use of multiple stripes for backing up and restoring SQL Server data. Stripes are one or more data streams that are extracted in parallel from a database and written in parallel to multiple media devices, such as tape drives. With NMM, striping can yield a significant performance advantage when a large amount of data is backed up and recovered with multiple tape drives.

Although the SQL Server supports 64 stripes, NMM supports only 32 stripes. This restriction applies to the NetWorker User for SQL Server program and the `nsrsqlsv` and `nsrsqlrc` commands.

When you perform striped backup and recovery operations, ensure that you assign the required user privileges. Read [Access privileges for backup and recovery](#) on page 44 for more information.

Before a striped backup begins, you must set the **Parallelism** field in the **Set Up Server** dialog box in NMC to at least one more than the number of stripes that the backup will use. For example, if you want to use three stripes, specify a value of four or more for the parallelism.

The total number of save streams (SS), which is the data and save set information that is being written to a storage volume during a backup, is the product of client parallelism (P) multiplied by the number of stripes (S), that is, $SS = P \times S$.

NMM supports interleaving for backing up multiple clients to a single backup device. Interleaving writes multiple stripes to the same volume. Recovering interleaved backups can require substantial recovery time. When restoring an entire interleaved backup, one stripe is recovered at a time. This recovery method requires multiple passes over the same areas on the backup medium. When you recover data through an interleaved backup, recover each client's backup separately.

You can specify striping from the following locations:

- In the NetWorker User for SQL Server GUI, start in the **Backup Options** dialog box.
- The NetWorker Administration GUI for a scheduled backup.
- The Windows client host command prompt.
For example, type:

```
nsrsqlsv.exe -s server_name -c client_name -l full -S n
"MSSQL:db_name"
```

where:

- Lowercase `-s` specifies the NetWorker server name.
- Uppercase `-S` with `n` specifies the number of stripes.

Performance considerations

Review the following suggestions in this section to improve striped backup performance.

- In the NetWorker Administration GUI, in the **Create Device** or **Edit Device** dialog box, configure target sessions for the NetWorker devices that are performing the striped backup.
Always specify one session per device when using striping. This configuration eliminates interleaving, yielding the best backup performance on the client host.

- Keep the following criteria in mind when you allot a separate backup device (such as a tape drive) for each strip in the backup or recovery operation:
 - Use a different NetWorker server for the backup and the recovery operations that is unrelated to the SQL Server.
 - Use a dedicated storage node for the striped backups.
- NMM automatically recovers data using the same number of stripes that were specified for the backup.
For best recovery performance, ensure that the same number of devices that are used for the backup are available at recovery time.
- Configure the computer to run the following components:
 - NetWorker client as a NetWorker storage node
 - NMM
 - SQL Server

Note

The NetWorker server cannot be installed on the same host as the NMM client software.

Performing striped backups

NMM supports manual and scheduled backup striping.

You must complete the required steps to perform a manual striped backup from the SQL Server host, which is a NetWorker client.

- Open the NetWorker User for SQL Server GUI:
 1. From the **Operation** menu, select **Backup**.
 2. Select an item to back up.
 3. From the **Options** menu, select **Backup Options**.
 4. In the **Backup Options** dialog box, select **Create a Striped Backup**.
 5. From the **Stripes** list, select the number of stripes.
- From the Windows command prompt, specify the **-S** option with the `nsrsqlsv` command. For example, at the command prompt, type:

```
nsrsqlsv.exe -s server_name -c client_name -l full -S n
"MSSQL:db_name"
```

where *n* specifies the number of stripes you want to use, for example, **-s3**.

- From the NetWorker Administration GUI:
 1. Click **Protection**.
 2. In the expanded left pane, click **Clients**.
 3. In the right pane, right-click the client, and then select **Modify Client Properties**.
 4. On the **Apps & Modules** tab, append **-S*n*** to the `nsrsqlsv` command, where *n* is the number of stripes to use.

5. Click **OK**.

Performing striped recovery

NMM performs striped recovery optimizations so striped recoveries proceed as quickly as possible.

Parallelism is the number of simultaneous data streams during a restore operation. Changing parallelism affects the volume of data that host can process at the same time to reduce restore times and optimize performance.

NMM distinguishes between disk type and tape type devices when determining parallelism. For disk type devices, NMM performs recovery operations with as many parallel streams as the NetWorker Server allows. For tape type devices, the parallelism value is limited to the number of tape devices where each stripe resides.

Refer to the *NetWorker Administration Guide* for details on the device parallelism that can be specified on the NetWorker Server.

Optimal striped recovery operations

Striped recovery assumes, by default, that enough NetWorker tape sessions are available to enable the striped recovery to proceed without blocking.

NMM then assumes the following:

- For tape devices, each stripe was written to a different volume during backup. Interleaving was not used during the striped backup.
- The same number of devices are available during both backup and restore operations. No device has failed since the backup.
- All devices that are used for backup are available. No other client sessions are assigned to the devices.

If the configuration meets all of these requirements, NMM restores the striped recovery. Backup data is delivered to the SQL Server when each tape session is established. This method provides the highest performance and lowest restore time.

Fail-safe striped recovery operations

If a striped backup is interleaved or if fewer tape sessions are available at the time of recovery than during backup, because of a failed or busy device, a fail-safe striped recovery is necessary.

NMM contacts the NetWorker server with the list of striped save sets to determine the available number of tape sessions before it starts striped recovery. If the number of tape sessions is smaller than the number of striped save sets, NMM uses a fail-safe striped recovery algorithm. The algorithm that NMM selects depends on the SQL Server version.

The following table provides guidelines for determining when to activate the fail-safe striped recovery algorithm.

Table 35 Guidelines for fail-safe striped recovery

If	Then
The backup was interleaved	Use NMC to determine if a second mount request has occurred for a volume that is already assigned tape sessions.

Table 35 Guidelines for fail-safe striped recovery (continued)

If	Then
Fewer devices are available	Use NMC to determine if a device has been deactivated or is performing an operation for another client.

About fail-safe striped recovery

When you enable a fail-safe striped recovery for a SQL Server client, additional configuration settings are recommended so the recovery can proceed at maximum throughput.

NMM takes advantage of the SQL Server VDI removable pipes feature. This feature allows third-party backup vendor products to recover a striped backup from fewer devices. However, NMM cannot accurately determine how many tape sessions the NetWorker server can assign. The NetWorker server does not support striped recoveries by default. Therefore, only one tape session is available to recover the striped save sets before the detection phase.

The SQL Server VDI removable pipes function allows a third-party backup vendor to start as many stripe recovery threads as there are tape sessions available. Otherwise, the recovery stops responding.

NOTICE

Because enabling striped recoveries on the NetWorker server might cause file system recoveries to fail, do not use this recovery method as a permanent solution.

Performing fail-safe striped recovery

The procedure in this section applies to normal and copy restore types.

You must complete the required steps to enable a fail-safe striped recovery.

Procedure

1. Open the NetWorker User for SQL Server GUI.
2. On the **Restore** page, in the left pane, mark the root item.
3. From the **Options** menu, select **Restore Options**.
4. Select the **Detect available tape sessions prior to restoring a striped backup** field.

This field is deprecated and this field will be removed from NMM in future releases. In NMM 9.1 and later, this field should always be selected if any backups are being restored from tape type devices. Selecting this field does not affect the performance of restore operations from disk type devices.

5. (Optional) Enable striped recoveries on the NetWorker server by creating the file `striped_recovery` in the `\nsr\debug` directory on the NetWorker server.

NOTICE

You do not need to restart the NetWorker services to activate this setting.

6. From the NetWorker User for SQL Server GUI or from the command prompt, start the restore operation.

7. If the volumes with the striped recovery are not managed by an autochanger, monitor events in NMC.

When a media wait event occurs, load the volume.

CHAPTER 9

Disaster Recovery

This chapter includes the following sections:

- [Overview of disaster recovery](#) 190
- [Disaster recovery features](#).....190
- [Performing disaster recovery](#) 191

Overview of disaster recovery

The examples that follow provide general principles and procedures for recovering data.

Before beginning a SQL Server disaster recovery, review the following:

- *NetWorker Server Disaster Recovery and Availability Best Practices Guide*
- *NetWorker Release Notes*
- Microsoft SQL Server Books Online

To rebuild system databases for SQL Servers with the command prompt, refer to the information provided at the following link: [http://msdn.microsoft.com/en-us/library/ms144259\(v=sql.100\).aspx](http://msdn.microsoft.com/en-us/library/ms144259(v=sql.100).aspx)

Note

The SQL Server databases must be recovered in the correct order. Follow the database recovery steps in the Data Recovery chapter.

Disaster recovery features

The following table lists the disaster recovery features that NMM supports.

Table 36 Disaster recovery features

Disaster recovery features	Descriptions
System database restore automation	<p>Certain system databases require SQL Server service control, including the master and the <code>msdb</code> databases. NMM automates the control of these services as follows:</p> <ul style="list-style-type: none"> • For the master database, the SQL Server restarts in single-user mode as required by the SQL Server. • For the <code>msdb</code> database, the SQL Agent shuts down to close the connections to the <code>msdb</code> database. <hr/> <p>Note</p> <p>NMM does not support snapshot backups or recoveries of the SQL Server master database. Recover from a traditional backup of the master system database during a disaster recovery.</p> <hr/>
Database restore order	<p>When you recover a complete backup of all databases or when you recover certain system databases, the recovery must occur in a specific order. If system databases are present in the recovery list, NMM ensures that the recovery order follows SQL Server procedures as follows:</p> <ol style="list-style-type: none"> 1. The master database is always recovered first. This order ensures that metadata present in the master database is correct for all subsequent restored databases. 2. The <code>msdb</code> database is always recovered after the master database and before all other databases. This sequence ensures that scheduling and other system data present in the <code>msdb</code> database are correct for all subsequent restored databases.

Table 36 Disaster recovery features (continued)

Disaster recovery features	Descriptions
	3. The model database is always recovered after the master and the <code>msdb</code> databases and before all other databases. This sequence ensures that the database configuration is correct for all subsequent databases.
SQL Server startup complete detection	<p>When the SQL Server is started, a database startup process is also started. The SQL Server enables user connections while the startup process is running. However, if the startup process is interrupted by a database <code>RESTORE</code> query, any database that has not yet started is marked as suspect. When the interruption occurs, subsequent restores of the <code>msdb</code> database fail. Subsequent recovery of any user database might also fail unless the Overwrite the Existing Database field is specified in the Properties dialog box.</p> <p>NMM waits for the SQL Server to complete the database startup process for all databases before starting a recovery. NMM waits for the recovery of the master database to complete. This wait process enables a proper restore of the <code>msdb</code> database and user databases following a SQL Server startup.</p> <p>Always use the <code>-f</code> option with the <code>nsrsqlrc</code> command for a recovery that follows a SQL Server startup.</p>
Overwriting capability	<p>Use the Files tab in the Properties dialog box of the NetWorker User for SQL Server GUI to specify that you want to overwrite all databases during a restore operation.</p> <p>NMM does not handle all dependent services. When recovering application services, such as Metadata or Replication services, as well as the databases on which these services depend, the services must be manually shut down.</p> <p>NMM does not ensure that all connections to a database are closed before it recovers a database as required by the SQL Server. Such open connections must be manually closed. The Microsoft SQL Server documentation provides more information.</p>

Performing disaster recovery

If NMM binary files or SQL Server binary files are damaged or lost, follow the instructions that are provided in this section.

The *NetWorker Server Disaster Recovery and Availability Best Practices Guide* provides more information about using NetWorker for disaster recovery.

When not to reinstall the SQL Server

During a disaster recovery of a SQL Server, you must not install the SQL Server in certain cases.

Do not reinstall the SQL Server if any of the following conditions exist:

- A recent NetWorker backup of the `ALL` save set exists.
- The SQL Server was installed before the `ALL` save set backup was performed.
- The SQL Server backups for all of the SQL Server databases exist.

If a SQL Server is installed and then the `ALL` backup is performed, critical state information for the SQL Server is backed up as part of the `ALL` save set. When the `ALL` save set contains the critical state information for the SQL Server and backups

for all of the SQL Server databases exist, you can recover the SQL Server by restoring the **ALL** save set and the SQL Server databases.

Recovering a damaged primary disk

Procedure

1. Shut down the SQL Server.
2. Rebuild the databases on the SQL Server. Information about rebuilding various SQL Servers is available on the Microsoft MSDN website.

The following table provides links to relevant Microsoft documentation for various versions of SQL Server.

Table 37 Rebuilding SQL Server system databases

SQL Server version	Microsoft article
SQL Server 2016	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.130).aspx
SQL Server 2014	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.120).aspx
SQL Server 2012	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.110).aspx
SQL Server 2008 R2 SQL Server 2008	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.100).aspx
SQL Server 2005	http://msdn.microsoft.com/en-us/library/ms144259(v=sql.90).aspx

3. Restart the SQL Server.
4. In the NetWorker User for SQL Server GUI, select **Replace** to recover the SQL Server master database and the `msdb` database.
5. Recover the other SQL application databases.

Note

To run steps 4 and 5 as a single operation, from a command prompt, type
`nsrsqlrc -s... -f MSSQL:`

6. Complete disaster recovery through the NetWorker User for SQL Server GUI:
 - a. Click **Restore**.
 - b. Select **Normal**, and then click **Continue**.
 The **Restore (Normal)** window appears.
 - c. Right-click the SQL Server root and select **Mark All Databases**.
 - d. From the **Options** menu, select **Restore Options**.
 The **Restore Options** dialog box appears.
 - e. Click **Automatically Overwrite Any Database Marked for Restore**, and then click **OK**.

f. Click **Start**.

Recovering a damaged binary disk

Procedure

1. Shut down the SQL Server.
2. Recover the `ALL` save set.

Note

When the `ALL` save set contains the critical state information for the SQL Server and backups for all of the SQL Server databases exist, you can recover the SQL Server by restoring the `ALL` save set and the SQL Server databases.

3. Restart the computer.
4. Shut down the SQL Server.
5. Rebuild the databases on the SQL Server. Information about rebuilding various SQL Servers is available on the Microsoft MSDN website.

The following table provides links to relevant Microsoft documentation for various versions of SQL Server.

Table 38 Rebuilding SQL Server system databases

SQL Server version	Microsoft article
SQL Server 2016	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.130).aspx
SQL Server 2014	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.120).aspx
SQL Server 2012	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.110).aspx
SQL Server 2008 R2 SQL Server 2008	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.100).aspx
SQL Server 2005	http://msdn.microsoft.com/en-us/library/ms144259(v=sql.90).aspx

6. Start the recovery in the NetWorker User for SQL Server GUI. Select **Replace** to recover the SQL Server master database and the `msdb` database.
7. Recover the other SQL application databases.

Note

You can run steps 6 and 7 as a single operation:
From a command prompt, type the following command:
`nsrsqlrc -s... -f MSSQL:`

8. Complete disaster recovery through the NetWorker User for SQL Server GUI:
 - a. Click **Restore**.

- b. Select **Normal**, and then click **Continue**.

The **Restore (Normal)** window appears.

- c. Right-click the SQL Server root and select **Mark All Databases**.

- d. From the **Options** menu, select **Restore Options**.

The **Restore Options** dialog box appears.

- e. Click **Automatically Overwrite Any Database Marked for Restore**, and then click **OK**.

- f. Click **Start**.

Recovering SQL Server and NetWorker server

Complete the required steps if the disks with the SQL Server binaries and the NetWorker online indexes (the `nsr` file system) or the SQL Server data are damaged.

Procedure

1. Recover the NetWorker binaries and online indexes. The *NetWorker Server Disaster Recovery and Availability Best Practices Guide* provides details.
2. Reinstall NetWorker:
 - If the computer that failed is the NetWorker server, reinstall the NetWorker server.
 - If the computer that failed is a NetWorker client, reinstall the NetWorker client.
3. Recover the `ALL` save set.

Note

When the `ALL` save set contains the critical state information for the SQL Server and backups for all of the SQL Server databases exist, you can recover the SQL Server by restoring the `ALL` save set and the SQL Server databases.

4. Restart the computer.
5. Shut down the SQL Server services if they are running.
6. Rebuild the databases on the SQL Server. Information about rebuilding various SQL Servers is available on the Microsoft MSDN website.

The following table provides links to relevant Microsoft documentation for various versions of SQL Server.

Table 39 Rebuilding SQL Server system databases

SQL Server version	Microsoft article
SQL Server 2016	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.130).aspx
SQL Server 2014	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.120).aspx
SQL Server 2012	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.110).aspx

Table 39 Rebuilding SQL Server system databases (continued)

SQL Server version	Microsoft article
SQL Server 2008 R2 SQL Server 2008	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.100).aspx
SQL Server 2005	http://msdn.microsoft.com/en-us/library/ms144259(v=sql.90).aspx

7. Restart the SQL Server.
8. In the NetWorker User for SQL Server GUI, select **Replace** to restore the SQL Server master database and the `msdb` database.
9. Recover the other SQL application databases.

Note

You can run steps 8 and 9 as a single operation:
 From a command prompt, type the following command:
`nsrsqlrc -s... -f MSSQL:`

10. Complete disaster recovery through the NetWorker User for SQL Server GUI:
 - a. Click **Restore**.
 - b. Select **Normal**, and then click **Continue**.
 The **Restore (Normal)** window appears.
 - c. Right-click the SQL Server root and select **Mark All Databases**.
 - d. From the **Options** menu, select **Restore Options**.
 The **Restore Options** dialog box appears.
 - e. Click **Automatically Overwrite Any Database Marked for Restore**, and then click **OK**.
 - f. Click **Start**.

Recovering the SQL Server without reinstalling the SQL Server

If a SQL Server is installed and then the `ALL` backup is performed, critical state information for the SQL Server is backed up as part of the `ALL` save set. When the `ALL` save set contains the critical state information for the SQL Server and backups for all of the SQL Server databases exist, you can recover the SQL Server by restoring the `ALL` save set and the SQL Server databases.

Procedure

1. Reinstall the operating system.
2. Reinstall NetWorker:
 - If the computer that failed is the NetWorker server, reinstall the NetWorker server.
 - If the computer that failed is a NetWorker client, reinstall the NetWorker client.

3. Recover the **ALL** save set.
4. Restart the computer.
5. Shut down the SQL Server.
6. Rebuild the databases on the SQL Server. Information about rebuilding various SQL Servers is available on the Microsoft MSDN website.

The following table provides links to relevant Microsoft documentation for various versions of SQL Server.

Table 40 Rebuilding SQL Server system databases

SQL Server version	Microsoft article
SQL Server 2016	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.130).aspx
SQL Server 2014	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.120).aspx
SQL Server 2012	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.110).aspx
SQL Server 2008 R2 SQL Server 2008	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.100).aspx
SQL Server 2005	http://msdn.microsoft.com/en-us/library/ms144259(v=sql.90).aspx

7. Restart the SQL Server.
8. In the NetWorker User for SQL Server GUI, select **Replace** to restore the SQL Server master database and the `msdb` database.
9. Recover the other SQL application databases. Recover the SQL Server to the most recent backup.

Note

To run steps 8 and 9 as a single operation from the command line, type the following command:

```
nsrsqlrc -s... -f MSSQL:
```

For information about how you can install various SQL Server versions from the command prompt, visit [http://msdn.microsoft.com/en-us/library/ms144259\(v=sql.100\).aspx](http://msdn.microsoft.com/en-us/library/ms144259(v=sql.100).aspx). Use the Other Versions list at the top of the page to view information for a particular SQL version.

10. Complete disaster recovery through the NetWorker User for SQL Server GUI:
 - a. Click **Restore**.
 - b. Select **Normal**, and then click **Continue**.
The **Restore (Normal)** window appears.
 - c. Right-click the SQL Server root and select **Mark All Databases**.

- d. From the **Options** menu, select **Restore Options**.

The **Restore Options** dialog box appears.

- e. Click **Automatically Overwrite Any Database Marked for Restore**, and then click **OK**.
- f. Click **Start**.

Recovering the SQL Server

Use the following procedure to recover the SQL Server.

Procedure

1. Reinstall the operating system.
2. Reinstall NetWorker:
 - If the computer that failed is the NetWorker server, reinstall the NetWorker server.
 - If the computer that failed is a NetWorker client, reinstall the NetWorker client.
3. Recover or reinstall the SQL Server software.
Do not reinstall the SQL Server if any of the following conditions apply:
 - A recent NetWorker backup of the **ALL** save set exists.
 - The SQL Server was installed before the **ALL** save set backup was performed.
 - The SQL Server backups for all of the SQL Server databases exist.
4. Start the SQL Server.
5. When recovery is started, in the NetWorker User for SQL Server GUI, select **Replace** to restore the SQL Server master database and the **msdb** database.
6. Recover the other SQL application databases.

Note

To run steps 5 and 6 as a single operation from the command line, type the following command:

```
nsrsqlrc -s... -f MSSQL:
```

7. Complete disaster recovery through the NetWorker User for SQL Server GUI:
 - a. Click **Restore**.
 - b. Select **Normal**, and then click **Continue**.
The **Restore (Normal)** window appears.
 - c. Right-click the SQL Server root and select **Mark All Databases**.
 - d. From the **Options** menu, select **Restore Options**.
The **Restore Options** dialog box appears.
 - e. Click **Automatically Overwrite Any Database Marked for Restore**, and then click **OK**.

f. Click **Start**.

CHAPTER 10

Bare Metal Recovery

This chapter includes the following sections:

- [Planning bare-metal recovery](#)..... 200
- [SQL Server in a cluster environment](#)..... 204
- [SQL Server in a stand-alone environment](#)..... 208

Planning bare-metal recovery

This section provides an overview of the Windows bare-metal recovery (BMR) solution and considerations when planning a BMR.

Overview

Bare-metal recovery (BMR) is a technique in the field of data recovery and restoration where the backed up data is available in a form that allows you to restore a system from bare metal, that is, without any requirements as to previously installed software or operating system.

Typically, the backed up data includes the necessary operating system, applications, and data components to rebuild or restore the backed up system to an entirely separate piece of hardware. The hardware receiving the restore should have a similar configuration as that of the hardware that was the source of the backup.

The basic BMR is the process of bringing up a server after a disaster and ensuring that the system recovers with the operating system, the applications, and the data as they were at the time of the failure.

Restoring a server to the exact configuration that it had at the time of its destruction can be a difficult task. When this restoration is performed to another hardware, it can add to the complexity of the process and can be time-consuming. Windows BMR solution provides a flexible and reliable method of restoring a server after a disaster.

System requirements

The following sections list requirements to perform Windows BMR. However, the *NetWorker Online Software Compatibility Matrix* at <http://compatibilityguide.emc.com:8080/CompGuideApp/> provides the latest information about the system requirements to perform Windows BMR by using NMM.

Supported operating systems

The *NetWorker Online Software Compatibility Matrix* at <http://compatibilityguide.emc.com:8080/CompGuideApp/> provides information about the operating systems and versions that NMM supports.

Supported Microsoft applications

The *NetWorker Online Software Compatibility Matrix* at <http://compatibilityguide.emc.com:8080/CompGuideApp/> provides information about the Microsoft applications and versions that NMM supports.

NetWorker software version requirements

The *NetWorker Online Software Compatibility Matrix* at <http://compatibilityguide.emc.com:8080/CompGuideApp/> provides information about the NetWorker software versions that NMM supports.

Use the NetWorker Windows BMR 32-bit ISO image to recover an x86 operating system on either an x86 or x64 computer.

Use the NetWorker Windows BMR 64-bit ISO image to recover only an x64 operating system on an x64 computer.

Microsoft BMR requirements

Perform BMR recoveries to the same or similar hardware, and physical to virtual environment.

The following Microsoft KB article provides the requirements to perform a BMR to similar hardware:

<http://support.microsoft.com/kb/249694>

CPU requirements

Consider the following CPU requirements:

- The operating system architecture and the processor architecture must match.
- Basic Input/Output System (BIOS) or Unified Extensible Firmware Interface (UEFI) must match.
- You can treat AMD or Intel processors as being the same if they follow the same architecture. You can recover the operating system backup of an AMD x64 computer to an Intel x64 computer. The process is reversible.
- You can restore the backup of an x86 operating system version only to an x86 processor computer.
- You can restore the backup of an x64 operating system version only to an x64 processor computer.

Hard disk requirements

Consider the following hard disk requirements:

- The disk or RAID drivers that are used in the old system must be compatible with the disk or RAID controllers in the new system.
- For each critical disk on the BMR target system, the startup hard disk capacity on the new system must be greater than or equal to the capacity on the old system. BMR fails if the capacity is smaller by even a single byte.
- Windows BMR supports IDE, SATA, or SCSI hard disks. You can perform the backup on one type of hard disk and recover on another type. For example, SAS to SATA is supported.
- You must restore the backup to the same logical disk numbers as on the original server. You cannot use different logical disk numbers on the target system to recover the critical volumes such as the operating system volume.
- Ensure that the RAID setup on the destination computer does not interfere with the disk order of the hard disks.

NIC driver requirements

To use different Network Interface Card (NIC) after a Windows BMR recovery, install new NIC drivers that match the NIC in the new computer after Windows starts on the new server.

Critical and noncritical volume requirements

Consider the following critical and noncritical volume requirements:

- Windows BMR backs up only critical volumes, and can be used for offline disaster recovery. Use the NetWorker client to back up the non-critical volumes.

Note

NetWorker considers only system volume as a critical volume. If you have installed a Microsoft application on a drive other than the system drive, the drive is not considered as critical. On Windows Server 2008 R2, a volume is critical if a Microsoft application has installed a Windows service on it, but on Windows Server 2012, a volume that has a Windows application service installed is not critical.

- To make a volume critical on Windows Server 2012, set the value of the `HKLM\System\CurrentControlSet\Control\SystemWriter\ReportWin32ServicesNonSystemState` registry key to 0. This ensures that BMR includes the Microsoft application binaries, and the volume on which they are installed is marked as critical.

Note

This registry key may not be in the registry by default. If it is not in the registry, create the `HKLM\System\CurrentControlSet\Control\SystemWriter\ReportWin32ServicesNonSystemState` registry key.

System Reserved Partition requirements

Ensure that System Reserved Partition (SRP) is online before you perform a BMR backup. Otherwise, the backup fails and displays the following error messages:

```
84687:save: Unable to get volume information of file system '\\
\?\Volume{245204f6-5ff7-11e2-a3ac-806e6f6e6963}\': The device
is not ready. (Win32 error 0x15). VSS OTHER: ERROR: VSS failed
to process snapshot: The shadow copy provider had an unexpected
error while trying to process the specified operation. (VSS
error 0x8004230f)
```

```
90108:save: Unable to save the SYSTEM STATE save sets: cannot
create the snapshot.
```

```
86024:save: Error occurred while saving disaster recovery save
sets.
```

If SRP is offline, perform the following steps to bring it online:

1. In the WinPE command prompt, type `diskpart` and press **Enter**.
2. Run the following command to display the list of volumes:
`DISKPART> list volume`
3. Run the following command to select the volume that is offline:
`DISKPART> select <volume_name>`
4. Run the following command to bring the selected volume online:
`DISKPART> online volume`

Protecting an environment before a disaster

To prepare for disaster recovery, back up application data and other necessary files.

Procedure

1. Use NetWorker client to back up non-application data.
The *NetWorker Administration Guide* provides details about using the NetWorker client to backup non-application files.

2. If you use a NetWorker server earlier than 9.0.x, create a NetWorker group without enabling the Snapshot option. Otherwise, create a policy.
3. Create a NetWorker client resource, and assign it to the group that you created in step 2.
4. Configure the NetWorker client resource by typing **ALL** in the **Save set** field, and clearing the **Backup command** and **Application information** fields.
5. Perform a backup.
6. Use NMM to perform a full backup of application data.

The NMM application specific user guides provide details about how to perform a full backup.

BMR by using NetWorker and NMM

Perform Windows file system backup and recovery by using the NetWorker client, and application-specific backup and recovery by using NMM.

Note

Specific information about how to backup and recover Microsoft applications in NMM 18.1 is provided in application-specific user guides.

NetWorker disaster recovery provides an automated BMR solution by using the Windows ASR writer and other VSS writers to identify critical volumes that are required to perform a disaster recovery on a disabled system. BMR is performed offline, that is, when the Windows operating system is inactive. This avoids the necessity to manually reinstall Windows, and the problems that occur when operating system files are being restored to an active Windows system.

NMM 18.1 is compatible with the NetWorker 18.1 client, which provides a true BMR capability for Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008. This capability is built into the NetWorker client and provides BMR support to the same or similar hardware. For Windows BMR, you should download the ISO recovery boot image file from the Online Support website (<https://support.emc.com>). The ISO image provides a wizard that allows you to format disks, recover Windows critical volumes from backup, and restart the server to bring it back online.

To support a NetWorker Windows BMR recovery, download the Windows BMR image from the Online Support website (<https://support.emc.com>). This download enables you to create a bootable Windows BMR ISO that contains NetWorker binaries and a wizard, which controls the recovery process.

The "Windows Bare Metal Recovery to Physical or Virtual Computers" section in the *NetWorker Administration Guide* provides detailed information about how to use the NetWorker Windows BMR image to perform a BMR recovery on protected hosts and VMware virtual machines.

NOTICE

For all the Microsoft applications, after performing Windows disaster recovery and restarting the system, check all the disk and volume configurations. Usually, the disks and volumes appear as on the original system.

However, it is possible, especially in BMR scenarios, that the volume or disk signatures do not match the original ones, and the non-critical volumes or disks are offline and not mounted. Use the Microsoft Disk Manager to bring the volumes and disks online, and then restart the system for drive letter reassignments. Assign the same drive letters that existed before the BMR. Non-critical volumes that the mount points access might have a similar issue.

SQL Server in a cluster environment

This section provides the bare-metal recovery (BMR) backup and recovery procedures for SQL Server by using the VSS technology in a cluster environment.

Backing up a SQL Server for BMR

Configure a SQL Server client resource to back up SQL Server for BMR.

Creating client resources by using the NetWorker Client Backup Configuration Wizard

Create separate client resources to back up the application data and the file system data. The client resource for the file system data must include the `ALL` save set.

The section [Access privileges for backup and recovery](#) on page 44 provides information about how to create a client resource to back up the application data.

The "Creating a client resource with the Client Backup Configuration Wizard" section in the *NetWorker Administration Guide* provides information about how to create a client resource to back up the file system data.

Creating client resources by using the Client Properties dialog box

Create separate client resources to backup the application data and the file system data. The client resource for the file system data must include the `ALL` save set.

This procedure uses the following sample SQL Server configuration:

- You have installed SQL Server on cluster volume `P :`.
- The cluster quorum drive is `E :`.
- You have clustered all the drives except the `C :` drive in the cluster.

Procedure

1. In the Administration view of NMC, create a policy, a workflow, a group, and an action for SQL Server.

The "Data Protection Policies" chapter in the *NetWorker Administration Guide* provides information.

2. Create a client resource to back up the `ALL` save set on the active node:
 - a. In the **NetWorker Administrator** window, click **Protection**.
 - b. In the expanded left panel, right-click **Clients** and select **New**.
 - c. In the **Create Client** dialog box:

- a. In the **Name** field, type the name for the active node.
- b. In the **Group** list, select the group that you created in step 1.
- c. In the **Save set** field, type **ALL**.
- d. Specify other fields according to the requirements.
- e. Click **OK**.

3. Perform the backup.

Ensure that the backup successfully completes.

4. Perform steps 2 and 3 for the passive node.

Note

When you perform step 2ca, in the **Name** field, type the name for the passive node.

5. Perform steps 2 and 3 for the quorum drive.

Note

When you perform step 2ca, in the **Name** field, type the name for the Windows cluster.

6. Perform steps 2 and 3 for the SQL cluster drives.

Note

When you perform step 2ca, in the **Name** field, type the SQL virtual name.

7. Create a client resource with SQL virtual name to back up the SQL application data:

- a. In the **NetWorker Administrator** window, click **Protection**.
- b. In the expanded left panel, right-click **Clients** and select **New**.

The **Create Client** dialog box appears.

c. On the **General** tab:

- a. In the **Name** field, type the SQL virtual name.
- b. In the **Group** list, select the group that you created in step 1.
- c. In the **Save set** field, type **APPLICATIONS:\SqlServerWriter**.

d. On the **Apps & Modules** tab, in the **Backup command** field, type **nsrnmmsv**.

e. Specify other fields according to the requirements.

f. Click **OK**.

8. Create dummy client resources for the physical nodes of the cluster.

9. Perform the backup.

Ensure that the backup successfully completes.

Performing BMR of a SQL Server cluster

The BMR process recovers C : drive only. It does not recover the SQL data, the quorum and SQL cluster drives, and the system databases such as master, msdb, and model.

Procedure

1. Separately recover the active and passive nodes by performing the procedures that the "Performing a Windows BMR recovery to physical or virtual computers" section in the *NetWorker Administration Guide* describes.
2. Recover the cluster quorum database:
 - a. After you have started the active and passive nodes, ensure that the cluster disks are online on the active node in the Disk Management Console of the Windows GUI.
 - b. Add the quorum drive to Failover Cluster Management of the Cluster Configuration UI.
 - c. Select the storage node in Failover Cluster Management and ensure that the quorum disk is online.
 - d. On the SQL Server active node, start the NetWorker User program.
 - e. Click **Recover**.
 - f. In the **Source Client** dialog box, select the source client with the Windows cluster name, and click **OK**.
 - g. In the **Destination Client** dialog box, select the destination client with the active node name, and click **OK**.
 - h. Select the quorum drive and recover the drive.
3. Recover the SQL cluster drives:
 - a. Add the SQL drives to Failover Cluster Management of the Cluster Configuration UI.
 - b. Ensure that the disk is online.
 - c. Configure the drive mount points between the drives so that the mount points are the same as the mount points during the backup.
 - d. On the SQL Server active node, start the NetWorker User program.
 - e. Click **Recover**.
 - f. In the **Source Client** dialog box, select the source client with the SQL virtual name and click **OK**.
 - g. In the **Destination Client** dialog box, select the destination client with the active node name and click **OK**.
 - h. Select all the SQL cluster drives and recover the drives.
4. Build the msdb and model databases:
 - a. Start the physical nodes of the cluster.
 - b. Ensure that you have the SQL Server installation software on the active node.
 - c. In the software, go to the folder where the SETUP file exists and run the following command:

```
Setup/QUIET/ACTION=REBUILDDATABASE/INSTANCENAME=SQL Server
instance name/SQLSYSADMINACCOUNTS=Domain\administrator
username/IAcceptSQLServerLicenseTerms
```

Refer to Microsoft documentation for information about how to rebuild the SQL Server system databases. The following table provides links to relevant Microsoft documentation for various versions of SQL Server.

Table 41 Rebuilding SQL Server system databases

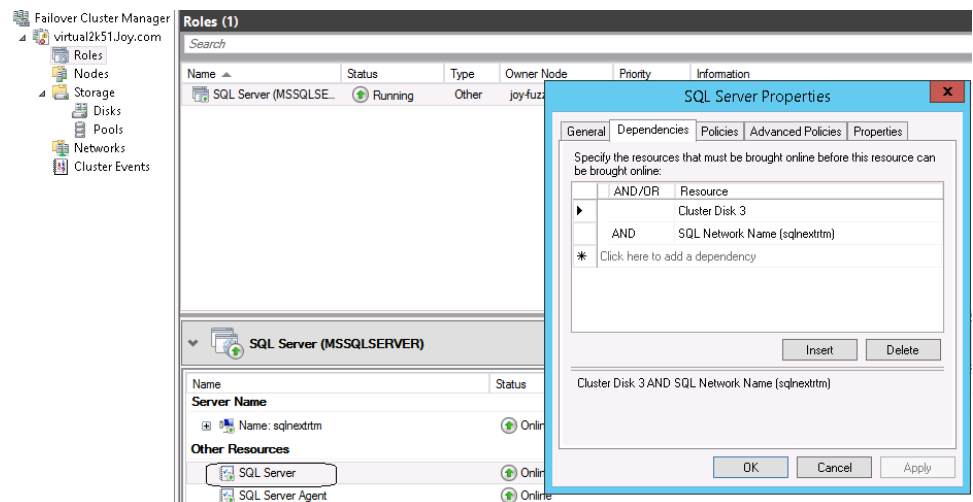
SQL Server version	Microsoft article
SQL Server 2016	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.130).aspx
SQL Server 2014	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.120).aspx
SQL Server 2012	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.110).aspx
SQL Server 2008 R2 SQL Server 2008	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.100).aspx

5. In the **Failover Cluster Management** window of the Cluster Configuration UI, perform the following steps:

- a. Open the **SQL Server Properties** dialog box.

The following figure shows the **SQL Server Properties** dialog box.

Figure 69 SQL Server Properties dialog box



- b. On the **Dependencies** tab, add all the SQL Server dependencies that you must bring online before you bring the SQL Server online. For example, add any dependent cluster disk.
 - c. Bring the SQL virtual resources online.
 - d. Bring the SQL Server online.
 - e. Perform the failover and ensure that the failover is successful.
6. Restore the databases by using the NetWorker User for Microsoft GUI:

- a. Open the NetWorker User for Microsoft GUI.
- b. Select the SQL Server and relevant SQL virtual server instance, and then click **Recover**.
- c. On the **Recover** tab page, click the **Browse** tab.
- d. Expand **APPLICATIONS**, and then select **SqlServerWriter**.
- e. Perform the restore operation.
- f. After the restore operation completes, ensure that the SQL Server service starts for the recovered instance.

SQL Server in a stand-alone environment

This section provides the BMR backup and recovery procedures for SQL Server in a stand-alone environment.

Backing up a SQL Server for BMR

Configure a SQL Server client resource to back up SQL Server for BMR.

Creating client resources by using the NetWorker Client Backup Configuration Wizard

Create separate client resources to back up the application data and the file system data. The client resource for the file system data must include the **ALL** save set.

[Access privileges for backup and recovery](#) on page 44 provides information about how to create a client resource to back up the application data.

The "Creating a client resource with the Client Backup Configuration Wizard" section in the *NetWorker Administration Guide* provides information about how to create a client resource to back up the file system data.

Creating client resources by using the Client Properties dialog box

To back up the application data and the file system data, create separate client resources. The client resource for the file system data must include the **ALL** save set.

Procedure

1. In the Administration view of the NetWorker Management Console, create a policy, a workflow, a group, and an action for Active Directory.

The Data Protection Policies chapter in the *NetWorker Administration Guide* provides information.

2. To back up the **ALL** save set, create a client resource:
 - a. In the **NetWorker Administrator** window, click **Protection**.
 - b. In the expanded left panel, right-click **Clients** and select **New**.

The **Create Client** dialog box appears.

- c. In the **Name** field, type the name of the client resource.
 - d. Click **OK**.
3. In the right panel, right-click the client resource and select **Modify Client Properties**.
4. In the **Client Properties** dialog box:

- a. On the **General** tab:
 - In the **Group** list, select the group that you have created in step 1.
 - In the **Save set** field, type **ALL**.
- b. On the **Apps & Modules** tab, ensure that the **Backup command** and **Application information** fields are clear.
- c. Specify other fields according to the requirements.
- d. Click **OK**.
5. Perform the backup.
Ensure that the backup successfully completes.
6. To back up the **MSSQL**: save set, create a client resource:
Perform step a to step d of step 2.
7. In the right panel, right-click the client resource and select **Modify Client Properties**.
8. In the **Client Properties** dialog box:
 - a. On the **General** tab:
 - In the **Group** list, select the group that you have created in step 1.
 - In the **Save set** field, type the name of the SQL Server instance that contains the databases you must restore:
 - **MSSQL**: for the default SQL Server instance.
 - **MSSQL\$<named_instance>** for a named SQL Server instance.
 - b. On the **Apps & Modules** tab, in the **Backup command** field, type **nsrsqlsv**.
 - c. Specify other fields according to the requirements.
 - d. Click **OK**.
9. Perform the backup.
Ensure that the backup successfully completes.

Performing BMR of a standalone SQL Server

The BMR process recovers only the **C:** drive. It does not recover the SQL data or the system databases, such as master, msdb, and model.

Procedure

1. Perform the procedures that the "Performing a Windows BMR recovery to physical or virtual computers" section in the *NetWorker Administration Guide* describes.
2. Rebuild the SQL Server system databases.

Refer to Microsoft documentation for information about how to rebuild the SQL Server system databases. The following table provides links to relevant Microsoft documentation for various versions of SQL Server.

Table 42 Rebuilding SQL Server system databases

SQL Server version	Microsoft article
SQL Server 2016	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.130).aspx
SQL Server 2014	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.120).aspx
SQL Server 2012	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.110).aspx
SQL Server 2008 R2 SQL Server 2008	https://msdn.microsoft.com/en-us/library/dd207003(v=sql.100).aspx
SQL Server 2005	http://msdn.microsoft.com/en-us/library/ms144259(v=sql.90).aspx

3. Restore the databases by using the NetWorker User for SQL Server GUI:
 - a. Start the NetWorker User for SQL Server program, and connect to the SQL server instance.
 - b. Browse the backed up databases.
 - c. Select the databases and restore the databases.

CHAPTER 11

Troubleshooting

This chapter includes the following sections:

- [Error logs for backup and recovery](#) 212
- [SQL savegroup notifications](#) 212
- [Troubleshooting general issues](#) 213
- [Troubleshooting GLR](#) 213

Error logs for backup and recovery

To help you diagnose problems, certain information is written to an application-specific log file during backup and restore operations.

The following types of information are available:

- Software configuration information
- Operation parameters
- Operation status and error messages

Application logs are generated for the `nsrsqlsv.exe` and `nsrsqlrc.exe` programs, and are created in the `nsr\applogs` folder on the SQL Server host. NMM logs are cumulative and are appended each time the `nsrsqlsv.exe` or `nsrsqlrc.exe` programs run. Debug logs are generated each time the `nsrsqlsv.exe` or `nsrsqlrc.exe` programs are run with debug logging enabled. A log is truncated when disk space is exhausted. The following tables cross-reference the program and log file names.

Table 43 Program and log file names without debug logging enabled

Program	Log file
<code>nsrsqlsv.exe</code>	<code>nsrsqlsv.log</code>
<code>nsrsqlrc.exe</code>	<code>nsrsqlrc.log</code>

Table 44 Program and log file names with debug logging enabled

Program	Log file
<code>nsrsqlsv.exe</code>	<code>nsrsqlsv_<timestamp>.<processid>.log</code>
<code>nsrsqlrc.exe</code>	<code>nsrsqlrc_<timestamp>.<processid>.log</code>

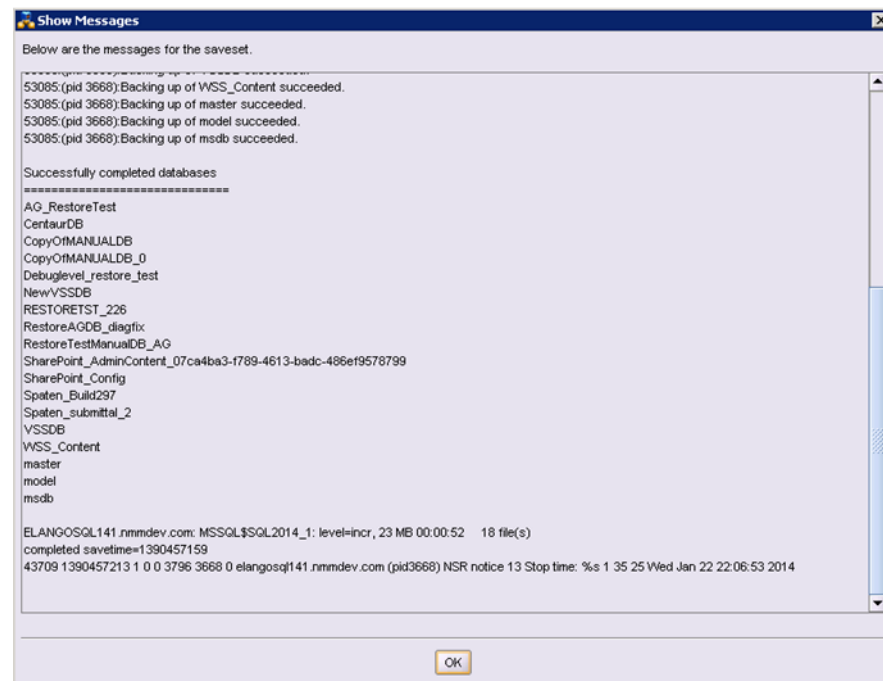
The logging capabilities of `nsrsqlsv` and `nsrsqlrc` are not cluster-aware. For both clustered and nonclustered configurations, the logs are stored on a local disk.

SQL savegroup notifications

A notification provides information about events that occur in the NetWorker and NMM environments. By default, the NetWorker server sends notifications to log files that are located in the `install_dir\applogs` directory.

NMM supports SQL savegroup notifications that list the backup success or failure and provides important details about any failure.

The following figure shows an example of SQL savegroup notification.

Figure 70 Example of SQL savegroup notification

Troubleshooting general issues

"An internal error occurred"

If this error message appears in the operational log, ensure that the user starting the backup or restore operation has the "Read & execute" and "List folder contents" permissions to the database data and log folder. The default data and log folder path is C:\Program Files\Microsoft SQL Server\MSSQL12.MSSQLSERVER\MSSQL\

"bsVdiAbort: SignalAbort failed: protocol error"

If this error message appears in the operational log, the user is not assigned the "Create global objects" Windows user policy. This policy is required to create named file mapping objects. [Access privileges for backup and recovery](#) on page 44 provides information on how to assign Windows user privileges.

Troubleshooting GLR

View the log files to troubleshoot GLR, which are by default located at C:\Program Files\EMC NetWorker\nsr\appllogs. Set a higher debug level in the **Options** view of the NMM SSMS plugin GUI for more detailed logs for troubleshooting.

This section contains solutions to common problems.

Mount failure because save set is not GLR-capable

While mounting a backup for granular-level recovery with the SSMS plugin, the **Monitor** page provides messages about the status of the mount operation. If the

backup selected for table-level restore is not a granular-level recovery capable backup, then the operation fails and the following error message appears:

```
Error: Saveset '<saveset_ID>' is not GLR capable.
```

SQL GLR failing to establish a Client Direct session

Client Direct is required for NMM SQL GLR. Ensure that Client Direct is enabled before continuing with GLR.

To verify that a Client Direct session is established, use the following steps:

1. Validate that the NetWorker device is enabled for Client Direct.
This verification must only be performed for AFTD devices. Data Domain is automatically enabled for Client Direct. The *NetWorker Administration Guide* provides more information about Client Direct.
2. Validate that the client has name resolutions for the systems.
If Data Domain is being used, ensure that the client has name resolution for the Data Domain device. If an AFTD storage node is being used, ensure that the client has name resolution for the storage node.
3. Check the application logs directory in the NetWorker Virtual File System (NWFS) log file, `nwfs.raw`, and look for messages confirming that a Client Direct session was established.
 - The message "Successfully established a client direct recover session for saveset" confirms that a Client Direct session is successfully established.
 - The message "Could not establish a client direct recover session for saveset '<saveset_ID>', exiting" indicates that a Client Direct session could not be established.

Cannot enable a Client Direct session and GLR failing as a result

Client Direct is required for GLR. However, if you cannot enable Client Direct for either policy or technical reasons, use the following workaround to allow GLR to continue without Client Direct.

1. On the NMM Client host, create a folder and name it "debug" in the `\nsr\` directory, if the folder does not already exist.
2. Within the "debug" folder, create an empty file and name it "nodirectfile" with no file name extension.
You may be required to create the "nodirectfile" file from a DOS Shell command line.

Note

This workaround disables Client Direct for all client operations, including subsequent backups. This workaround is against NMM best practices and you may run into timeout and other restore issues if you do not enable Client Direct.

"Saveset '<saveset_ID>' is not GLR capable"

This error message is logged in the `nwfs.raw` file when the save set you are trying to perform GLR on is not GLR capable. The [Considerations](#) on page 172 section provides information on supported backup types.

Table Restore tab does not appear in the NMM SSMS plugin GUI

If the **Table Restore** tab is not visible in the **NetWorker** window of the NMM SSMS plugin GUI, verify that ItemPoint is installed in Programs and Features. To install ItemPoint, the SQL GLR option must be selected while installing or changing installation of NMM with the Install Wizard.

Large databases cause slow loading times

ItemPoint may appear to stop responding while reading a backup of a very large database. The time for ItemPoint to read this data from the backups may take hours because of suboptimal Virtual File System performance.

If this issue occurs, perform a SQL flat file restore of the same backup to a server with enough space, and then use ItemPoint on the files on the local disk.

The [Considerations](#) on page 172 section provides more information.

