

Dell EMC NetWorker Module for Microsoft for Hyper-V

Version 18.2

User Guide

302-005-255

REV 01

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

Published December, 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 Technologies, Dell, EMC, 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	Introduction 15
	Overview.....16
	NMM protection methods for Hyper-V.....16
	Guest-level backup and recovery.....16
	Image-level backup and recovery.....17
	Granular level recovery.....18
	Comparing NMM protection methods for Hyper-V.....19
	Using NMM with Hyper-V.....20
	Using NMM in a stand-alone environment.....21
	Using NMM in a Clustered Shared Volumes environment.....21
	Using NMM with Hyper-V virtual machines over SMB 3.0.....22
	Supported backup and recovery workflows.....23
	Required Hyper-V privileges.....24
	Required SMB privileges.....25
	NMM 18.2 compatibility with NetWorker 8.2.3 or 8.2.4 servers26
Chapter 2	Best Practices and Recommendations 27
	Non-ASCII files and directories.....28
	Required AFTD DFA device settings for Hyper-V environments.....28
	Hyper-V Server backup and recovery best practices.....28
	Microsoft application backup and recovery within a Hyper-V virtual machine best practices.....29
	Improving VSS-based backup performance in Windows clusters with CSV....29
	Data mining using Hyper-V granular level recovery.....30
	Restrictions and requirements for relocating and recovering data to a different location.....30
	Restrictions for GLR of Hyper-V virtual machines.....31
	Hyper-V Server disaster recovery best practices.....31
Chapter 3	Backups 33
	Backup overview.....34
	Block based backups.....34
	Supported Hyper-V backup types and levels.....35
	Files included in backups.....35
	Supported and unsupported features for Hyper-V backups.....36
	Requirements and considerations for Hyper-V virtual machine backups.....42
	VSS backups.....45

	VSS backup consistency types with NMM.....	46
	Configuring backups.....	46
	Creating a client resource for a VSS-based backup by using the Client Backup Configuration wizard.....	47
	Manually creating a client resource for a VSS-based backup by using the Client Properties dialog box.....	52
	Improving performance of VSS CSV backups by using multiple cluster nodes as proxies.....	57
	RCT backups.....	60
	Configuring backups.....	61
	Creating a client resource for an RCT-based backup by using the Client Backup Configuration wizard.....	62
	Manually creating a client resource for an RCT-based backup by using the Client Properties dialog box.....	67
	Improving performance of RCT-based backups by using multiple cluster nodes as proxies.....	71
	Manually configuring highly available backups (cluster-aware backups).....	71
	Viewing backup status and summary.....	72
Chapter 4	Recoveries	73
	Overview of recoveries.....	74
	Recovery scenarios and GUIs to use for various types of recoveries.....	74
	Recovering Hyper-V virtual machines.....	78
	Recovering Hyper-V standalone server virtual machines.....	79
	Recovering Hyper-V clustered server virtual machines.....	84
	Recovering Hyper-V Server virtual machines at granular level.....	89
	Using NMM 9.1 or later to recover the backups that were performed by using NMM 8.2.x.....	94
Chapter 5	File Level Recoveries	95
	Introduction.....	96
	Required ports for Hyper-V File Level Restore GUI	96
	Performing a browser-based file level restore.....	97
	Performing a directed file level restore.....	98
	Monitoring file level restores.....	99
	Hyper-V FLR web UI log files.....	100
Chapter 6	Data Protection Add-in for SCVMM	101
	Overview of the Data Protection Add-in for SCVMM.....	102
	Recoveries.....	102
	Backups.....	102
	Supported versions.....	102
	Software dependencies.....	103
	Required privileges.....	103
	Installation and configuration overview.....	104
	How the Data Protection Add-in works with SCVMM.....	104
	Workflows overview.....	105
	GUI overview.....	106
	SCVMM user roles and allowed actions.....	106
	Supported scopes and contexts.....	107
	Installation and uninstallation.....	107
	Installing SCVMM and the SCVMM console.....	108
	Installing the Data Protection Add-in.....	108
	Importing the Data Protection Add-in.....	109

	Activating the Data Protection Add-in.....	109
	Uninstalling the Data Protection Add-in.....	110
	Upgrading the Data Protection Add-in.....	111
	Preferences.....	111
	Adding NetWorker servers.....	112
	Removing NetWorker servers.....	113
	Setting the refresh interval.....	113
	Including debug output for logging purposes.....	113
	Using multiple NetWorker servers that define the same clients and virtual machine save sets.....	113
	Data Protection Add-in overview data.....	114
	SCVMM Recoveries.....	120
	Viewing available virtual machines.....	122
	Recovering virtual machines to the original location.....	123
	Redirected recoveries.....	123
	Recovering a deleted virtual machine.....	126
	Monitoring.....	126
	Troubleshooting.....	127
	Recovered virtual machine does not start.....	127
	Installation fails due to access issue.....	128
	The Data Protection Add-in for SCVMM displays an incorrect NetWorker Server version.....	128
	Importing fails due to access issue.....	128
	Virtual machine attributes might display incorrect values.....	128
	Redirected recovery appears to succeed but no virtual machine appears in Hyper-V Manager.....	129
	Checks for redirected recovery failures	129
	Avoid virtual machine names with the same name within an SCVMM context.....	129
	Cluster virtual machine backups do not display on the Recover page	129
	Redirected recovery is not supported when the virtual machine name or virtual machine configuration path contains special characters.....	130
Chapter 7	Windows Bare Metal Recovery Solution	131
	Microsoft Hyper-V Backup and BMR.....	132
	Backing up Hyper-V for BMR.....	132
	Performing BMR of Hyper-V.....	132
	Microsoft System Center Virtual Machine Manager backup and BMR.....	133
	Backing up System Center Virtual Machine Manager for BMR.....	133
	Performing BMR of a System Center Virtual Machine Manager...	135
Chapter 8	Troubleshooting	137
	Troubleshooting generic issues.....	138
	Troubleshooting backups issues.....	138
	Troubleshooting recovery issues.....	141
Appendix A	Recovering SQL Server, Exchange Server, and SharePoint Server Items from a Hyper-V Virtual Machine	145
	Overview.....	146
	Recovering items that are stored on a Hyper-V virtual machine.....	146
	Recovering SQL Server items from a Hyper-V virtual machine....	148

CONTENTS

Recovering Exchange Server items from a Hyper-V virtual machine149

Recovering SharePoint Server items.....150

FIGURES

1	Guest-level backup and recovery environment.....	17
2	Image-level backup and recovery environment.....	17
3	RCT-based full backup environment.....	18
4	RCT-based incremental backup environment.....	18
5	GLR environment.....	19
6	Two-node Hyper-V failover cluster.....	22
7	Image-level backups in a Windows Server cluster with SMB.....	23
8	Specify Backup Options page for a standalone setup.....	50
9	Specify Backup Options page for a clustered setup.....	50
10	Specify Backup Options page for a standalone setup.....	65
11	Specify Backup Options page for a clustered setup.....	65
12	Standalone - Hyper-V virtual machine recovery options page.....	82
13	Standalone - Destination Hyper-V Server and path page.....	83
14	Clustered - Hyper-V virtual machine recovery options page.....	87
15	Clustered - Destination cluster node or Hyper-V Server, and path page.....	88
16	GLR mount point location.....	91
17	Hyper-V virtual machine GLR options page.....	93
18	Data Protection Add-in architecture.....	105
19	Data Protection Add-in Preferences page.....	110
20	Data Protection Add-in.....	112
21	Data Protection Add-in Overview page for Administrator, Fabric Administrator, and Read-Only Administrator user roles.....	115
22	Virtual machine Protection Details tooltip for Administrator, Fabric Administrator, and Read-Only Administrator user roles.....	116
23	Virtual machine Protection Details window for Administrator, Fabric Administrator, and Read-Only Administrator user roles.....	116
24	Data Protection Add-in Overview page for Tenant Administrator and Application Administrator user roles	118
25	Virtual Machine Backup Status tooltip for Tenant Administrator and Application Administrator user roles.....	119
26	Virtual machine Protection Details window for Tenant Administrator and Application Administrator user roles.....	120
27	Data Protection Add-in for SCVMM Recover page.....	121
28	Data Protection Add-in for SCVMM Monitoring page.....	127
29	Change NetWorker Server window.....	147
30	Select Viewable Clients window.....	147
31	Selecting granular level recovery.....	148
32	Selecting SQL Server items for recovery from a Hyper-V virtual machine.....	148
33	Selecting Exchange Server items for recovery from a Hyper-V virtual machine.....	149

FIGURES

TABLES

1	Revision history.....	12
2	Style conventions.....	12
3	Comparison of guest and image-level backup and recovery.....	19
4	VSS Writer and VSS Provider used.....	20
5	Supported backup and recovery workflows.....	24
6	Access privileges needed for SMB backup and recovery	25
7	Supported backup types.....	35
8	Supported backup levels.....	35
9	Virtual machine files included in backups.....	35
10	Supported features.....	37
11	Supported special characters.....	39
12	Unsupported features.....	40
13	Backup tasks for Hyper-V	46
14	Supported application information variables.....	47
15	Hyper-V save set syntax.....	52
16	Hyper-V application information variables.....	53
17	Backup tasks for Hyper-V	61
18	Supported application information variables.....	62
19	Hyper-V save set syntax.....	67
20	Hyper-V application information variables.....	68
21	Recovery scenarios.....	74
22	GUIs used for various types of recoveries.....	78
23	Recovery options for virtual machines.....	78
24	SCVMM user roles and actions allowed by the Data Protection Add-in.....	107
25	Virtual machine IDs after redirected recovery.....	124
26	Backup types.....	146

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

Purpose

This guide contains information about using the NetWorker Module for Microsoft (NMM) 18.2 software to back up and recover Hyper-V virtual machines by using the Microsoft Volume Shadow Copy Service (VSS) and Resilient Change Tracking (RCT) technologies.

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
- NetWorker data protection policy management
- NetWorker block based backup (BBB) technology
- Microsoft VSS technology
- Microsoft RCT technology
- Microsoft Hyper-V Server technology
- Microsoft Failover Cluster technology

Revision history

The following table presents the revision history of this document.

Table 1 Revision history

Revision	Date	Description
01	December 14, 2018	First release of this document for the NetWorker Module for Microsoft 18.2 release.

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

Table 2 Style conventions (continued)

	<ul style="list-style-type: none"> • Pathnames, file names, file name extensions, prompts, and syntax • 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://www.dell.com/support>
- <https://community.emc.com>

Where to get support

The Support website <https://www.dell.com/support> provides access to product licensing, documentation, advisories, downloads, and how-to and troubleshooting information. The information can enable you to resolve a product issue before you contact Support.

To access a product-specific page:

1. Go to <https://www.dell.com/support>.
2. In the search box, type a product name, and then from the list that appears, select the product.

Knowledgebase

The Knowledgebase contains applicable solutions that you can search for either by solution number (for example, KB000xxxxxx) or by keyword.

To search the Knowledgebase:

1. Go to <https://www.dell.com/support>.
2. On the **Support** tab, click **Knowledge Base**.
3. In the search box, type either the solution number or keywords. Optionally, you can limit the search to specific products by typing a product name in the search box, and then selecting the product from the list that appears.

Live chat

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

1. Go to <https://www.dell.com/support>.
2. On the **Support** tab, click **Contact Support**.
3. On the **Contact Information** page, click the relevant support, and then proceed.

Service requests

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

1. Go to <https://www.dell.com/support>.
2. On the **Support** tab, click **Service Requests**.

Note

To create a service request, you must have a valid support agreement. For details about either an account or obtaining a valid support agreement, contact a sales representative. To get the details of a service request, in the *Service Request Number* field, type the service request number, and then click the right arrow.

To review an open service request:

1. Go to <https://www.dell.com/support>.
2. On the **Support** tab, click **Service Requests**.
3. On the **Service Requests** page, under **Manage Your Service Requests**, click **View All Dell Service Requests**.

Online communities

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

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

Introduction

This chapter includes the following sections:

- [Overview](#) 16
- [NMM protection methods for Hyper-V](#) 16
- [Using NMM with Hyper-V](#) 20
- [Supported backup and recovery workflows](#) 23
- [Required Hyper-V privileges](#) 24
- [NMM 18.2 compatibility with NetWorker 8.2.3 or 8.2.4 servers](#) 26

Overview

Microsoft Hyper-V is a hypervisor-based server virtualization product for Microsoft Windows Server. Hyper-V enables you to create multiple virtual machines (VMs) on a standalone server or Windows cluster server to consolidate workloads.

NetWorker Module for Microsoft (NMM) supports VSS-based backup, recovery, and granular recovery of Hyper-V virtual machines that run on the Hyper-V role that is installed on Windows Server 2012, 2012 R2, 2016, and on Server Core installations for Windows Server 2012, 2012 R2, and 2016.

NMM supports RCT-based backup, recovery, and granular recovery of Hyper-V virtual machines that run on the Hyper-V role that is installed on Windows Server 2016, and on Server Core installations for Windows Server 2016.

Note

The Microsoft Hyper-V documentation provides a complete and updated list of system requirements and supported guest operating system versions. The *NetWorker E-LAB Navigator*, which is available at <https://elabnavigator.emc.com/elab/elhome>, provides the most up-to-date and accurate listing of hardware, operating system, service pack, and application versions that the NMM client supports.

NMM protection methods for Hyper-V

You can perform Hyper-V guest-level or image-level backup and recovery depending on certain criterion, such as the Windows operating system that is running on the guest and where the NMM software is installed.

Guest-level backup and recovery

For guest-level backup and recovery, install an NMM client on each virtual machine that hosts databases or specific applications, for example, Exchange Server or SharePoint Server. To NMM, each virtual machine is a separate client, and you can perform individual backups of each virtual machine and Microsoft application.

The following figure illustrates Hyper-V guest-level backup and recovery with NMM.

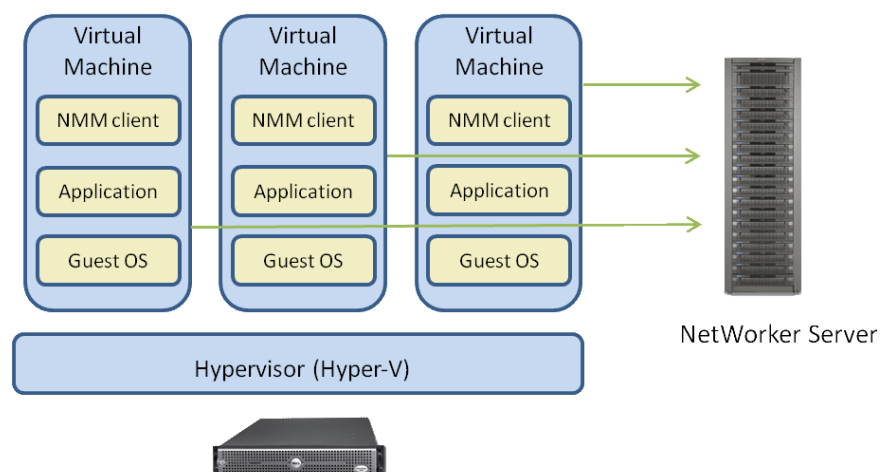
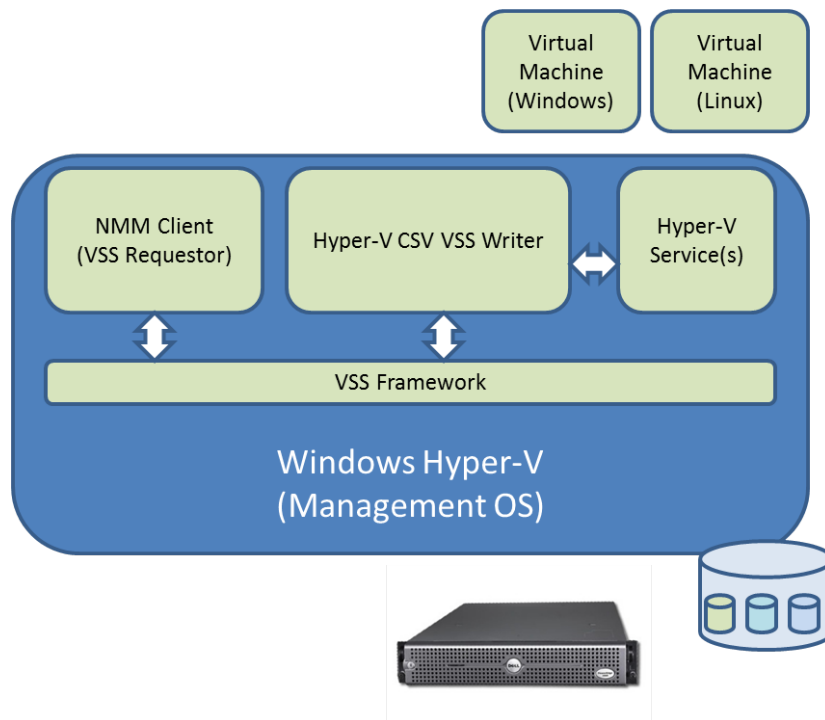
Figure 1 Guest-level backup and recovery environment

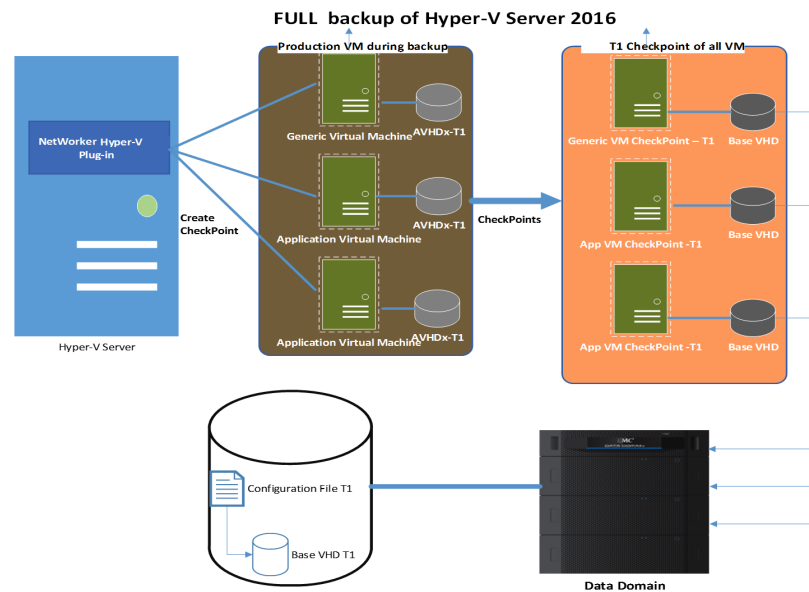
Image-level backup and recovery

For image-level backup and recovery, install an NMM client on the Hyper-V Server. You can perform a full and incremental image-level backup of the virtual machines.

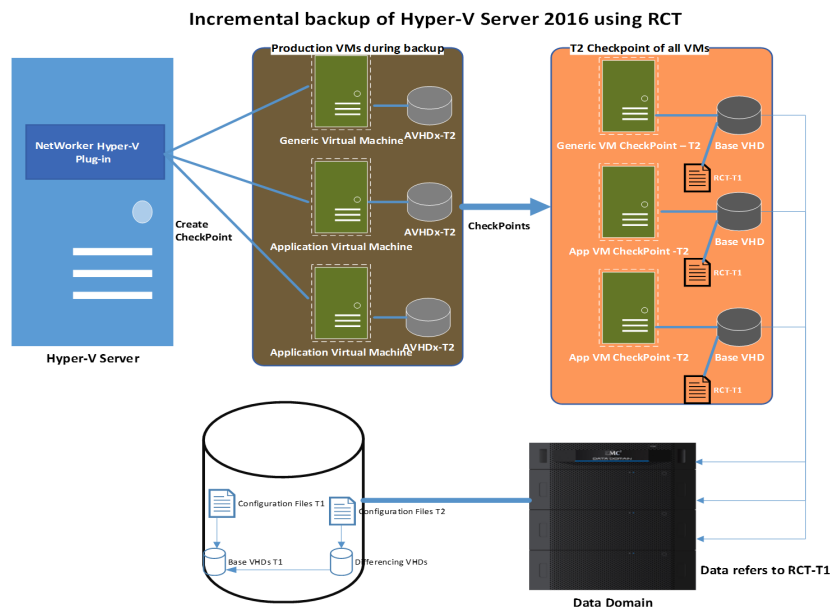
The following figure illustrates the VSS-based image-level backup and recovery environment.

Figure 2 Image-level backup and recovery environment

The following figure illustrates the RCT-based image-level full backup environment.

Figure 3 RCT-based full backup environment

The following figure illustrates the RCT-based image-level incremental backup environment.

Figure 4 RCT-based incremental backup environment

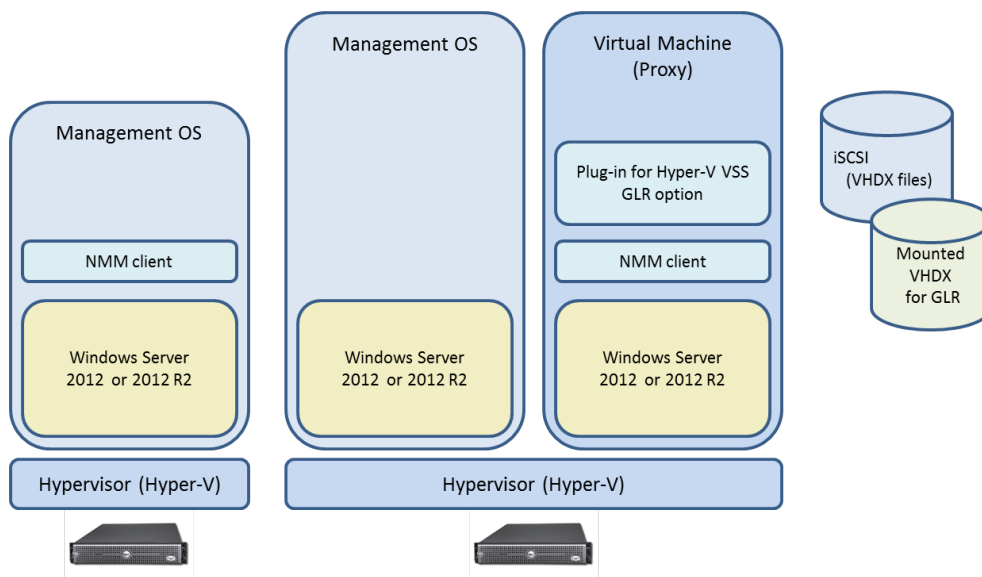
Granular level recovery

Granular level recovery (GLR) enables you to recover specific files from a single backup without recovering the full virtual machine. GLR reduces the recovery time. You can perform GLR by using the NMM client software.

You can perform GLR of NMM backups of Hyper-V virtual machine that has a Windows operating system. NMM does not support GLR of non-Windows virtual machines on Hyper-V, for example Linux virtual machines. The NMM GLR feature mounts the virtual machine that contains the items to recover.

The following figure illustrates the GLR environment.

Figure 5 GLR environment



Comparing NMM protection methods for Hyper-V

The following table compares guest-level and image-level backup and recovery.

Table 3 Comparison of guest and image-level backup and recovery

Requirement	Guest-level	Image-level
NMM	Install on each virtual machine	Install only on the Hyper-V Server
NetWorker client	Install on each virtual machine	Install only on the Hyper-V Server
NetWorker server network connection	Required for each virtual machine	Required only on the Hyper-V Server
iSCSI/pass-through disk support	Required	Not required
Windows Bare Metal Recovery (BMR) backup	Can be performed (by using the NetWorker client DISASTER RECOVERY save set)	Can be performed
Virtual machine status for backup	Virtual machine must be online	Virtual machine can be online, offline, or saved state
Customized backup, including exclusion of certain files or file types	Can be performed (by using the NetWorker client DISASTER RECOVERY save set)	Cannot be performed
Application-aware backup and recovery	Can be performed, for applications such as: <ul style="list-style-type: none"> Exchange Server SharePoint Server SQL Server Active Directory 	Cannot be performed

Table 3 Comparison of guest and image-level backup and recovery (continued)

Requirement	Guest-level	Image-level
Application-consistent backup and recovery	Can be performed	Can be performed (for applications with VSS Writer and VSS integration component)
Individual files and folders recovery	Can be performed	Can be performed using granular-level recovery
Disaster recovery	Can be performed by: <ul style="list-style-type: none"> Recovering the operating system state critical volumes through the NetWorker client BMR wizard Recovering applications and non-critical volume data with NMM 	Can be performed with NMM

Using NMM with Hyper-V

You can use NMM with Hyper-V in the stand-alone and clustered environments, and over Server Message Block (SMB) 3.0.

NMM supports the following Hyper-V configurations:

- Local volumes on a stand-alone server with Windows Server 2012, 2012 R2, 2016
- Cluster Shared Volumes (CSV) on a cluster with Windows Server 2012, 2012 R2, 2016
- Server Message Block (SMB) 3.0 file shares on the following Windows Server 2012, 2012 R2, 2016 file servers:
 - Stand-alone file server
 - Scale-Out File Server (SOFS)

The *NetWorker Module for Microsoft Installation Guide* lists the Hyper-V hardware requirements.

The following table lists the type of environment and the corresponding VSS Writer and Provider used.

Table 4 VSS Writer and VSS Provider used

Type of environment	VSS Writer used	VSS Provider used
Stand-alone	Hyper-V VSS Writer	Microsoft Software Shadow Copy Provider
CSV	Hyper-V VSS Writer and CSV VSS Writer	Microsoft CSV Shadow Copy Provider
Server Message Block (SMB) 3.0	Hyper-V VSS Writer	Microsoft File Share Shadow Copy Provider

Using NMM in a stand-alone environment

The Hyper-V role can be installed on a stand-alone machine. On a stand-alone Hyper-V Server, NMM uses the Hyper-V VSS Writer to take VSS snapshots of virtual machines.

Using NMM in a Clustered Shared Volumes environment

Cluster Shared Volumes (CSV) is a feature in Windows Server which allows all nodes within a failover cluster to concurrently access to shared disks.

Microsoft and NMM refer to the node in the cluster where a CSV is locally mounted as the "coordinating node." NMM provides the option to move the CSV ownership among the various nodes to provide the best backup and recovery performance.

NMM supports physical and virtual proxy nodes for Hyper-V CSV backups. When you specify a Preferred Server Order List (PSOL) in the Application Information attribute when configuring a client resource for the Cluster Server Name, NMM performs a snapshot on the cluster primary node, and then each proxy node backs up the snapshots in parallel.

With NMM, you can use virtual machines on a Hyper-V cluster as proxy nodes for CSV backups. NMM reassigns and distributes the backup workload to selected virtual proxy nodes. All physical proxy nodes perform backups in parallel, which increases backup performance. At the same time, because the proxies are virtual machines, the hypervisor governs their compute, disk, and network resource utilization to ensure they do not unnecessarily impact other virtual machine workloads.

You may also run NMM as a proxy on the physical Hyper-V node where the system resource utilization may not be governed by the Hypervisor. These virtual proxies must be highly available Hyper-V virtual machines, and the virtual proxies must be connected to the same domain as the physical nodes. A mix of physical and virtual machines is supported.

Note

For Hyper-V cluster and CSV environments, including proxy environments, it is recommended that you install the NetWorker client and NMM on all nodes in the cluster and the virtual proxies because the cluster ownership (and cluster alias) can failover to any cluster node within the cluster and NMM schedules backups against the cluster alias.

In a failover cluster, all servers (nodes) run Hyper-V and can host one or more virtual machines. A clustered virtual machine is only active on one node but can be configured to fail over to other nodes. Microsoft supports failover clustering for Hyper-V through CSV.


NMM implements the single snapshot and multiple snapshot features to protect the Hyper-V CSV environment during a backup:

- **Single snapshot feature**—A snapshot of each CSV is created on the active (master) node of the cluster and data is rolled over from proxy nodes.
- **Multiple snapshot feature**—A snapshot is taken of one or more CSV at a time.

Hyper-V CSV in a failover cluster

The following figure illustrates a Hyper-V failover cluster with two nodes. There are four virtual machines that can fail over between the nodes, and a fifth virtual machine that runs exclusively on the second node.

Figure 6 Two-node Hyper-V failover cluster



Using NMM with Hyper-V virtual machines over SMB 3.0

You can use NMM to back up Hyper-V virtual machines over SMB 3.0 file shares on Windows Server 2012, 2012 R2, and 2016 file servers.

The SMB file shares are supported on the following file servers:

- Stand-alone file server
- Scale-Out File Server (SOFS)
- Clustered file server

[Required SMB privileges](#) on page 25 describes the required permissions for SMB backup and recovery.

Note

NMM for Hyper-V is qualified for file shares hosted on Windows file server, VNX file server, NetApp file server, and Nutanix file server. All file servers must support SMB 3.0. All file shares must adhere to Microsoft specifications when used for Hyper-V Server.

Windows Server Hyper-V stand-alone configurations with SMB file shares

For stand-alone Hyper-V Server with virtual machines on SMB storage, install NMM on the Hyper-V Server.

Windows Server Hyper-V clusters with SMB file shares

If you store virtual machines on SMB 3.0 file shares for use by a Hyper-V cluster, you can configure NMM to perform federated backup and restore of the virtual machines.


The SMB file shares can be on either a Scale-Out File Server (SOFS) or on a clustered file server.

To perform backups of virtual machines on SMB file shares, install NMM on each node in the Hyper-V cluster.

The federated cluster client receives backup and restore requests from the NetWorker server and forwards the requests to the NMM client on the cluster nodes. The NMM client performs the backup and returns the backup data and metadata to the NetWorker server. The federated cluster client manages all NMM client requests and ensures that you can back up all virtual machines on SMB file shares in the cluster.

The following figure illustrates a Windows Server cluster where the federated cluster client enables image-level backups of virtual machines on SMB file shares.

Figure 7 Image-level backups in a Windows Server cluster with SMB



Note

The Microsoft TechNet website provides instructions on configuring an SMB file share.

Supported backup and recovery workflows

The following table lists the Hyper-V backup and recovery workflows that this release of NMM supports.

Table 5 Supported backup and recovery workflows

Workflow	VSS support	RCT support
Standalone full and incremental backups	Yes	Yes
CSV backup and recovery	Yes	Yes
SMB backup and recovery	Yes	Yes
Standalone Host Components' backup	Yes, for only Windows Servers 2012 and 2012 R2	No
Federated backup with PSOL	Yes, physical and virtual machines	Yes, only physical machines
Partial backup	Yes	Yes
Granular Level Recovery	Yes	Yes
Alternate recovery or redirected recovery	Yes	Yes
Support for special characters and unicode characters	Yes	Yes
Support for virtual machines with the same name, shared VHD or VHDx, guest cluster, pass through disk, and dynamic disk	No	No
Support for virtual machines with a configuration version earlier than 6.2	Yes	No

Required Hyper-V privileges

Review the *NetWorker Module for Microsoft Administration Guide* for the required privileges for all Microsoft applications that are backed up and recovered through NMM. For additional privileges required for backup and recovery of Hyper-V virtual machines in a CSV environment or over SMB 3.0, review this section.

Note

SMB and CSV backups and recoveries in NetWorker Restricted Datazone (RDZ) environments require additional permissions and configuration. The *NetWorker Module for Microsoft Administration Guide* provides details about the required permissions and configuration. The *NetWorker Administration Guide* provides detailed information about the NetWorker RDZ feature.

During client resource configuration for Hyper-V backups, you must provide the Domain User account and password, instead of providing a Domain Administrator account and password. Perform the following steps to create a domain user and provide the required privileges.

Procedure

1. Create a Domain User.
2. Add the following Active Directory security groups to the newly created Domain User:
 - Backup Operators
 - Hyper-V Administrator
 - Windows Authorization Access Group
 - Users
 - Remote Desktop Users
 - Group Policy User Control
 - Group Policy Creator Owners

Note

To remotely manage Hyper-V hosts, you need the Remote Desktop Users, Group Policy User Control, and Group Policy Creator Owners privileges. For example, to remotely view errors, events, configuration, and so on. Otherwise, you do not need these privileges.

3. On each cluster node, log in, and then add the Domain User account to the following groups on the local node:
 - Users
 - Administrators
 - Hyper-V Administrators
4. Provide local administrator privileges to the Domain User.
5. Provide access for cluster management to each group. Open Windows PowerShell and type the following command:


```
PS C:\...\NMMEMC> Grant-ClusterAccess -User domain\user -Full
```

This command grants the Domain User account full access to the cluster, which provides access for cluster management to each group.

Required SMB privileges

SMB backup and recovery require privileges beyond Hyper-V backup and recovery privileges.

The following table describes the required privileges for SMB backup and recovery.

Table 6 Access privileges needed for SMB backup and recovery

SMB configuration	Required privileges
Stand-alone file server Scale-Out File Server Cluster File Server	Perform one of the following steps: <ul style="list-style-type: none"> • Add backup permissions for the backup user on all file servers in the cluster. • Add the backup user as the cluster administrator (domain administrator). • When Windows is used as the file server, add server accounts for all cluster nodes and virtual proxies to the local Administrators account of each server in the file server tier.

Table 6 Access privileges needed for SMB backup and recovery (continued)

SMB configuration	Required privileges
Scale-Out File Server	In the Local Backup operator group of each SMB node, configure the application server as a member of the Backup Operators group.
Cluster File Server	Add each Cluster File Server node to the SMB nodes of the Local Backup Operator group.

Verify that the Hyper-V Server and the file server are in the same domain. Recoveries require the same permissions as backups.

To enable communication between the SMB host and clients, install the File Share Shadow Copy Agent on the file server that hosts the SMB file shares.

NMM 18.2 compatibility with NetWorker 8.2.3 or 8.2.4 servers

NMM supports backup and recovery with NetWorker client version 18.2 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. For more details, see the individual NMM release sections of the *NetWorker E-LAB Navigator*, which is available at <https://elabnavigator.emc.com/elab/elhome>.

Note the following limitations when you configure NMM backup and recovery with an NMM 18.2 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 18.2. As a result, you cannot configure a dedicated storage node when you use NetWorker 18.2 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.

CHAPTER 2

Best Practices and Recommendations

This chapter includes the following sections:

• Non-ASCII files and directories.....	28
• Required AFTD DFA device settings for Hyper-V environments.....	28
• Hyper-V Server backup and recovery best practices.....	28
• Microsoft application backup and recovery within a Hyper-V virtual machine best practices.....	29
• Improving VSS-based backup performance in Windows clusters with CSV.....	29
• Data mining using Hyper-V granular level recovery.....	30
• Restrictions and requirements for relocating and recovering data to a different location.....	30
• Restrictions for GLR of Hyper-V virtual machines.....	31
• Hyper-V Server disaster recovery best practices.....	31

Non-ASCII files and directories

If you create a client resource by using the **Client Properties** dialog box and the **Save set** field contains non-ASCII characters, you must edit the **Save operations** field on the **Apps & Modules** tab for the client resource.

To access the **Save operations** field, in the **NetWorker Administration** window, click **View > Diagnostic Mode**.

In the **Client Properties** dialog box, on the **Apps & Modules** tab, in the **Save operations** field, specify `I18N:mode=utf8path`

Required AFTD DFA device settings for Hyper-V environments

For Hyper-V environments, when creating a NetWorker AFTD DFA device on an NTFS or ReFS volume, Microsoft requires certain settings.

If the NetWorker AFTD DFA device is created on an NTFS volume, virtual hard disk (VHD/VHDX) files must be uncompressed and unencrypted. If the NetWorker AFTD DFA device is created on an ReFS volume, virtual hard disk (VHD/VHDX) files must not have the integrity bit set.

Hyper-V Server backup and recovery best practices

When you perform Hyper-V backups and recoveries, consider the following best practices:

- NMM skips virtual machine pass-through disks during Hyper-V Server backups.
- NMM supports Windows Server Failover Clustering, which allows you to configure a failover of virtual machine.
- To host virtual machines, you can use the storage device that is connected to Fibre Channel or iSCSI storage.
- Do not take a Hyper-V VSS server snapshot of Hyper-V virtual machines that are part of a guest cluster, such as SharePoint farm, SQL AAG, and Exchange DAG. To back up a guest cluster on the Hyper-V virtual machine:
 1. Install the NMM client on the virtual machine.
 2. Perform the guest cluster backup within the virtual machine.

The Microsoft website provides recommendations and requirements about using SharePoint, SQL AAG, or Exchange DAG and Hyper-V together.

Microsoft application backup and recovery within a Hyper-V virtual machine best practices

Microsoft applications backup and recovery are performed within the Hyper-V virtual machine and use application and system component writers that are available on the virtual machine.

- Microsoft recommends using backups within the virtual machine as the preferred method for Exchange backup and recovery.
- NMM skips virtual machine pass-through disks in Hyper-V backups. NMM supports pass-through disks backups within the virtual machine.
- NMM supports Windows Server Failover Clustering with iSCSI storage.
- NMM does not support virtual machine Windows Server failover clustering with Fibre Channel storage because SCSI-3 is not supported in Hyper-V virtual machines.

Improving VSS-based backup performance in Windows clusters with CSV

When a cluster node is used as the proxy client, and you use the cluster proxy client to perform image-level backups of the virtual machines in a Windows Server cluster, NMM backs up the data from the cluster node that owns the CSV, on which the virtual machine files are present.

For example, Cluster Node 1 owns CSV1 on which the virtual machine 1 files reside, and the cluster proxy client runs on Cluster Node 2. When the cluster proxy node backs up virtual machine 1, the backup process:

- Creates a shadow copy of CSV1.
- Streams the backup data from Cluster Node 1 to Cluster Node 2.
- Routes the backup data to the NetWorker server.

In this example, the backup performance depends on network performance between the cluster nodes. The performance for this backup is slower than backups where the CSV node ownership is co-resident with the cluster proxy client.

When using a cluster proxy client for virtual machine backups, consider the following recommendations to improve performance for image-level backups:

- Maximize the network bandwidth between the cluster nodes.
- Move the CSV ownership to the proxy cluster client so that the shadow copies of these volumes are local to the backup process.

To maximize the backup performance, ensure that the cluster node that runs the NMM proxy cluster client owns the targeted CSVs. Before you move the CSVs, consider the following recommendations:

- Ensure that the cluster node with the cluster client proxy has the capacity to own all physical nodes.
The CSV owner node is responsible for file system metadata updates to the NTFS partition. If you change the ownership of a CSV volume to a single node, the performance of all the associated virtual machines on the CSV may be impacted.

The cluster proxy client node must have the capacity to be the owner of all the CSVs.

- Ensure that any CSV you move is in the “healthy state”, online, and in full access mode.

There are two ways to change the CSV ownership to the proxy node:

- Use the Failover Cluster Manager GUI.
- Use the PowerShell Module `FailoverClusters` cmdlet `Move-ClusterSharedVolume`.

The Microsoft Failover Cluster document provides additional instructions for moving CSV ownership.

Data mining using Hyper-V granular level recovery

NMM can perform granular level recovery for backups of Hyper-V virtual machines that were created with NMM 3.0 or above. NMM with Hyper-V also supports data mining the information from the virtual machine image drives by using a tool such as EMC ItemPoint.

To prepare to mine the data, use the NMM GUI to mount the virtual machine, attach the VHDs, and load the virtual machines.

For example, if the virtual machine guest is running SharePoint, first use the NMM GUI to mount the Hyper-V virtual machine image, attach the VHDs, and load the virtual machines. Then use EMC ItemPoint for SharePoint to recover SharePoint sites, lists, libraries, and items.

You must keep the NMM GUI open while you explore and recover files on the mounted virtual machine VHDs. If you close or change the focus of the NMM GUI, you lose access to the mounted VHDs. A warning is displayed when a closure or focus change causes loss of access to a mounted virtual machine image.

[Recovering SQL Server, Exchange Server, and SharePoint Server Items from a Hyper-V Virtual Machine](#) on page 145 provides more information.

Restrictions and requirements for relocating and recovering data to a different location

Hyper-V has several restrictions on relocating and recovering to other locations.

NMM does not support directed recoveries of virtual machines to a Hyper-V Server that is a later version than the source Hyper-V Server. Although you can restore a virtual machine to a destination Hyper-V Server that runs a Hyper-V Server version later than the source Hyper-V server, the virtual machine may not fully function on that server. For mixed environments, you might be unable to perform a redirected restore of a virtual machine from one type of environment to another. Mixed environments include the following configurations:

- Environments with both stand-alone and clustered Hyper-V Servers
- Cluster environments with different operating systems and types of virtual machine storage (CSV and SMB file shares)

Hyper-V does not support:

- Recovering Hyper-V virtual machines to non-Hyper-V Servers.
- Recovering the Host Component file to a different location.

- Relocating or redirecting Hyper-V backups that were taken before an NMM upgrade.

Before you relocate or recover Hyper-V backups, review these requirements:

- The parent partition must run Windows Server 2012 or later to recover, with relocation of files, a virtual machine that has Hyper-V snapshots.
- The destination host must have the NMM client installed.
- When you perform a directed recovery of a virtual machine to a Hyper-V Server that differs from the source, you must update the Network Adapter settings of the virtual machine with the Hyper-V Manager before you start the virtual machine.

Restrictions for GLR of Hyper-V virtual machines

The following restrictions apply when you perform a GLR of Hyper-V virtual machines:

- Windows Server does not support recovery of deduplicated data. To recover deduplicated volume data, enable the Deduplication role.
- Hyper-V GLR does not support differencing disk with parent and child hard disk on different hard drives.

Hyper-V Server disaster recovery best practices

When you perform a disaster recovery, consider the following best practices:

- To get the maximum benefit from the Hyper-V Role, create separate virtual machines for each application, so that the application-type backup and recovery that are performed at the host level is only for Hyper-V.
- After performing disaster recovery of the Hyper-V Server, you might need to recover applications within each virtual machine if:
 - Separate virtual machine backups are performed.
 - The backups are more recent than the complete Hyper-V server backups.
- This type of backup is best used for Bare Metal Recovery of a guest and for recovery of operating system roles.
- For Host Component file backups, perform a back up whenever Hyper-V configuration changes are made. You do not need to back up the Host Component file each time a virtual machine guest is backed up.
- In NMM, the Hyper-V Writer does not support backup of the Host Configuration file to a proxy client.
- The primary purpose for recovering the Host Component file in NMM is for disaster recovery of the Hyper-V Server.
- Roll-forward recovery is not available for virtual machine level disaster scenarios. From a Hyper-V server, a roll-forward recovery of a virtual machine is not possible. Recoveries from a Hyper-V server are point-in-time (disaster recovery).

CHAPTER 3

Backups

This chapter includes the following sections:

• Backup overview	34
• VSS backups	45
• RCT backups	60
• Manually configuring highly available backups (cluster-aware backups)	71
• Viewing backup status and summary	72

Backup overview

You can perform full and incremental backups of Hyper-V virtual machines by using the Microsoft's VSS and RCT backup technologies, and NetWorker block based backup feature. The backup strategy for a Hyper-V environment must include:

- Stand-alone Hyper-V Servers
- Clustered Shared Volumes
- Hyper-V virtual machines over SMB 3.0

NMM does not back up the management operating system. To protect the Hyper-V management operating system, perform a disaster recovery backup with the NetWorker client. The "Windows Bare Metal Recovery Solution" chapter provides details.

Block based backups

To perform Hyper-V backups by using the VSS and RCT backup technologies, NMM supports only NetWorker block based backup (BBB) as the backup type or option. The *NetWorker Module for Microsoft Administration Guide* provides information about BBB.

NMM 9.0.1 and later support synthetic full and incremental forever BBB backups:

- **Synthetic full backup:** This backup combines a full backup and subsequent incremental backups to form a new full backup. A synthetic full is equivalent to a traditional full backup and can be used in the same ways as a traditional full backup.
- **Incremental forever backup:** This backup backs up the files that have changed since the last full or incremental backup.
Incremental backups reduce storage consumption, network traffic, and backup time. NMM Hyper-V incremental backups rely on changed block tracking (CBT) in the virtual machine. To enable incremental backup of a virtual machine, Windows Server 2012 and later use the "IncrementalBackupEnabled" property of the "Msvm_VirtualSystemSettingData" data type in WMI v2. Setting this property to "true" enables incremental backups of the virtual machines.

It is a best practice to create the following schedule policies to leverage these backup levels:

- Incremental forever, which is used to perform BBB backups to Double Data Rate (DDR) target.
- Incremental forever full on the first day of the month, which performs a full backup on the first day of each month. This is a best practice for a BBB to an AFTD target because it limits the depth of the differencing chain.

When you create a client, NMM enables incremental backups by default. NMM automatically enables incremental backups for a virtual machine the first time it is backed up.

NMM promotes incremental backups to full backup under the following circumstances:

- Unable to query and get differential changes of disks by using Microsoft API, in the case of an RCT-based backup.
- NMM does not find a previous full backup of the virtual machine.
- NMM does not find a previous RCT-based backup of the virtual machine.

- The last backup of the virtual machine is not RCT-based.
- New disks are added to the virtual machine.
- The required Recovery Snapshot cannot be validated.
- Incremental backup is disabled for the virtual machine.

When one virtual machine backup is promoted to full, NMM does not promote other virtual machines in the backup set. NMM uses multiple shadow copy sessions by a single backup operation to perform full and incremental backups.

Supported Hyper-V backup types and levels

The following tables list the backup types and levels that NMM supports for Hyper-V.

Table 7 Supported backup types

Type of backup	Description
Federated Hyper-V CSV backup	The backup includes virtual machines that are stored on clustered shared volumes.
Federated Hyper-V over SMB 3.x backup	The backup includes all cluster virtual machines that are stored on SMB 3.x file servers.
Hyper-V Servers standalone backup	The backup includes all virtual machines and the host component.
Hyper-V Servers standalone over SMB 3.0 backup	The backup includes all standalone virtual machines that are stored on SMB 3.0 file servers.

Note

NMM 9.1 and later supports all these backup types for Hyper-V Server 2016.

Table 8 Supported backup levels

Target device	Supported backup level	Resulting backup level
Data Domain	Full	Full
	Incremental	Full
AFTD	Full	Full
	Incremental	Incremental

Files included in backups

NMM backs up the virtual machine files that are listed in the following table:

Table 9 Virtual machine files included in backups

File type	File name extension	Description
Virtual hard disk files	.VHDX	For the virtual machines that are created with Windows Server 2012 and later, Hyper-V uses the Microsoft Virtual Hard Disk (VHDX) specification to

Table 9 Virtual machine files included in backups (continued)

File type	File name extension	Description
		store virtual hard disks for virtual machines. A virtual machine can have one or more virtual disks.
Differencing virtual hard disk files	.AVHDX	A virtual machine snapshot creates one differencing VHDX file per VHDX.
Virtual machine configuration	.XML	Hyper-V uses a virtual machine configuration file in XML format to store virtual machine settings (for example, CPU, memory, VHDXs). Note The .XML file type does not apply to Windows Server 2016.
Virtual machine running state files	.BIN .VSV	Hyper-V uses a virtual machine configuration file in XML format to store virtual machine running state (memory) files. Note The .BIN file type does not apply to Windows Server 2016.
Virtual machine configuration snapshots	.XML	A virtual machine snapshot creates a copy of the current virtual machine configuration and saves it to enable rollback. Note The .XML file type does not apply to Windows Server 2016.
Virtual machine configuration data files	.vmcx	On Windows Server 2016, virtual machine configuration data files use a .vmcx file name extension.
Runtime state data files	.vmrs	On Windows Server 2016, runtime state data files use a .vmrs file name extension.
Bitmap files	.mrt .rct	On Windows Server 2016, bitmap files use .mrt and .rct file name extensions.

NOTICE

For a federated backup to succeed, ensure that all the virtual machine files are present on either CSV or SMB storage, but not on the local disk.

Supported and unsupported features for Hyper-V backups

The following sections provide information about supported features, supported special characters, and unsupported features for Hyper-V backups.

Supported features

The following table provides information about supported features for Hyper-V backups.

Table 10 Supported features

Feature or functionality	Description
Disabling incremental backups and enabling only full backups	<p>By default, when Hyper-V backups are performed by using NMM, the first instance of the backup is always full and subsequent backups are incremental. Hyper-V enables Change Block Tracking (CBT) for all the VHDs associated with a virtual machine. After the first instance of the backup, that is full backup, NMM performs an incremental backup and Hyper-V creates a differencing Virtual Hard Disk (AVHD). All data changes to the virtual machine during the backup and after the backup go to the differencing disk. The AVHD can grow in size if there are a number of changes in the virtual machine between two backups, which may require extra space on the production storage, for example, CSV, SMB shares, or local disk.</p> <p>The "Hyper-V incremental backup and differencing disks" section describes how differencing disks are created between backups when CBT is enabled.</p> <p>NMM provides the user with the option to perform only full backups, instead of full and incremental backups. To ensure that only full backups are performed, the user must disable CBT. When CBT is disabled, the differencing disk is not created. A user has the option to enable CBT or disable CBT at anytime as per their business needs.</p> <p>When a user disables CBT in an existing environment, the previously created differencing disks are merged after the backup.</p> <p>To enable or disable incremental backups for Hyper-V when creating a client resources for backups, use the Application information variable NSR_DISABLE_INCREMENTAL. The "Configuring a client resource manually by using the NetWorker Management Console" provide details.</p>
Client Direct to AFTD or Data Domain Boost storage devices	<p>The Client Direct feature reduces bandwidth usage and bottlenecks at the storage node and provides highly efficient backup data transmission.</p> <p>NMM performs Hyper-V Server backups by using Client Direct, which enables clients with network access to AFTD or Data Domain Boost storage devices to send their backup data directly to these devices, bypassing the NetWorker storage node. The storage node manages the devices for the NetWorker clients, but does not handle the backup data. Ensure that the clients have the required permissions or accessibility to the target device, otherwise backups fail.</p> <hr/> <p>Note</p> <p>When using a AFTD device, configure the device using UNC path.</p> <hr/> <p>When you create a client resource, NetWorker enables the Client Direct feature by default.</p>

Table 10 Supported features (continued)

Feature or functionality	Description
	The <i>NetWorker Administration Guide</i> provides details about the Client Direct to AFTD and Data Domain Boost storage devices.
CSV cluster level and individual CSV virtual machine backups	<p>NetWorker performs CSV virtual machine backups through a client resource that is created for the cluster virtual server only. You create client resources for all the nodes in the cluster and for the cluster server. However, the backup is scheduled against the cluster virtual server client resource only. NetWorker indexes the backup against the cluster server name.</p> <hr/> <p>Note</p> <p>A Hyper-V CSV distributed backup supports only conventional backups from a temporary shadow copy (rollover). NMM does not support proxy host backups and instant backups that use persistent point-in-time shadow copies.</p> <hr/> <ul style="list-style-type: none"> To perform a CSV-level federated backup, you must specify the application information variables for the cluster virtual server and the client resource in the Client Properties dialog box of NMC. <ul style="list-style-type: none"> In the Save set field, type <code>Applications: \Microsoft Hyper-V\</code> to back up all the CSV virtual machines in the cluster. In the Application information field, type <code>NSR_FEDERATED_BACK=YES</code> and <code>NSR_FEDERATED_TYPE=CSV</code>. To perform a CSV virtual machine backup, in the Save set field, type <code>APPLICATIONS: \Microsoft Hyper-V \</code> to back up the CSV virtual machine in the cluster. NMM supports backup of non-clustered virtual machines that run on specific cluster nodes. NMM excludes virtual machines that do not reside on the CSV from the CSV backup. Backup and recovery of non-clustered virtual machines is managed through the individual physical node name, not the cluster virtual server client resource. The physical node name is the client resource name.
Backing up Hyper-V virtual machines residing on Windows Server 2012, 2012 R2, and 2016 over SMB 3.0	You must back up stand-alone servers and non-CSV failover clusters over SMB the same way you back up local virtual machines. To backup SMB 3.0 Hyper-V cluster, NMM uses a federated backup architecture.
Special characters in virtual machine names and virtual machine configuration paths	You can use all foreign language characters and certain special characters in virtual machine names and virtual machine configuration paths, for stand-alone, CSV, and SMB 3.0 configurations. Use the NetWorker User for Microsoft GUI to recover backups of virtual machines with special characters and unicode characters in virtual machine name,

Table 10 Supported features (continued)

Feature or functionality	Description
	<p>VHD Paths, VHD name, and so on. The NMM Hyper-V File Level Restore (FLR) GUI and Data Protection Add-in for SCVMM GUI cannot be used for this.</p> <hr/> <p>Note</p> <p>NMM has been qualified with German, Spanish, French, Simplified Chinese, Traditional Chinese, and Japanese.</p> <hr/> <p>The Table 11 on page 39 table lists the special characters that NMM supports.</p> <p>Examples of usage of special characters:</p> <ul style="list-style-type: none"> Virtual machine name: VM@emc%1% VHD paths: C:\@folder\&^% VHD names: C:\folder\VM@emc(1).vhd Files and folders inside virtual machine: samplefile!@#.txt Folder paths: C:\@folder\&^\sample.txt Mount paths: C:\mountSplcharvms\!@#\$\$% <hr/> <p>Note</p> <p>Backups are skipped in the following scenarios:</p> <ul style="list-style-type: none"> The virtual machine names contain Double quotes ("), single quotes ('), square brackets ([]), forward slash (/), or question mark (?). The path of a mount point location contains unicode characters, such as Chinese or Korean characters, and so on. <hr/> <p>Note</p> <p>The system hangs if back slash (\) is used in virtual machine names.</p> <hr/>

Supported special characters

The following table provides information about supported special characters for Hyper-V backups.

Table 11 Supported special characters

Special character	Description	Special character	Description
A–Z, a–z, 0–9	Alphanumeric	\$	Dollar
-	Dash	%	Percentage
.	Period)	Right parenthesis
_	Underscore	(Left parenthesis

Table 11 Supported special characters (continued)

Special character	Description	Special character	Description
{ }	Curly brackets	&	Ampersand
+	Plus	^	Carat
=	Equal	@	At sign
~	Tilde	<	Less-than sign
!	Exclamation	>	Greater-than sign
#	Hash	*	Asterisks
	Space	,	Comma
;	Semi colon		Vertical bar
:	Colon	`	Backtick

Unsupported features

The following table provides information about unsupported features for Hyper-V backups.

Table 12 Unsupported features

Feature or functionality	Description
Hyper-V node backups of Windows Server 2012 R2 and 2016 in the same cluster	To perform a backup of Windows Server 2012 R2 and 2016, you must have separate clusters for Windows Server 2012 R2 and 2016. Microsoft provides limited support for server versions of different operating systems in the same cluster, and especially during migration of Windows Server 2012 R2 to Windows Server 2016. Microsoft recommends you not to plan for any backup and recovery activity during an operating system migration. A full backup is recommended before and after migration of operating system.
VHD and VHDX backups when the Microsoft option "Enable virtual hard disk sharing" is enabled	On Windows Server 2012 R2 and 2016, NMM does not support VHD and VHDX backups when the "Enable virtual hard disk sharing" Microsoft option is enabled. Backups ignore the virtual machines with shared VHD or VHDX, and proceed with the rest of the save set. To exclude VHDs and VHDXs that are sharing-enabled from the backup, specify the NSR_EXCLUDE_SUBCOMPONENTS application information variable with the list of VHDs and VHDXs when you configure the client resource. To protect the data on a VHD or VHDX that is sharing-enabled, install the NMM software on the guest virtual machine, and then run the backup inside the guest virtual machine.
Hyper-V virtual machines that contain dynamic disks and vice versa	Microsoft does not support the Hyper-V virtual machines that contain dynamic disks and the Hyper-V virtual machines on dynamic disks. So, NMM also does not support these configurations.
RCT-based backups of Hyper-V virtual machines with a configuration version earlier than 6.2	NMM does not support RCT-based backups of Hyper-V virtual machines with a configuration version earlier than 6.2.

Table 12 Unsupported features (continued)

Feature or functionality	Description
	<p>After you upgrade the operating system from 2012 R2 to 2016, upgrade the virtual machine configuration version to 6.2 or later.</p> <p>To upgrade multiple virtual machines, run the following PowerShell command:</p> <pre>get-vm Update-VMVersion -A</pre>
RCT-based backups of Hyper-V proxy virtual machines in a federated environment	In a federated environment, when you perform an RCT-based backup of a Hyper-V proxy virtual machine by using the Client Backup Configuration wizard, the Available Servers list on the Specify Backup Options page does not display the proxy virtual machines.
Cloning Hyper-V save sets if the "Delete source savesets after clone completes" option in the NetWorker Administration window is selected	In the NetWorker Policy Action Wizard , on the Specify the Clone Options page, if you select Delete source savesets after clone completes , and then clone Hyper-V save sets, the operation fails.
Cloning Hyper-V incremental backups if the source or primary device is either AFTD or CloudBoost	Cloning Hyper-V block based incremental backup save sets fails if the source backup device is either AFTD or CloudBoost.
Backup of a virtual machine that has its VHDs stored on multiple and mixed types of storage	<p>NMM does not support backup of a virtual machine that has its VHDs stored on multiple and mixed types of storage, such as CSV, SMB, standalone, and so on, or a combination of any of these configurations. If the backup includes such multi-storage virtual machines and normal virtual machines, you see the following results:</p> <ul style="list-style-type: none"> • If the backup of normal virtual machines succeeds, the status of the entire backup on NMC appears as succeeded with a green tick mark. However, the backup skips the multi-storage virtual machines. • The Savesets excluded from backup section in the <code>nsrnmmsv.raw</code> log file on the respective node lists the skipped multi-storage virtual machines.
Redirected recovery of a Hyper-V virtual machine from a cluster to an alternate cluster or cluster node in a cross-domain environment	<p>NMM does not support the following redirected recovery scenarios:</p> <ul style="list-style-type: none"> • Recovery of a Hyper-V virtual machine from a cluster to an alternate cluster or cluster node that belongs to a different domain. • Recovery of a Hyper-V virtual machine from a cluster to a Hyper-V standalone server, which is outside the cluster.

Requirements and considerations for Hyper-V virtual machine backups

Before you back up Hyper-V virtual machines, review the following requirements and considerations:

Using NMM 18.2 with NetWorker 8.2.3 or later

The procedure to create a client resource when you use NetWorker server 8.2.3 is different from the procedure to create a client resource when you use NetWorker server 18.2. When you use NMM 18.2 and NetWorker 8.2.3 or later:

- Configure a regular NetWorker backup group instead of configuring a data protection policy. Do not enable the **Snapshot** option in the Group properties page.
- Type `nsrnmmsv.exe` in the **Backup Command** field.

Viewing valid application data save sets

When manually configuring a client resource in the NMC, you are required to type the save sets in the Save Set attribute. The "Configuring a client resource manually by using the NetWorker Management Console" section provides information about how to manually create client resources and the save sets to use for various data types.

To view the list of the application data save sets that is available for backup, open the command prompt on the Hyper-V Server, and then type any of the following commands according to your requirement:

- For a Hyper-V standalone server, run the following command:
`nsrnmmsv -P -A NSR_APP_TYPE=HYPERV`
- For a Hyper-V clustered server, run the following command:
`nsrnmmsv -P -A NSR_FEDERATED_BACKUP=yes -A NSR_APP_TYPE=HYPERV`

Using the correct Integration services components version

Ensure that the IC version that runs inside the virtual machine is the same as the version of Hyper-V on the host. A mismatch in versions may lead to backup failures.

Check the Hyper-V version on the server by starting the **Hyper-V Manager** and then selecting **About Hyper-V Manager** from the **Help** menu.

Check the IC version:

1. In the Device Manager application inside the guest virtual machine, on **System Devices**, select **Device Manager**.
2. Right-click the entry **Microsoft Hyper-V Volume Shadow Copy**.
3. Select **Properties**.
4. Check the IC version on the **Driver** tab.

If the IC version does not match the Hyper-V version, insert the integration services disk by choosing that option under the **Action** menu in the virtual machine console. Install the integration components, and then restart the virtual machine.

Note

On Windows Server 2016, the IC is updated through Windows update.

Adding a virtual machine to an existing incremental backup

The first instance of all virtual machines backups must be a full backup. When you add a virtual machine to a Hyper-V Server or a failover cluster that is scheduled for writer-level incremental backups, NMM determines that the newly added virtual machine is not enabled for incremental backups because a full backup of the virtual machine does not exist. During the next incremental backup of the writer-level incremental save set,

NMM splits the virtual machines into two sets, creates a snapshot for each set, and then backs up each set separately.

- Set of virtual machines for which full backups are available, only incremental backup is performed.
- Set of virtual machines that contains the newly added virtual machine, first a full backup is performed, followed by incremental backup. The backup for this set of virtual machines might take longer to complete.

Special requirements and APIs to support VSS-based backup applications for Windows Server 2012, 2012 R2, and 2016 Hyper-V CSVs

For Windows Server 2012, 2012 R2, and 2016, Microsoft has released several special requirements and APIs to support VSS-based backup applications. A backup application can back up all the CSVs from a single node. CSVs are not required to be put in I/O Redirection Mode as with previous Microsoft OS releases, and CSVs can be backed up in parallel.

The Windows Server 2012, 2012 R2, and 2016 interoperability backup application is CSV-aware because the CSV writer metadata information must be updated to its component name by querying the primary server for CSV resources.

In Windows Server 2012, 2012 R2, and 2016, the new CSV VSS writer can report the backup components on the behalf of a different cluster node. This CSV VSS Writer can also take the snapshots of volumes on the remote node. These features enable NMM to back up not only the local image of a Hyper-V virtual machine, but also to back up the image that is on a remote node. This allows for more configuration options. For example, you can dedicate a single node to back up the cluster.

Creating multiple snapshots of CSV volumes in Hyper-V scale-out environments

NMM 9.0 and earlier uses the "single snapshot" feature as the default method for protecting Hyper-V CSV environments. The Client Backup Configuration wizard enables the user to create client resources for the single snapshot feature.

In Hyper-V scale-out environments, the CSV Writer might fail to create a single snapshot of all the virtual machines. In single snapshot, the CSV Writer from the cluster owner node performs a snapshot of all virtual machines running on different nodes of the cluster.

In NMM 9.0.1 and later, a number of snapshots can be created instead of a single snapshot of the entire CSV environment. To achieve this, the virtual machines are grouped by the CSV volumes on which they are present. The CSV Writer then creates a snapshot for each group of virtual machines.

Estimating the required additional primary storage for differencing disks

Hyper-V VSS Writer uses differencing disks to capture changed blocks between backups. A differencing disk is a VHD that contains changes for another VHD or the guest operating system. Differencing disks are stored in the same subdirectory as the parent VHDx for the virtual machine. This location is not configurable.

Incremental changes since the last checkpoint or backup are written to a new differencing disk. This differencing disk gets merged to its parent disk after the backup completes. When a snapshot occurs as part of a backup, a new differencing disk is created. This new differencing disk receives all the writes until the next backup starts. The backup saves the differencing disk as part of an incremental backup. If you create checkpoints, then multiple differencing disks can be present on the system. If you do not create checkpoints, there is one differencing disk present along with parent disk.

To estimate the required additional primary storage for differencing disks, consider the rate of changes happening inside the virtual machine. As a virtual machine grows,

it requires more storage space on the primary disk. NMM logs how much data it has backed up for each backup.

Removing and merging stale recovery checkpoints

Stale checkpoints are orphaned checkpoints that are not deleted due to some failures in the backup process. For example, stale checkpoints may be created from a failed backup or when the Windows Server leaves a virtual machine in a locked or backup state even after backup is complete.

To remove and merge stale recovery checkpoints, use the following command at the Command Interface: `nsrnmmsv -H -A NSR_FEDERATED_BACKUP=<yes/no> -A NSR_MERGE_STALE_CHECKPOINTS=<yes>`

The `NSR_MERGE_STALE_CHECKPOINTS=yes` attribute with `-H` option merges all stale checkpoints. The `nsrnmmsv` command checks each virtual machine in the Hyper-V host or Hyper-V cluster for stale checkpoints and merges them, except for the last recovery checkpoint. The next backup is incremental for the virtual machine that contains the last recovery checkpoint. The stale recovery checkpoints are removed from clustered and non-clustered virtual machines by using `NSR_FEDERATED_BACKUP=yes/no` option.

The "Configuring a client resource manually by using the NetWorker Management Console" section provides details.

Excluding VHDs and VHDXs from Hyper-V VSS and RCT backups in standalone, federated, and SMB configurations

You can use the Application Information `NSR_EXCLUDE_SUBCOMPONENTS` variable during client resource configuration to exclude VHDs and VHDXs from a Hyper-V backup in standalone, federated, and SMB configurations. The "Configuring a client resource manually by using the NetWorker Management Console" section provides details.

Note

You cannot recover virtual machines from which the VHD and VHDX are excluded during backup by using the NetWorker User for Microsoft GUI. Perform a flat file recovery of the VHD and VHDX, and then attach the recovered VHD and VHDX files to the virtual machine.

Merging overgrowing AVHDX and AVHD files when you perform VSS backups

You may face the issue of the AVHDX file growing fast during incremental backups. AVHDX is a differencing file created during incremental backup, and keeps growing to the maximum size of its parent VHD till the next incremental backup. An oversized AVHDX file is often created as a result of longer backup frequency, which leads to the consumption of a large amount of space on production storage. You are also unable to resize the base VHD due to the presence of the AVHDX file.

There are two solutions to prevent the AVHDX file from growing fast :

- Backup the virtual machines more frequently to prevent the AVHDX file from growing to its parent VHD size. For example, backup the virtual machines twice a day if the current backup frequency is once in a day, or twice a week if the current backup frequency is once a week.
- Merge and remove the overgrowing AVHDX file by using the following command at the Command Line Interface:
 - Use `NSR_MERGE_CHECKPOINTS_THRESHOLD =<%value>` attribute and `-H` option with the `nsrnmmsv` command, where %value is the threshold percentage value set for the AVHDX file.

- Based on the threshold percentage, the overgrown recovery checkpoints are removed from clustered and non-clustered virtual machines with the `NSR_FEDERATED_BACKUP=<yes/no>` option.

For example:

- For clustered virtual machines: `nsrnmmsv -H -A`
`NSR_FEDERATED_BACKUP=yes -A`
`NSR_MERGE_CHECKPOINTS_THRESHOLD=20`
- For non-clustered virtual machines: `nsrnmmsv -H -A`
`NSR_FEDERATED_BACKUP=no -A`
`NSR_MERGE_CHECKPOINTS_THRESHOLD=20`

The command checks all the virtual machines on a Hyper-V host or Hyper-V cluster for the current file size of a recovery checkpoint (AVHD/AVHDX) and merges all recovery checkpoints if the current file size is exceeding threshold percentage. The subsequent backup for the virtual machine will be a full backup.

The "Configuring a client resource manually by using the NetWorker Management Console" section provides details.

Creating checkpoints for a virtual machine to capture the virtual machine state

Hyper-V users can use the Hyper-V management interface to create checkpoints for a virtual machine to capture the virtual machine state at strategic points of their own choice. Hyper-V checkpoints are of two types: Backup checkpoints and user created checkpoints. Both these types of checkpoints create differencing disks that provide roll back capabilities to these strategic points.

When NMM performs VSS full or incremental backup of the Hyper-V virtual machine, the differencing disks that are created for the backup checkpoints are merged by NMM on the backup media. The Hyper-V management interface continues to display the checkpoints even after the backup is complete. Because the differencing disks are merged on the backup media after backup is complete, the checkpoints are not displayed in Hyper-V management interface after restore.

Current versions of Hyper-V do not recommend user-created checkpoints for production virtual machines.

Note

When you disable CBT, the differencing disks that are created for the virtual machine checkpoints are not merged on the backup media after the backup.

VSS backups

Volume Shadow Copy Service (VSS) is a Windows service that enables administrators to ensure backup transaction consistency. It consistently coordinates the whole process of data copy and concurrently ensures that applications remain online. It does not create a copy of data, but saves the current data state and ensures that this state is consistent within a point in time.

If you install the backup (volume snapshot) integration services on the guest operating system, a VSS requester is also installed. The VSS requester enables VSS writers in the guest operating system to participate in the backup of the virtual machine.

Hyper-V uses one of the following VSS methods to back up a virtual machine:

- **Saved State:** Before the process takes a snapshot, the process puts the virtual machine to back up into a saved state. The process takes snapshots of the volumes, and then puts the virtual machine into the previous state.

- **Child VM Snapshot:** The process uses VSS that is inside the child virtual machine to participate in the backup.

VSS backup consistency types with NMM

The Hyper-V Writer in the management operating system determines if the backup image is application-consistent or crash-consistent. You do not need to select these backup types when performing scheduled backups of virtual machines with NMM.

A crash-consistent backup is performed when Microsoft VSS Integration component (IC) is not installed on the virtual machine. In a crash-consistent backup, the virtual machine is paused before shadow copy creation and resumed after the shadow is created.

An application-consistent backup is performed when IC is installed on the virtual machine. Ensure that the virtual machine is online and VSS-capable. In an application-consistent backup, the IC runs in-guest and freezes the operating system and all application states.

The Microsoft Technet website provides more information about Microsoft application-consistent or crash-consistent backups.

NOTICE

Create and maintain separate policies for VSS and RCT backups. Configure and schedule either VSS backups or RCT backups at a time. Do not mix VSS backups with RCT backups.

Configuring backups

When configuring backups, the backup tasks differ depending on the items to back up.

The following table describes the backup tasks you must perform when using NetWorker server 18.2 and NMM 18.2 to back up Hyper-V parent and virtual machines. The *NetWorker Administration Guide* provides detailed information about how to perform the tasks in the table.

Table 13 Backup tasks for Hyper-V

Items to back up	Backup tasks to perform
<p>On the server</p> <p>The Hyper-V role can coexist with other Microsoft applications, such as:</p> <ul style="list-style-type: none"> • SQL Server • SharePoint Server • Exchange Server • Windows Server Cluster 	<ol style="list-style-type: none"> 1. Configure the NetWorker devices for backup storage. 2. Configure a backup group. 3. Configure one or more client resources for each client. 4. Configure a data protection policy for scheduled backups, including selecting a group, policy, policy workflow, and backup action. 5. Configure required NetWorker privileges. 6. Configure backup proxies.
<p>Hyper-V on the server</p> <p>Hyper-V virtual machines and Host Component file</p>	<ol style="list-style-type: none"> 1. Configure the backup storage resources. 2. Configure a backup group. 3. Configure one or more client resources for each client. 4. Configure a data protection policy for scheduled backups, including selecting a backup group, policy, policy workflow, and backup action.

Table 13 Backup tasks for Hyper-V (continued)

Items to back up	Backup tasks to perform
Hyper-V virtual machine applications Microsoft application data, such as: <ul style="list-style-type: none"> SQL Server SharePoint Server Exchange Server Windows Server Cluster 	Install NMM on the virtual machine operating system and configure application backups with NMM installed within the virtual machine operating system: <ul style="list-style-type: none"> Configure Windows application backups. Configure Windows Server cluster backups. Specific instructions for the Microsoft application are provided in the application-specific user guides.

Note

The "Using NMM 18.2 with NetWorker 8.2.3 or later" section provides the considerations to follow when using NetWorker 8.2.3 and later and NMM 18.2.

The following table lists the application information variables that NMM supports when you manually configure client resources:

Table 14 Supported application information variables

Client or workflow	Hyper-V Servers 2012 and 2012 R2 VSS	Hyper-V Server 2016 VSS	Hyper-V Server 2016 RCT
Standalone		NSR_RCT_BACKUP=No	NSR_RCT_BACKUP=Yes
Clustered	NSR_FEDERATED_BACKUP=Yes NSR_FEDERATED_TYPE=CSV/SMB NSR_MOVE_CSV_OWNERSHIP=Yes/No	NSR_RCT_BACKUP=No NSR_FEDERATED_BACKUP=Yes NSR_FEDERATED_TYPE=CSV/SMB NSR_MOVE_CSV_OWNERSHIP=Yes/No	NSR_RCT_BACKUP=Yes NSR_FEDERATED_BACKUP=Yes

Creating a client resource for a VSS-based backup by using the Client Backup Configuration wizard

The Client Backup Configuration wizard for Hyper-V simplifies configuration of scheduled backups for Hyper-V servers. The NMM client must be installed on all Hyper-V cluster nodes for the Client Backup Configuration wizard to function correctly. The wizard automatically configures Hyper-V save sets, backup commands, application information variables, and backup options. Use the wizard to configure client resources for stand-alone and federated environments.

Before you begin

Before you use the wizard, review the following requirements:

- Ensure that the NetWorker server host is listed in the `servers` file on the client computer.

- Ensure that the communication between the NMC server, NetWorker client, and NetWorker server uses `nsrauth` strong authentication.
- Ensure that the user who runs the wizard meets the following requirements:
 - Root (UNIX) or Administrator (Windows) privileges.
 - A member of a User Group on the NetWorker server that has Configure NetWorker privileges.
- Ensure that multiple wizard hosts are not trying to access the same client computer simultaneously.

Note

If you use NetWorker server 8.2.3 or later and NMM 18.2:

- For the Client Backup Configuration wizard to properly function, ensure that JRE is installed on the host, where NMC is used. The NMC for NetWorker 8.2.3 supports JRE 7, and the NMC for NMM 18.2 supports JRE 8 and 9.
 - The procedure used to create a client resource when you use NetWorker server 8.2.3 and later is different from the procedure used to create a client resource when using NetWorker server 18.2.
The "Scheduled Backup" chapter in the *NetWorker Module for Microsoft Release Administration Guide* provides more information about editing client resources that were created by using NMM 8.2.x, and information about the NMC bulk edit feature.
 - Before you start the NMM 18.2 Client Backup Configuration wizard to modify a client resource that was created by using NMM 8.2.x, ensure that the **Snapshot** option of the NetWorker group that this client resource belongs to is clear. If the **Snapshot** option is selected, you cannot select the NetWorker group in the wizard, and you are prompted to create or select another group.
 - Do not use NetWorker strong authentication (`nsrauth`) to communicate with other hosts, such as NMC server, NetWorker server, and NetWorker client. The NetWorker 18.2 Administration Guide provides more information about `nsrauth`.
-

Procedure

1. Use NMC to connect to the NetWorker server.
2. In the **NetWorker Administration** window, click **Protection**.
3. In the expanded left panel, right-click **Clients**, and then select **New Client Wizard**.

The **Client Backup Configuration** wizard appears.

4. On the **Specify Client Information** page:
 - a. In the **Client Name** field, type either the hostname or the Fully Qualified Domain Name (FQDN) of the client.
For federated backups, type the cluster server name
For non-federated backups, type the stand-alone Hyper-V server name.
-

Note

Do not type the IP address of the client or the server.

- b. Optionally, in the **Comment** box, type a description of the client.

If you are creating multiple client resources for the same NetWorker client host, specify this field to differentiate the purpose of each resource.

- c. In the **Tag** box, type one or more tags to identify this Client resource for the creation of dynamic client groups for data protection policies.

Place each entry on a separate line.

- d. In the **Type** box, select **Traditional NetWorker client**.

- e. Optionally, from the **Group** list, select a group for the Client resource.

The group to which the client belongs determines the workflow that is used to back up the client.

Note

You can also assign the client to one or more groups after you create the Client resource.

- f. Click **Next**.

5. On the **Specify the Backup Configuration Type** page, select **Hyper-V Server** and click **Next**.

6. On the **Select the NetWorker Client Properties** page:

- a. Select the priority level in the **Priority** field.

- b. Select the level of parallelism in the **Parallelism** field.

- c. Leave the **Remote Access** field empty.

- d. Select the device type from the **Data Domain Interface** list.

- e. Select the **Parallel Save Streams** option to enable multiple save streams for each save set during backup.

- f. Click **Next**.

Note

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

7. On the **Select the Hyper-V Backup Objects** page, to exclude a save set from the backup, perform the following steps:

- a. Select the top level save set.

- b. Select the **Exclude Component List** option that appears at the bottom of the page.

- c. Clear the save set that you want to exclude from the backup.

- d. Click **Next**.

According to your setup — standalone or federated, the **Specify Backup Options** page accordingly appears as shown in one of the following figures:

Figure 8 Specify Backup Options page for a standalone setup

Figure 9 Specify Backup Options page for a clustered setup

8. On the **Specify Backup Options** page:

a. To perform federated backups, specify the following fields:

- In the **Remote username** field, type the clustered server username that has the administrator privileges.
- In the **Password** field, specify the corresponding password for the username that you have specified in the **Remote username** field.

For federated setups, the wizard creates empty (placeholder) client resources for all the nodes that you do not use to perform backups. The wizard also creates a client resource with the cluster name, and specifies the *NSR_FEDERATED_BACKUP=yes* variable.

For standalone setups, the wizard creates a client resource with the physical name by using the selected save sets.

b. Under **Backup Type**, select **VSS-based**.

This type of backup applies to Hyper-V Server 2012 and later.

To perform federated backups:

- For CSV configurations:
 - Select **Back up VMs over CSV**. By default, federated backups are performed as CSV backups.
 - To enable or disable CSV ownership change, select **Move CSV ownership**. For optimal performance of multi-proxy backups, you can change the CSV ownership. To enable or disable CSV ownership changes, you must have created the client resource.
 - Do not select CSV and non-CSV virtual machines in one backup instance. Otherwise, the backup fails.
- For SMB configurations, select **Back up VMs over SMB3**.

c. Select the backup options:

- **Perform partial backup**: This option is selected by default. It enables backing up other virtual machines when the snapshot operation fails for some of the virtual machines that you have selected for the backup. The Hyper-V VSS Writer reports the virtual machines that are not backed up.
- **Disable incremental backup**: This option disables incremental backups so that only full backups are performed.

d. (Optional) To distribute the backup workload across physical and virtual servers, select proxy servers for the Preferred Server Order List (PSOL):

- To add a server from the **Available Servers** list to the **Proxy Servers** list, select the server, and then click the right arrow.
Repeat this step for each server that you want to add to the **Proxy Servers** list.

To remove a server from the **Proxy Servers** list, select the server, and then click the left arrow.
- To add all the servers from the **Available Servers** list to the **Proxy Servers** list at a time, click the right double arrow.
To remove all the servers from the **Proxy Servers** list, click the left double arrow.
- To change the order or position of the servers in the **Proxy Servers** list, click either the up arrow or the down arrow.

If a server is unavailable, NMM performs the backup from the node, to which the cluster server name resolves. The node can be a proxy virtual machine.

Note

Ensure that NetWorker client and NMM are installed on each cluster node.

e. Click **Next**.

9. On the **Backup Configuration Summary** page, perform one of the following steps:

- Click **Back** to go to the previous pages to modify the configuration.
 - Click **Create** to configure the client resources.
The **Client Configuration Results** page appears. This page provides details about the client resources that have been created.
10. To modify the previous configuration, in NMC, select the client, right-click, and select **Modify Client Wizard**.
 11. To verify the details for the client, in the **NetWorker Administration** window, right-click the client, and then select **Modify Client Properties**.
 12. (Optional) To modify remote user access:
 - a. In the **NetWorker Administration** window, right-click the client, and then select **Modify Client Properties**.
 - b. On the **Globals (2 of 2)** tab, type the hostname of the proxy client in the **Remote Access** field.

If the NMM client is part of a cluster, type the names of the physical nodes of the cluster in the **Remote Access** field.

Manually creating a client resource for a VSS-based backup by using the Client Properties dialog box

You can manually create a Hyper-V client resource by using the **Client Properties** dialog box.

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
2. In the expanded left panel, select **Clients**.
3. Right-click **Clients** and select **New**.
The **Client Properties** dialog box appears.
4. On the **General** tab:
 - a. In the **Name** field, type either the hostname or the FQDN of the client.
The client must be a fully qualified host to be a NetWorker client.

Note

In this field, do not type the IP address of the client.

- b. In the **Comment** field, type a description. If you are creating multiple client resources for the same NetWorker client host computer, specify this field to differentiate the purpose of each resource.
- c. In the **Save Set** field, specify the components that you want to back up, each in a line. The following table lists the type of backup data and the corresponding save set syntaxes that you must specify in this field.

Table 15 Hyper-V save set syntax

Type of backup data	Save set syntax
Hyper-V Manager The Hyper-V Host Component file and each virtual machine.	APPLICATIONS:\Microsoft Hyper-V

Table 15 Hyper-V save set syntax (continued)

Type of backup data	Save set syntax
	The Hyper-V Writer does not support offline backup of the configuration file.
Hyper-V Host Component file There is one configuration file in the Hyper-V Manager installation. This file lists the Hyper-V settings for the host operating system and the guest operating systems.	APPLICATIONS:\Microsoft Hyper-V\Host Component The Hyper-V Writer does not support offline backup of the configuration file. You cannot use the APPLICATIONS:\Microsoft Hyper-V\Host Component save set in a proxy backup group.
Hyper-V virtual machine There are usually multiple virtual machines on the host operating system.	APPLICATIONS:\Microsoft Hyper-V\virtual_machine_name Child pertains or virtual machines can be in a proxy backup group.

To perform a CSV-level federated backup, you must set the application information variables for the cluster virtual server and the client resource.

- d. Select the appropriate option in the **Protection group** list field.
5. Click the **Apps & Modules** tab and do the following:
 - a. In the **Remote user** and **Password** fields respectively, type the domain administrator username and password.

For guest virtual machines that are hosted over SMB 3.0 and CSV, the backup fails if you do not provide the domain administrator credentials.
 - b. In the **Backup command** field, type the following backup command:

`nsrnmmsv.exe`
 - c. In the **Application Information** field, specify the application information variables that you require to perform the backup. The following table lists the variables that you can specify in this field.

Table 16 Hyper-V application information variables

Variable name	Description	Values
NSR_RCT_BACKUP	Performs RCT-based backups.	<ul style="list-style-type: none"> Yes No
NSR_FEDERATED_BACKUP	Marks the backup for CSV and SMB recovery.	Yes
NSR_FEDERATED_TYPE	Specifies whether the backup is federated CSV or SMB.	<ul style="list-style-type: none"> csv smb <p>If this variable is not specified, NMM applies the CSV value by default.</p>
NSR_FEDERATED_PSOL	Optional. Distributes the backup workload across physical and virtual servers in the PSOL. If a server is not available or down, then	Type a comma-separated list of the server names. For example: <code>NSR_FEDERATED_PSOL=FQDN_of_s</code>

Table 16 Hyper-V application information variables (continued)

Variable name	Description	Values
	<p>NMM performs the backup from the node to which the cluster server name resolves, including virtual proxies.</p> <p>RCT backups support only cluster nodes, but not virtual machines as proxy servers.</p> <p>If you specify this application information variable with a list of servers (cluster nodes) as values, and perform an RCT backup, the backup operation performs rollover on all the specified servers.</p> <p>If you do not specify this application information variable, and perform an RCT backup, the backup operation performs rollover of the virtual machines from the respective nodes, on which the virtual machines are present.</p> <hr/> <p>Note</p> <p>Ensure that NetWorker client and NMM are installed on each cluster node.</p>	<p>erver1, FQDN_of_server2, FQDN_of_server3,...</p> <hr/> <p>Note</p> <p>A Hyper-V CSV backup fails if the same host name is added twice in PSOL list.</p>
NSR_EXCLUDE_SMB	Optional. Excludes virtual machines that have data that is stored on SMB file servers. By default, SMB virtual machines are in the writer level backup.	Yes
NSR_VSS_FULL_BACKUP	<p>Optional. The default value is Yes.</p> <p>If the value is Yes, NMM initially performs a full backup and later performs incremental backups.</p> <p>If the value is No, then NMM performs a copy full backup and does not merge the recovery snapshot with the base VHDX.</p>	<ul style="list-style-type: none"> • Yes • No
NSR_VM_COPY_BACKUP	<p>Optional. If the value is Yes, NMM performs a VSS copy backup in guest virtual machines, which means there is no log truncation for applications running in guest virtual machines.</p> <p>If the value is No, NMM performs a VSS full backup in guest virtual machines.</p> <p>By default, all Hyper-V virtual machine backups are VSS copy type.</p>	<ul style="list-style-type: none"> • Yes • No
NSR_PARTIAL_BACKUP	Optional. Supports partial backup failure. If snapshot creation fails for some virtual machines, the backup continues for other virtual machines. The Hyper-V VSS Writer reports the virtual machines that failed to backup. The default value is Yes.	<ul style="list-style-type: none"> • Yes • No

Table 16 Hyper-V application information variables (continued)

Variable name	Description	Values
NSR_EXCLUDE_COMPONENTS	Optional. Excludes a virtual machine from the backup. Specify the writer level save set and the components to exclude from the backup. NMM logs the excluded components in the <code>nsrnmmsv.raw</code> log file for references.	Type a comma-separated list of the server names. For example: <code>NSR_EXCLUDE_COMPONENTS=VM1, VM2, VM3</code>
NSR_EXCLUDE_SUBCOMPONENTS	Excludes VHDs and VHDXs from a Hyper-V backup in standalone, federated, and SMB configurations. Note Ensure that the VHDs and VHDXs do not have special characters comma (,), colon (:), equal to (=), and vertical bar () in their names because the special characters are used as separators.	Use comma (,), colon (:), and vertical bar () as separators to exclude VHDs and VHDXs from the backup. For example: <code>NSR_EXCLUDE_SUBCOMPONENTS = VM1:VHD1.vhdx, VHD2.vhdx VM2:VHD3.vhdx</code> <code>NSR_EXCLUDE_SUBCOMPONENTS = VM1:VHD1.vhdx, VHD2.vhdx</code> <code>NSR_EXCLUDE_SUBCOMPONENTS =VM2:VHD3.vhdx</code> where, <ul style="list-style-type: none"> The colon (:) is used as the separator between the virtual machine name and VHD name. The comma (,) is used as the separator between two VHD names. The vertical bBar () is used as the separator between two VMs.
NSR_MAX_CSV_PER_SNAPSHOT	Creates multiple snapshots for scale-out Hyper-V environments.	The maximum value of the variable is the number of CSVs in the scale-out environment.
NSR_DISABLE_INCREMENTAL	Optional. By default, when Hyper-V backups are performed by using NMM, the first instance of the backup is always full and subsequent backups are incremental. The <code>NSR_DISABLE_INCREMENTAL=yes</code> variable disables incremental backups and enables only full backups. When incremental backups are disabled using this parameter, the recovery snapshots, that is .AVHDX files, are merged prior to the full backup.	Use the value is Yes to disable change block tracking (CBT) and perform only full backups. Use the value is No to enable change block tracking (CBT) and perform both full and incremental backups. You have the option to enable CBT or disable CBT at any time depending on your business needs.

- Click the **Globals (1 of 2)** tab and, in the **Aliases** field, type the NETBIOS name for the client.

NMM client uses the host computer NETBIOS or short name when connecting to the NetWorker server to browse backups. If the NETBIOS name is not found, NMM cannot display backups.

- To set up remote access, click the **Globals (2 of 2)** tab and do one of the following:

- If you are setting up a proxy client for the NMM client, type the hostname of the proxy client in the **Remote Access** field.
 - If the NMM client is part of a cluster, type the names of the physical nodes of the cluster in the **Remote Access** field.
8. Click **OK**.

Configuring NMM 8.2.x client resources after upgrading NMM from 8.2.x to 9.1 or later

After you upgrade NMM from 8.2.x to 9.1 or later, to use the same 8.2.x client resources to perform backups, you must perform the following changes to all the 8.2.x client resources:

Procedure

1. Either ensure that the **Snapshot** option of the NetWorker group that a client resource belongs to is clear or create a new group and ensure that the **Snapshot** option is clear.
2. In the **NetWorker Administration** window, click **Protection**.
3. Either right-click **Clients** in the navigation tree or right-click the required client in the **Clients** table.
4. Select **Modify Client Properties**.
5. In the **Backup Command** field, delete the `nsrsnap_vss_save` command, and then type the `nsrnmmsv` command.

Perform this step for all client resources.

6. To enable Hyper-V virtual machine full backups, in the **Application Information** field, type the following variables each in a line:

- `NSR_VSS_FULL_BACKUP=yes`
- `NSR_RCT_BACKUP=no`

To perform Hyper-V virtual machine copy full backups, in the **Application Information** field, perform one of the following steps:

- Do not specify the `NSR_VSS_FULL_BACKUP` variable.
- Type `NSR_VSS_FULL_BACKUP=no`.

7. Make other changes according to your requirement.
8. Click **OK**.

NOTICE

The "Scheduled Backup" chapter in the *NetWorker Module for Microsoft Release Administration Guide* provides more information about editing client resources that were created by using NMM 8.2.x, and information about the NMC bulk edit feature.

Improving performance of VSS CSV backups by using multiple cluster nodes as proxies

In larger Hyper-V environments, you can improve performance by scaling out the Hyper-V CSV backups to multiple cluster nodes or proxies.

When you create multiple proxies, NMM should be installed on all cluster nodes and virtual proxies. There must be client resources for all the cluster nodes, virtual proxies, and cluster alias.

You add secondary roles, which must be physical or virtual cluster nodes. You specify the proxy host by setting `NSR_FEDERATED_PSOL` in the Application Information for the client resource of the cluster, or by using the Preferred Server Order List (PSOL) in the Client Configuration Wizard. The PSOL distributes the backup workload across all servers in the PSOL. You schedule the backups against the cluster alias, and the primary role runs on the cluster node that owns the cluster alias. The recovery process for virtual machines that are backed up as part of a multi-proxy setup is the same as the recovery process for traditional backups.

In a multi-proxy setup, you can select multiple cluster nodes to act as proxy nodes to perform parallel backups on all proxy nodes. An NMM CSV algorithm is used to intelligently reassign virtual machines to proxies and, if the `NSR_MOVE_CSV_OWNERSHIP` application information attribute is set to `yes`, assign CSVs to proxies. The backup load is evenly split between multiple nodes. All proxy nodes perform backups in parallel, significantly increasing backup performance compared with single proxy backups in a normal distributed CSV environment. You can add or remove proxy nodes as needed.

NMM takes a single snapshot of the Windows Server 2012, 2012 R2, and 2016 Hyper-V cluster from the controlling node. NMM mounts the snapshot on the primary node and shares it among the secondary nodes. This makes the snapshot process faster.

If all the selected proxy nodes are unavailable, then NMM performs the backup on the node to which the cluster alias is resolved.

Hyper-V intelligent proxies

NMM allows you to use virtual machines on a Hyper-V cluster as proxy nodes for CSV backups. NMM intelligently reassigns and load balances the virtual machines in the backup to the selected virtual proxy nodes. All virtual proxy nodes will perform backups in parallel, thus increasing backup performance. These virtual proxies must be highly available cluster Hyper-V virtual machines, and they must be connected to the same domain as the physical nodes. A mix of physical and virtual machines is supported. NMM automatically excludes the virtual machines for the virtual proxies from backups.

The NMM controller node is the cluster node on which the cluster alias is active. All servers in the PSOL are treated as proxies. Before starting the distribution of CSVs, the controller validates all the servers in the PSOL and excludes any unavailable nodes from the PSOL. It is a best practice to configure more proxies in the PSOL so that a spare node will always be available. Because NMM uses the cluster alias to schedule cluster backups, NMM is immune to node failures. If the currently active node fails, the cluster alias moves to a different node.

The CSV ownership distribution algorithm logic tries to distribute the CSVs among the proxies equally. For example, in an 8-node cluster with 4 CSVs:

- If two servers are in the PSOL, then NMM backs up two CSVs per proxy.
- If four servers are in the PSOL, then NMM moves the CSV ownership for maximum performance and backs up one CSV per proxy.

- If five servers are in the PSOL, then NMM backs up one CSV per proxy for four of the servers.

Multi-proxy backup components

In addition to the existing single proxy client components, NMM uses the following software components in multi-proxy backups:

- Main proxy client—NMM schedules and browses backups against the cluster server name (cluster alias). The NMM process that starts on the active node of the cluster alias is the main proxy client. The main proxy client node acts as the primary node in the cluster.
- Client software—You must install the NetWorker and NMM client software on all secondary proxy nodes. This ensures tolerance for node failures, because the cluster alias ownership can change.
- CSV ownership distribution algorithm—When the `NSR_MOVE_CSV_OWNERSHIP` value is `yes`, NMM changes the coordinator node of the CSV.
 - The CSV coordinator node owns the storage stack for the CSV. Although other nodes might read and write files on the CSV through the SCSI stack, all metadata operations go to the coordinator node.
 - The CSV coordinator node also owns the VSS software shadow copy for the CSV. In this case, I/O on the coordinator node is local, but I/O of the shadow copy volume from other nodes is redirected to the coordinator node over the network.
 - When CSV ownership moves, the software shadow copies for that CSV also move.
 - NMM evenly distributes the backup load across the proxy nodes you select, and the proxy nodes perform backups in parallel to maximize the backup performance. Also, if you want to have an optimized I/O path to the CSV shadow copy for the proxy nodes, select to move CSV ownership.
- Application Information attributes—You can configure Hyper-V CSV options by using the Client Configuration wizard or by using the NetWorker Management Console to edit the client resource for the cluster alias. To enable multi-proxy backups by using the NetWorker Management Console, add the following Application Information attributes on the main proxy client:
 - `NSR_FEDERATED_PSOL`—Enables multi-proxy backups and distributes the backup workload across all servers in the PSOL.
For example:

```
NSR_FEDERATED_PSOL=server1, server2, server 3
where server1, server2, and server3 acts as proxy servers.
```

Note

If you do not specify `NSR_FEDERATED_PSOL`, NMM performs the backup from the current active node and substitutes the cluster master node as the proxy node. If you specify values for `NSR_FEDERATED_PSOL`, NMM performs backups from all the valid, available nodes in the list. If the number of nodes is greater than the number of CSVs, NMM excludes the nodes that exceed this number.

- `NSR_MOVE_CSV_OWNERSHIP`—Uses the CSV ownership distribution algorithm to allow or disallow CSV ownership change during multi-proxy backups for best backup performance. After you initially create the client resource, you can allow or disallow CSV ownership changes as needed.
For example:

`NSR_MOVE_CSV_OWNERSHIP=Yes`

`NSR_MOVE_CSV_OWNERSHIP=No`

The default value is `Yes`. If you set the value of this attribute to `Yes`, then NMM changes the CSV ownership. If you set this attribute value to `No`, then NMM does not change the CSV ownership.

Note

To achieve the highest levels of performance with software snapshots on a CSV, it is recommended to scale up the bandwidth of intra-node communication.

Configuring multi-proxy backups in an SMB configuration

To achieve better backup performance for Hyper-V over SMB Cluster environments on Windows Server 2012 R2, you can use parallel federated, multiple proxy-based backups.

In a parallel federated multi-proxy architecture, you can select multiple SMB cluster nodes to act as proxy nodes. These nodes perform parallel VSS snapshots and parallel federated data backups on all the selected proxy nodes. The backup load is evenly distributed between all proxy nodes (as much as possible) to perform backups in parallel, thus significantly increasing backup performance compared to the existing "federated or single proxy" backup method. NMM also supports virtual machines running on a Hyper-V cluster as proxy nodes for performing Hyper-V over SMB cluster backups on Windows Server 2012 R2.

To use this backup method, add `NSR_FEDERATED_BACKUP=yes`, `NSR_FEDERATED_TYPE=SMB`, and `NSR_FEDERATED_PSOL=node1, node2, node3` to the client resource of the cluster.

Best practices for configuring multi-proxy backups

You can improve multi-proxy backup performance by following best practices for configuring and allocating Hyper-V proxies.

The following section describes the components that affect multi-proxy backup performance. This section also describes best practices for configuring these components to achieve best backup performance.

Load balance virtual machine data on CSVs

To attain maximum backup performance, load balance virtual machine data on all the available CSVs as much as possible.

NMM performs load balancing by running a correlational factor of the number of virtual machines residing on the CSVs and the common size share. NMM calculates a 'set of CSV' for a 'set of virtual machine' whose maximum common share resides on those CSVs.

NMM moves this set of CSV to one proxy and backs up the set of corresponding virtual machines from that node. Maintaining proper CSV load balance results in fewer CSV ownership changes.

Allocate the number of proxy nodes and CSVs

When determining how many proxy nodes to use, you should allocate the maximum number of proxy nodes to gain maximum backup performance. Increasing the number

of proxy nodes can improve backup performance. However, this maximum number of proxy nodes should be less than or equal to the maximum number of CSVs.

For best performance, the number of CSVs should be multiples of the number of proxy nodes. Each virtual machine should be contained within a single CSV only, rather than distributed across multiple CSVs.

Allow CSV ownership change

To achieve maximum backup performance, change CSV ownership to the nodes with less CSV data. CSV ownership change allows the NMM CSV algorithm to intelligently change CSV ownership to the proxy nodes you selected, thus correctly load balancing the backup data. NMM performs the data split per CSV, not per virtual machine.

Select proxy nodes with good system resources

Backups can be an intensive operation on system resources. To attain the best backup performance, select cluster nodes with maximum resources available as proxy nodes. Select nodes with minimal live data movement, so that the backup operation does not disturb the day-to-day production activities.

RCT backups

Microsoft Windows Hyper-V Server 2016 has built-in system of Resilient Change Tracking (RCT). This RCT feature of Hyper-V 2016 enables easier and faster incremental backups. With this feature, you can back up Hyper-V virtual machines without using VSS framework.

RCT backups ensure better results in terms of scalability, performance, and reliability.

RCT backups provides the following benefits:

- Storage-efficient backups
- No I/O overhead
- In-built Change Block Tracking (CBT) in memory and on disk
- No requirement for additional storage for incremental backups
- No requirement for VSS framework, CSV Writers and Providers, and File Server Providers
- No requirement for CSV-based or SMB-based workflow
- Supports VHD or VHDX that is used for guest clusters

NMM supports image-level full and incremental RCT backups of Hyper-V 2016 standalone and federated virtual machines to the AFTD, Data Domain, or CloudBoost storage device that you configure.

Before you configure RCT backups, consider the following notes:

- RCT backups do not include VSS writers because they do not use VSS framework.
- Create and maintain separate policies and groups for RCT and VSS backups. Configure and schedule either RCT backups or VSS backups at a time. Do not mix RCT backups with VSS backups.
- RCT backups of the virtual machines that contain any user checkpoints or recovery checkpoints fail. Before you back up such virtual machines, merge the checkpoints by running the following PowerShell command:

```
Get-VM -name <virtual_machine_name> -ComputerName <node_name> | Get-
VMSnapshot | Remove-VMSnapshot
```

RCT backups of the virtual machines that do not contain any checkpoints proceed.

- When you switch from the VSS-based backup to the RCT-based backup, merge all the checkpoints of the virtual machines for the RCT backups to succeed.
- You can back up together the virtual machines that are configured on CSV and SMB.
- You cannot back up the virtual machines with a configuration version earlier than 6.2.
- You cannot back up the virtual machines that are present on the local storage node, and added to a cluster.
- You cannot back up the virtual machines that contain shared disks.

To configure RCT backups, use one of the following applications:

- **Client Backup Configuration** wizard of NMC
[Creating a client resource for a VSS-based backup by using the Client Backup Configuration wizard](#) on page 47 provides information.
- **Client Properties** dialog box of the NetWorker Administration program
[Manually creating a client resource for a VSS-based backup by using the Client Properties dialog box](#) on page 52 provides information.

Configuring backups

When configuring backups, the backup tasks differ depending on the items to back up.

The following table describes the backup tasks you must perform when using NetWorker server 18.2 and NMM 18.2 to back up Hyper-V parent and virtual machines. The *NetWorker Administration Guide* provides detailed information about how to perform the tasks in the table.

Table 17 Backup tasks for Hyper-V

Items to back up	Backup tasks to perform
<p>On the server</p> <p>The Hyper-V role can coexist with other Microsoft applications, such as:</p> <ul style="list-style-type: none"> • SQL Server • SharePoint Server • Exchange Server • Windows Server Cluster 	<ol style="list-style-type: none"> 1. Configure the NetWorker devices for backup storage. 2. Configure a backup group. 3. Configure one or more client resources for each client. 4. Configure a data protection policy for scheduled backups, including selecting a group, policy, policy workflow, and backup action. 5. Configure required NetWorker privileges. 6. Configure backup proxies.
<p>Hyper-V on the server</p> <p>Hyper-V virtual machines and Host Component file</p>	<ol style="list-style-type: none"> 1. Configure the backup storage resources. 2. Configure a backup group. 3. Configure one or more client resources for each client. 4. Configure a data protection policy for scheduled backups, including selecting a backup group, policy, policy workflow, and backup action.
<p>Hyper-V virtual machine applications</p> <p>Microsoft application data, such as:</p> <ul style="list-style-type: none"> • SQL Server • SharePoint Server • Exchange Server 	<p>Install NMM on the virtual machine operating system and configure application backups with NMM installed within the virtual machine operating system:</p> <ul style="list-style-type: none"> • Configure Windows application backups. • Configure Windows Server cluster backups.

Table 17 Backup tasks for Hyper-V (continued)

Items to back up	Backup tasks to perform
<ul style="list-style-type: none"> Windows Server Cluster 	Specific instructions for the Microsoft application are provided in the application-specific user guides.

Note

The "Using NMM 18.2 with NetWorker 8.2.3 or later" section provides the considerations to follow when using NetWorker 8.2.3 and later and NMM 18.2.

The following table lists the application information variables that NMM supports when you manually configure client resources:

Table 18 Supported application information variables

Client or workflow	Hyper-V Servers 2012 and 2012 R2 VSS	Hyper-V Server 2016 VSS	Hyper-V Server 2016 RCT
Standalone		NSR_ RCT_BACKUP=No	NSR_ RCT_BACKUP=Yes
Clustered	NSR_FEDERATED_B ACKUP=Yes NSR_FEDERATED_T YPE=CSV/SMB NSR_MOVE_CSV_O WNSHIP=Yes/No	NSR_ RCT_BACKUP=No NSR_FEDERATED_B ACKUP=Yes NSR_FEDERATED_T YPE=CSV/SMB NSR_MOVE_CSV_O WNSHIP=Yes/No	NSR_ RCT_BACKUP=Yes NSR_FEDERATED_B ACKUP=Yes

Creating a client resource for an RCT-based backup by using the Client Backup Configuration wizard

The Client Backup Configuration wizard for Hyper-V simplifies configuration of scheduled backups for Hyper-V servers. The NMM client must be installed on all Hyper-V cluster nodes for the Client Backup Configuration wizard to function correctly. The wizard automatically configures Hyper-V save sets, backup commands, application information variables, and backup options. Use the wizard to configure client resources for stand-alone and federated environments.

Before you begin

Before you use the wizard, review the following requirements:

- Ensure that the NetWorker server host is listed in the `servers` file on the client computer.
- Ensure that the communication between the NMC server, NetWorker client, and NetWorker server uses `nsrauth` strong authentication.
- Ensure that the user who runs the wizard meets the following requirements:
 - Root (UNIX) or Administrator (Windows) privileges.
 - A member of a User Group on the NetWorker server that has Configure NetWorker privileges.

- Ensure that multiple wizard hosts are not trying to access the same client computer simultaneously.

Note

If you use NetWorker server 8.2.3 or later and NMM 18.2:

- For the Client Backup Configuration wizard to properly function, ensure that JRE is installed on the host, where NMC is used. The NMC for NetWorker 8.2.3 supports JRE 7, and the NMC for NMM 18.2 supports JRE 8 and 9.
 - The procedure used to create a client resource when you use NetWorker server 8.2.3 and later is different from the procedure used to create a client resource when using NetWorker server 18.2.
The "Scheduled Backup" chapter in the *NetWorker Module for Microsoft Release Administration Guide* provides more information about editing client resources that were created by using NMM 8.2.x, and information about the NMC bulk edit feature.
 - Before you start the NMM 18.2 Client Backup Configuration wizard to modify a client resource that was created by using NMM 8.2.x, ensure that the **Snapshot** option of the NetWorker group that this client resource belongs to is clear. If the **Snapshot** option is selected, you cannot select the NetWorker group in the wizard, and you are prompted to create or select another group.
 - Do not use NetWorker strong authentication (nsrauth) to communicate with other hosts, such as NMC server, NetWorker server, and NetWorker client. The NetWorker 18.2 Administration Guide provides more information about nsrauth.
-

Procedure

1. Use NMC to connect to the NetWorker server.
2. In the **NetWorker Administration** window, click **Protection**.
3. In the expanded left panel, right-click **Clients**, and then select **New Client Wizard**.

The **Client Backup Configuration** wizard appears.

4. On the **Specify Client Information** page:
 - a. In the **Client Name** field, type either the hostname or the Fully Qualified Domain Name (FQDN) of the client.
For federated backups, type the cluster server name
For non-federated backups, type the stand-alone Hyper-V server name.
-

Note

Do not type the IP address of the client or the server.

- b. Optionally, in the **Comment** box, type a description of the client.
If you are creating multiple client resources for the same NetWorker client host, specify this field to differentiate the purpose of each resource.
- c. In the **Tag** box, type one or more tags to identify this Client resource for the creation of dynamic client groups for data protection policies.
Place each entry on a separate line.

- d. In the **Type** box, select **Traditional NetWorker client**.
- e. Optionally, from the **Group** list, select a group for the Client resource.

The group to which the client belongs determines the workflow that is used to back up the client.

Note

You can also assign the client to one or more groups after you create the Client resource.

- f. Click **Next**.
5. On the **Specify the Backup Configuration Type** page, select **Hyper-V Server** and click **Next**.
6. On the **Select the NetWorker Client Properties** page:
 - a. Select the priority level in the **Priority** field.
 - b. Select the level of parallelism in the **Parallelism** field.
 - c. Leave the **Remote Access** field empty.
 - d. Select the device type from the **Data Domain Interface** list.
 - e. Select the **Parallel Save Streams** option to enable multiple save streams for each save set during backup.
 - f. Click **Next**.

Note

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

7. On the **Select the Hyper-V Backup Objects** page, to exclude a save set from the backup, perform the following steps:
 - a. Select the top level save set.
 - b. Select the **Exclude Component List** option that appears at the bottom of the page.
 - c. Clear the save set that you want to exclude from the backup.
 - d. Click **Next**.

According to your setup — standalone or federated, the **Specify Backup Options** page accordingly appears as shown in one of the following figures:

Figure 10 Specify Backup Options page for a standalone setup

Figure 11 Specify Backup Options page for a clustered setup

8. On the **Specify Backup Options** page:

a. To perform federated backups, specify the following fields:

- In the **Remote username** field, type the clustered server username that has the administrator privileges.
- In the **Password** field, specify the corresponding password for the username that you have specified in the **Remote username** field.

For federated setups, the wizard creates empty (placeholder) client resources for all the nodes that you do not use to perform backups. The wizard also creates a client resource with the cluster name, and specifies the *NSR_FEDERATED_BACKUP=yes* variable.

For standalone setups, the wizard creates a client resource with the physical name by using the selected save sets.

b. Under **Backup Type**, select **RCT-based**.

This type of backup applies to Hyper-V Server 2016 and later.

c. Select the backup options:

- **Perform partial backup:** This option is selected by default. It enables backing up other virtual machines when the snapshot operation fails for some of the virtual machines that you have selected for the backup.
- **Disable incremental backup:** This option disables incremental backups so that only full backups are performed.

d. (Optional) To distribute the backup workload across physical and virtual servers, select proxy servers for the Preferred Server Order List (PSOL):

- To add a server from the **Available Servers** list to the **Proxy Servers** list, select the server, and then click the right arrow.
Repeat this step for each server that you want to add to the **Proxy Servers** list.

To remove a server from the **Proxy Servers** list, select the server, and then click the left arrow.
- To add all the servers from the **Available Servers** list to the **Proxy Servers** list at a time, click the right double arrow.
To remove all the servers from the **Proxy Servers** list, click the left double arrow.
- To change the order or position of the servers in the **Proxy Servers** list, click either the up arrow or the down arrow.

If a server is unavailable, NMM performs the backup from the node, to which the cluster server name resolves.

Note

RCT backups support only cluster nodes, but not virtual machines as proxy servers.

If you specify proxy servers (cluster nodes), and perform an RCT backup, the backup operation performs rollover on all the specified servers.

If you do not specify proxy servers, and perform an RCT backup, the backup operation performs rollover of the virtual machines from the respective nodes, on which the virtual machines are present.

Ensure that NetWorker client and NMM are installed on each cluster node.

e. Click **Next**.

9. On the **Backup Configuration Summary** page, perform one of the following steps:

- Click **Back** to go to the previous pages to modify the configuration.
- Click **Create** to configure the client resources.
The **Client Configuration Results** page appears. This page provides details about the client resources that have been created.

10. To modify the previous configuration, in NMC, select the client, right-click, and select **Modify Client Wizard**.

11. To verify the details for the client, in the **NetWorker Administration** window, right-click the client, and then select **Modify Client Properties**.
12. (Optional) To modify remote user access:
 - a. In the **NetWorker Administration** window, right-click the client, and then select **Modify Client Properties**.
 - b. On the **Globals (2 of 2)** tab, type the hostname of the proxy client in the **Remote Access** field.

If the NMM client is part of a cluster, type the names of the physical nodes of the cluster in the **Remote Access** field.

Manually creating a client resource for an RCT-based backup by using the Client Properties dialog box

You can manually create a Hyper-V client resource by using the **Client Properties** dialog box.

Procedure

1. In the **NetWorker Administration** window, click **Protection**.
 2. In the expanded left panel, select **Clients**.
 3. Right-click **Clients** and select **New**.
- The **Client Properties** dialog box appears.
4. On the **General** tab:
 - a. In the **Name** field, type either the hostname or the FQDN of the client.

The client must be a fully qualified host to be a NetWorker client.

Note

In this field, do not type the IP address of the client.

- b. In the **Comment** field, type a description. If you are creating multiple client resources for the same NetWorker client host computer, specify this field to differentiate the purpose of each resource.
- c. In the **Save Set** field, specify the components that you want to back up, each in a line. The following table lists the type of backup data and the corresponding save set syntaxes that you must specify in this field.

Table 19 Hyper-V save set syntax

Type of backup data	Save set syntax
Hyper-V virtual machine There are usually multiple virtual machines on the host operating system.	APPLICATIONS:\Microsoft Hyper-V \virtual_machine_name Child pertains or virtual machines can be in a proxy backup group.

To perform a CSV-level federated backup, you must set the application information variables for the cluster virtual server and the client resource.

- d. Select the appropriate option in the **Protection group** list field.

5. Click the **Apps & Modules** tab and do the following:

a. In the **Remote user** and **Password** fields respectively, type the domain administrator username and password.

b. In the **Backup command** field, type the following backup command:

```
nsrnmmsv.exe
```

c. In the **Application Information** field, specify the application information variables that you require to perform the backup. The following table lists the variables that you can specify in this field.

Table 20 Hyper-V application information variables

Variable name	Description	Values
NSR_RCT_BACKUP	Performs RCT-based backups.	<ul style="list-style-type: none"> Yes No
NSR_FEDERATED_BACKUP	Marks the backup for CSV and SMB recovery.	Yes
NSR_FEDERATED_TYPE	Specifies whether the backup is federated CSV or SMB.	<ul style="list-style-type: none"> csv smb <p>If this variable is not specified, NMM applies the CSV value by default.</p>
NSR_FEDERATED_PSOL	<p>Optional. Distributes the backup workload across physical and virtual servers in the PSOL. If a server is not available or down, then NMM performs the backup from the node to which the cluster server name resolves, including virtual proxies.</p> <p>RCT backups support only cluster nodes, but not virtual machines as proxy servers.</p> <p>If you specify this application information variable with a list of servers (cluster nodes) as values, and perform an RCT backup, the backup operation performs rollover on all the specified servers.</p> <p>If you do not specify this application information variable, and perform an RCT backup, the backup operation performs rollover of the virtual machines from the respective nodes, on which the virtual machines are present.</p> <hr/> <p>Note</p> <p>Ensure that NetWorker client and NMM are installed on each cluster node.</p>	<p>Type a comma-separated list of the server names. For example:</p> <pre>NSR_FEDERATED_PSOL=FQDN_of_server1, FQDN_of_server2, FQDN_of_server3,...</pre> <hr/> <p>Note</p> <p>A Hyper-V CSV backup fails if the same host name is added twice in PSOL list.</p>
NSR_EXCLUDE_SMB	Optional. Excludes virtual machines that have data that is stored on SMB file servers. By	Yes

Table 20 Hyper-V application information variables (continued)

Variable name	Description	Values
	default, SMB virtual machines are in the writer level backup.	
NSR_PARTIAL_BACKUP	Optional. Supports partial backup failure. If snapshot creation fails for some virtual machines, the backup continues for other virtual machines. The Hyper-V VSS Writer reports the virtual machines that failed to backup. The default value is Yes.	<ul style="list-style-type: none"> • Yes • No
NSR_EXCLUDE_COMPONENTS	Optional. Excludes a virtual machine from the backup. Specify the writer level save set and the components to exclude from the backup. NMM logs the excluded components in the <code>nsrnmmsv.raw</code> log file for references.	Type a comma-separated list of the server names. For example: <code>NSR_EXCLUDE_COMPONENTS=VM1 , VM2 , VM3</code>
NSR_EXCLUDE_SUBCOMPONENTS	<p>Excludes VHDs and VHDXs from a Hyper-V backup in standalone, federated, and SMB configurations.</p> <hr/> <p>Note</p> <p>Ensure that the VHDs and VHDXs do not have special characters comma (,), colon (:), equal to (=), and vertical bar () in their names because the special characters are used as separators.</p> <hr/>	<p>Use comma (,), colon (:), and vertical bar () as separators to exclude VHDs and VHDXs from the backup. For example:</p> <p><code>NSR_EXCLUDE_SUBCOMPONENTS = VM1 : VHD1 . vhd x , VHD2 . vhd x VM2 : VHD3 . vhd x</code></p> <p><code>NSR_EXCLUDE_SUBCOMPONENTS = VM1 : VHD1 . vhd x , VHD2 . vhd x</code></p> <p><code>NSR_EXCLUDE_SUBCOMPONENTS = VM2 : VHD3 . vhd x</code></p> <p>where,</p> <ul style="list-style-type: none"> • The colon (:) is used as the separator between the virtual machine name and VHD name. • The comma (,) is used as the separator between two VHD names. • The vertical bar () is used as the separator between two VMs.
NSR_DISABLE_INCREMENTAL	Optional. By default, when Hyper-V backups are performed by using NMM, the first instance of the backup is always full and subsequent backups are incremental. The <code>NSR_DISABLE_INCREMENTAL=yes</code> variable disables incremental backups and enables only full backups. When incremental backups are disabled using this parameter, the recovery snapshots, that is .AVHDX files, are merged prior to the full backup.	<p>Use the value is Yes to disable change block tracking (CBT) and perform only full backups.</p> <p>Use the value is No to enable change block tracking (CBT) and perform both full and incremental backups.</p> <p>You have the option to enable CBT or disable CBT at any time depending on your business needs.</p>

6. Click the **Globals (1 of 2)** tab and, in the **Aliases** field, type the NETBIOS name for the client.

NMM client uses the host computer NETBIOS or short name when connecting to the NetWorker server to browse backups. If the NETBIOS name is not found, NMM cannot display backups.

7. To set up remote access, click the **Globals (2 of 2)** tab and do one of the following:
 - If you are setting up a proxy client for the NMM client, type the hostname of the proxy client in the **Remote Access** field.
 - If the NMM client is part of a cluster, type the names of the physical nodes of the cluster in the **Remote Access** field.
8. Click **OK**.

Configuring NMM 8.2.x client resources after upgrading NMM from 8.2.x to 9.1 or later

After you upgrade NMM from 8.2.x to 9.1 or later, to use the same 8.2.x client resources to perform backups, you must perform the following changes to all the 8.2.x client resources:

Procedure

1. Either ensure that the **Snapshot** option of the NetWorker group that a client resource belongs to is clear or create a new group and ensure that the **Snapshot** option is clear.
2. In the **NetWorker Administration** window, click **Protection**.
3. Either right-click **Clients** in the navigation tree or right-click the required client in the **Clients** table.
4. Select **Modify Client Properties**.
5. In the **Backup Command** field, delete the `nsrsnap_vss_save` command, and then type the `nsrnmmSV` command.

Perform this step for all client resources.

6. To enable Hyper-V virtual machine full backups, in the **Application Information** field, type the following variables each in a line:
 - `NSR_VSS_FULL_BACKUP=yes`
 - `NSR_RCT_BACKUP=no`

To perform Hyper-V virtual machine copy full backups, in the **Application Information** field, perform one of the following steps:

 - Do not specify the `NSR_VSS_FULL_BACKUP` variable.
 - Type `NSR_VSS_FULL_BACKUP=no`.
7. Make other changes according to your requirement.
8. Click **OK**.

NOTICE

The "Scheduled Backup" chapter in the *NetWorker Module for Microsoft Release Administration Guide* provides more information about editing client resources that were created by using NMM 8.2.x, and information about the NMC bulk edit feature.

Improving performance of RCT-based backups by using multiple cluster nodes as proxies

In larger Hyper-V environments regardless of CSV or SMB, you can improve backup performance by using multiple cluster nodes as proxies.

To improve backup performance by using multiple cluster nodes as proxies, ensure that you meet the following requirements:

- Install NMM on all cluster nodes.

Note

RCT-based backups do not support virtual machines as proxies.

- Create client resources for all the cluster nodes and cluster aliases.
- Add secondary roles, which must be physical or virtual cluster nodes.
- To distribute backup work load and improve performance, select proxy hosts for Preferred Server Order List (PSOL) by using one of the following methods:
 - Using the **Select proxy servers for Preferred Server Order List (PSOL)** field on the **Specify backup options** page in the **Client Backup Configuration** wizard
 - Specifying the *NSR_FEDERATED_PSOL* application information variable in the **Client Properties** dialog box of NMC

If you specify proxy hosts for PSOL, and perform an RCT backup, the backup operation performs rollover on all the specified servers.

If you do not specify proxy hosts for PSOL, and perform an RCT backup, the backup operation performs rollover of the virtual machines from the respective nodes, on which the virtual machines are present.

- Schedule backups against the cluster alias so that the primary role runs on the cluster node that owns the cluster alias.

In a multi-proxy setup, to perform concurrent backups on all proxy nodes, you can select multiple cluster nodes as proxy nodes. Concurrent backups increase the backup performance. According to your requirement, you can add or remove proxy nodes.

The procedure to recover the virtual machines that are backed up as part of a multi-proxy setup is the same as the procedure to recover traditional backups.

Manually configuring highly available backups (cluster-aware backups)

The following procedure describes how to manually configure the backup of CSV virtual machines as part of a highly available (cluster-aware) backup or a physical proxy node backup. Cluster-aware backups are highly available because you install NMM on each node in the cluster. If a node is not available, NMM starts the backup from the node that resolves to the cluster server name at runtime.

Procedure

1. Install NMM on each node in the cluster.
2. Create an empty (placeholder) NetWorker client resource for each node in the cluster.

3. Create a client resource for the cluster server name and specify the save sets to back up. Add this client to a backup group.
4. At runtime, the cluster server name resolves to one of the nodes in the cluster. This node becomes the master backup node.

Viewing backup status and summary

Besides log files, you can use the NetWorker Administration window to view the status of backups.

The **Policies** section on the **Monitoring** tab in the NetWorker Administration window displays the status of the backups.

In the case of a backup failure, you can also view the summary by expanding the failed backup and double-clicking **pseudo_saveset**. The **Show Messages** dialog box displays the backup summary, which is case-insensitive.

CHAPTER 4

Recoveries

This chapter includes the following sections:

- [Overview of recoveries](#).....74
- [Recovery scenarios and GUIs to use for various types of recoveries](#)..... 74
- [Recovering Hyper-V virtual machines](#)..... 78
- [Using NMM 9.1 or later to recover the backups that were performed by using NMM 8.2.x](#).....94

Overview of recoveries

Depending on what you specified in the backup save set, you can recover the following from a Hyper-V virtual machine backup by using the NetWorker User for Microsoft GUI:

- All Hyper-V components
- The Host Component file
Applies only to VSS-based backups
- Individual or multiple virtual machines
- Granular level recoveries for individual files and folders

NOTICE

To recover backups that were created using an NMM release earlier than 9.0, click **Start > EMC NetWorker > NetWorker Tools > Restore previous NMM release backups** to start the NetWorker Module for Microsoft GUI. Browse the backups and perform the recovery from the GUI that appears.

Recovery scenarios and GUIs to use for various types of recoveries

The tables in this section list various recovery scenarios that NMM supports, and the corresponding GUIs that you must use to perform various types of recoveries.

Table 21 Recovery scenarios

Recovery scenario	Host, on which GUI and recovery binary starts	Host, to which virtual machine recovers	Supported on Hyper-V 2012 R2	Supported on Hyper-V 2016
Standalone - Recover virtual machine to the same host	<ul style="list-style-type: none"> • Start GUI on the same host • Start recovery binary on the same host 	Virtual machine recovers to the same host	Yes	Yes
Standalone - Recover virtual machine to an alternate location	<ul style="list-style-type: none"> • Start GUI on the same host • Start recovery binary on the selected host 	Virtual machine recovers to the selected host	Yes	Yes

Table 21 Recovery scenarios (continued)

Recovery scenario	Host, on which GUI and recovery binary starts	Host, to which virtual machine recovers	Supported on Hyper-V 2012 R2	Supported on Hyper-V 2016
Standalone - Recover virtual machine to Hyper-V core server	<ul style="list-style-type: none"> Start GUI on the Windows host (Desktop Experience) Start recovery binary on the Hyper-V core server 	Virtual machine recovers to Hyper-V core server	Yes	Yes
Standalone Hyper-V core server - Recover virtual machine to an alternate host	<ul style="list-style-type: none"> Start GUI on the Windows host (Desktop Experience) Start recovery binary on the selected host 	Virtual machine recovers to the selected host	Yes	Yes
Cluster - Recover virtual machine to the primary node, and run the recovery process on the host, on which virtual machine was present during backup	<ul style="list-style-type: none"> Start GUI on the primary node Virtual machine to recover is present on a node in the cluster Start recovery binary on the primary node 	Virtual machine recovers to the primary node	Yes	No - Respective node option is disabled
Cluster - Recover virtual machine to a secondary node, and run the recovery process on the host, on which virtual machine was present during backup	<ul style="list-style-type: none"> Start GUI on the secondary node Virtual machine to recover is present on a node in the cluster 	Virtual machine recovers to the secondary node	Yes	No - Respective node option is disabled

Table 21 Recovery scenarios (continued)

Recovery scenario	Host, on which GUI and recovery binary starts	Host, to which virtual machine recovers	Supported on Hyper-V 2012 R2	Supported on Hyper-V 2016
	<ul style="list-style-type: none"> Start recovery binary on the secondary node 			
Cluster - Recover virtual machine to the primary node, and run the recovery process on an alternate host (a secondary host)	<ul style="list-style-type: none"> Start GUI on the primary node Virtual machine to recover is present on a node in the cluster Start recovery binary on the selected secondary node 	Virtual machine recovers to the primary node	Yes	Yes
Cluster - Recover virtual machine to a secondary node, and run the recovery process on an alternate host (the primary host)	<ul style="list-style-type: none"> Start GUI on the secondary node Virtual machine to recover is present on a node in the cluster Start recovery binary on the selected primary node 	Virtual machine recovers to the secondary node	Yes	Yes
Cluster - Recover virtual machine to the primary node, and run the recovery process on the respective node	<ul style="list-style-type: none"> Start GUI on the secondary node Virtual machine to recover is present on 	Virtual machine recovers to the primary node	Yes	No

Table 21 Recovery scenarios (continued)

Recovery scenario	Host, on which GUI and recovery binary starts	Host, to which virtual machine recovers	Supported on Hyper-V 2012 R2	Supported on Hyper-V 2016
	<ul style="list-style-type: none"> the primary node in the cluster Start recovery binary on the primary node 			
Cluster - Recover multiple virtual machines, and run the recovery process on the respective nodes	<ul style="list-style-type: none"> Start GUI on the primary node Select two virtual machines - one on the primary node and the other on a secondary node in the cluster Start recovery binary first on the primary node Start recovery binary next on the secondary node 	<ul style="list-style-type: none"> Virtual machine 1 recovers to the primary node Virtual machine 2 recovers to the secondary node 	Yes	Yes
Cluster - Recover virtual machine to a different node of a different cluster (redirected recovery)	<ul style="list-style-type: none"> Start GUI on any node in the cluster Virtual machine to recover is present on its respective node Start recovery binary on the 	Virtual machines recover to the selected node	Yes	Yes

Table 21 Recovery scenarios (continued)

Recovery scenario	Host, on which GUI and recovery binary starts	Host, to which virtual machine recovers	Supported on Hyper-V 2012 R2	Supported on Hyper-V 2016
	selected node			

Table 22 GUIs used for various types of recoveries

Type of recovery	GUIs
Recovery of virtual machines: <ul style="list-style-type: none"> Multiple virtual machines are recovered to the original (source) location Only one virtual machine is recovered at a time to an alternate location 	NetWorker User for Microsoft GUI Recovering Hyper-V standalone server virtual machines on page 79 and Recovering Hyper-V clustered server virtual machines on page 84 provide information.
Recovery of files and folders in a virtual machine (Granular level recoveries)	NetWorker User for Microsoft GUI Recovering Hyper-V Server virtual machines at granular level on page 89 provides information.
Recovery of files and folders in a virtual machine (File level recoveries)	NetWorker User for Microsoft GUI File Level Recoveries on page 95 provides information.
Recovery of all System Center Virtual Machine Manager (SCVMM) managed virtual machines that have NMM conventional backups	Data Protection Add-in for SCVMM GUI SCVMM Recoveries on page 120 provides information.

Recovering Hyper-V virtual machines

The following recovery options are available when you perform regular backups of virtual machines:

Table 23 Recovery options for virtual machines

Recovery option	Description
Recover a virtual machine to the source (original) location, that is the location on the same Hyper-V Server, on which the backup was performed.	The original files are overwritten. Recover to the original location when: <ul style="list-style-type: none"> You must roll back the virtual machine because a patching or virus issue occurred. You must perform disaster recovery of the virtual machine after a disk failure. The virtual machine was accidentally deleted.
Recover a virtual machine to an alternate (different) location on same Hyper-V Server, on which the backup was performed.	The original files are not overwritten. You can recover a virtual machine to the original Hyper-V Server, but move the virtual machine files to different file system locations. Perform this type of recovery if the virtual machine files were moved after the selected backup time, and you want to preserve the new locations. If the original virtual machine is present, the recovery operation overwrites the virtual machine.

Table 23 Recovery options for virtual machines (continued)

Recovery option	Description
Recover a virtual machine to an alternate Hyper-V Server.	<ul style="list-style-type: none"> Select an alternate Hyper-V server for the virtual machine recovery. Select a location on the destination Hyper-V Server. In a clustered environment, select the CSV, where the files will be placed during a recovery. <p>NMM supports redirected recoveries to a host that runs the same or later operating system version. For example: NMM supports redirected recovery from a Windows Server 2012 R2 source host to a Windows Server 2012 R2 destination host, but NMM does not support redirected recovery from a Windows Server 2012 R2 source host to a Windows Server 2012 destination host.</p>
Recover individual files and folders (Granular Level Recovery).	Hyper-V Granular Level Recovery (GLR) provides the ability to recover specific files from a virtual machine image backup without recovering the full virtual machine, reducing the recovery time. The NMM GLR feature mounts the virtual machine that contains the items to recover.

NOTICE

Before you perform a recovery, start the Microsoft iSCSI Initiator and Target services. Otherwise, the recovery fails.

Recovering Hyper-V standalone server virtual machines

NMM enables you to recover one virtual machine to either the source location or an alternate location, and multiple virtual machines to only their source locations. The alternate location includes a different location on the source Hyper-V Server and a different Hyper-V Server.

Before you begin

- To recover virtual machines to the source Hyper-V server, ensure that the original drive letters or mount points for the virtual machines exist on the source server. The directory paths are automatically created. Recovering virtual machines to the source server overwrites the source virtual machines.
- Because Hyper-V recognizes virtual machines by using an internal GUID, ensure that you do not either move or rename the virtual machines during the recovery operation.
- The Host Component file contains the authorization configuration for Hyper-V. If either the file is corrupted or you want to roll back the authorization settings, ensure that you recover the Host Component to the source Hyper-V Server. The NMM System Component backups also include the Host Component.

Note

This requirement applies to only Hyper-V 2012 R2 standalone VSS-based backups.

Procedure

- Open the NetWorker User for Microsoft GUI.
- Click the icon that is beside the **NetWorker server** field, and specify the NetWorker server.
- From the **Client** list, select the Hyper-V standalone server that contains the virtual machines that you want to recover.

If the Hyper-V server does not appear in the **Client** list, add it to the list:

- a. On the menu bar, click **Options > Configure Options**.
 - b. In the **Configuration Options** dialog box, click the icon beside the **Client name** field.
 - c. In the **Select Viewable Clients** dialog box:
 - a. From the **Available clients on <NetWorker_server_name>** list, select the Hyper-V Server, and then click **Add**.
The selected Hyper-V Server appears in the **Clients to list on menu bar** list.
 - b. Click **OK**.
 - d. In the **Configuration Options** dialog box, click **OK**.
4. From the left panel, click **Recover > Hyper-V Recover Session > Image Recovery**.
 5. In the middle panel, on the **Browse** tab:
 - a. Expand **Microsoft Hyper-V**.
 - b. To select all the virtual machines to recover, select at root-level (**Microsoft Hyper-V**). Otherwise, select individual virtual machines to recover.
 6. To view required volumes of a selected virtual machine either at the root-level or the individual-level to recover, right-click the virtual machine, and then select **Required volumes**.
In the **Required NetWorker Volumes** dialog box, review the list of volumes, and then click **OK**.
 7. To select a particular version or backup time of a virtual machine:
 - a. Right-click the virtual machine, and then select **Versions**.
 - b. In the **NetWorker Versions** dialog box:
 - a. Select the backup time.
 - b. Select **Use selected item backup time as new browse time**.
 - c. Click **OK**.
 8. To search for either a virtual machine or a VHD or VHDx that is attached to a virtual machine, in the middle panel:
 - a. Perform one of the following steps:
 - To search for a virtual machine, perform one of the following steps:
 - On the **Browse** tab, select **Microsoft Hyper-V**. Click the **Search** tab.
 - On the **Browse** tab, right-click **Microsoft Hyper-V**, and then select **Search for**.
The **Path** field displays the Microsoft Hyper-V path.
 - To search for a VHD or VHDx that is attached to a virtual machine, perform one of the following steps:
 - On the **Browse** tab, expand **Microsoft Hyper-V**, and then select the virtual machine. Click the **Search** tab.
 - On the **Browse** tab, expand **Microsoft Hyper-V**, right-click the virtual machine, and then select **Search for**.

The **Path** field displays the virtual machine path.

b. (Optional) In the **Name** field, type the name of the search item. You can refine the search by using any of the following types of search:

- Literal match (case-insensitive): Type `abc` to return `abc`, `ABC`, or `AbC` but not `abcd` or `ABCD`.
- Literal match (case-sensitive): Type `"abc"` to return `abc`, but not `AbC` or `abcd`.
- Name contains (case-insensitive): Type `%abc%` to return `abc`, `abcd`, `ABCD`, or `xyzABCde`.
- Name starts with (case-insensitive): Type `abc%` to return `abcd` or `ABCde`, but not `xyzABCde`.
- Name ends with (case-insensitive): Type `%abc` to return `xyzAbc`, but not `ABCde`.
- Single-character match search by using the `?` wildcard:
 - Type `?` to return single character entries and drive volumes, such as `C` or `D`.
 - Type `<writer_name>?Writer` to return the `<writer_name> Writer`, for example, `abc Writer`.
- Multiple-character match search by using the `*` wildcard:
 - Type `*.txt` to return all entries with a `.txt` extension.
 - Type `*` to return all items within the selected container.
 - Type `*writer*` to return all writers.
- Search by using the `*` and `?` wildcards: Type `*???*writer*` to return the `abc Writer`.

c. Click **Search**.

The **Result** panel displays the search results.

9. From the **Hyper-V Recover Session** toolbar, click **Recover**.

The **Hyper-V Virtual Machine Recovery** wizard appears with the **Hyper-V Virtual Machine Recovery Options** page.

Figure 12 Standalone - Hyper-V virtual machine recovery options page

Hyper-V Virtual Machine Recovery

Hyper-V Virtual Machine Recovery Options
Specify Hyper-V virtual machine recovery options.

Recovery Options

☒ Recover the virtual machines to the respective source (original) locations
(This option overwrites the source virtual machines.)

☐ Recover the virtual machine to an alternate location
(This option is enabled only when you have selected one virtual machine to recover.)

Diagnostic Level

Select the diagnostic output level: Normal

Host to Run the Recovery Process

☐ Run the recovery process on the respective hosts

☒ Run the recovery process on an alternate host

WIN-8RQ5K2Q75C0.WORKGROUP

< Back Next > Cancel Help

10. Perform one of the following types of recovery:

- To recover the selected virtual machines to their source locations:
 - a. On the **Hyper-V Virtual Machine Recovery Options** page:
 - a. Select **Recover the virtual machines to the respective source (original) locations**.

Note

This option overwrites the source virtual machines.

- b. From the **Select the diagnostic output level** list, select a debug level, **Normal** (or 1) through **9** according to the amount of recovery debug information that you want the GUI to log. The debug level **9** logs more information.
- c. Click **Next**.
- b. On the **Hyper-V Virtual Machine Recovery Options Summary** page:
 - a. Review the settings that you have configured.
To modify the settings, click **Back**.
 - b. To start the recovery operation, click **Start**.

Note

When the recovery is in progress, you cannot perform other tasks in the **Hyper-V Virtual Machine Recovery** wizard.

- To recover the selected virtual machine to an alternate location:
 - a. On the **Hyper-V Virtual Machine Recovery Options** page:
 - a. Select **Recover the virtual machine to an alternate location**.

Note

This option is enabled when you have selected only one virtual machine to recover.

- b. From the **Select the diagnostic output level** list, select a debug level, **Normal** (or 1) through **9** according to the amount of recovery debug information that you want the GUI to log. The debug level **9** logs more information.
- c. Click **Next**.
The **Destination Hyper-V Server and Path** page appears.

Figure 13 Standalone - Destination Hyper-V Server and path page

Hyper-V Virtual Machine Recovery

Destination Hyper-V Server and Path
Specify the destination Hyper-V Server and path, to which you want to recover the virtual machine.

Original Hyper-V Server: WIN-8RQ5K2Q75C0

Destination Hyper-V Server

☒ Recover the virtual machine to an alternate path on the source (original) Hyper-V Server

☐ Recover the virtual machine to an alternate Hyper-V Server

Destination Path

Specify the destination path for the virtual machine configuration files:

Select the destination path for the virtual machine:

VM	VHD	Destination
VM4	E:\HyperV\VM4\Virtual ...	

- b. On the **Destination Hyper-V Server and Path** page:
 - a. Select one of the following alternate location options:
 - **Recover the virtual machine to an alternate path on the source (original) Hyper-V Server:** This option recovers the selected virtual machine to an alternate path on the source Hyper-V Server.
 - **Recover the virtual machine to an alternate Hyper-V Server:** This option enables you to select a Hyper-V Server to recover the selected virtual machine.
From the list, select the destination Hyper-V Server.
 - b. In the **Specify the destination path for the virtual machine configuration files** field, click **Browse**, and then select the destination path to recover the virtual machine configuration files.
 - c. From the **Select the destination path for the virtual machine** table or list, select the destination path to recover the selected virtual machine.
To change the destination path, click **Change Destination**, and then select the path.

- d. Click **Next**.
- c. On the **Hyper-V Virtual Machine Recovery Options Summary** page:
 - a. Review the settings that you have configured.
To modify the settings, click **Back**.
 - b. Click **Start** to start the recovery operation.

Note

When the recovery is in progress, you cannot perform other tasks in the **Hyper-V Virtual Machine Recovery** wizard.

11. To view the status of the recovery, in the left panel of the NetWorker User for Microsoft GUI, click **Monitor**.

Recovering Hyper-V clustered server virtual machines

NMM enables you to recover Hyper-V virtual machines at the cluster and individual levels. The host that runs the Hyper-V service can be outside a cluster.

For a deleted CSV virtual machine, that is, if a virtual machine does not exist in the cluster:

- NetWorker User for Microsoft GUI is started within the cluster, and the recovery operation starts on the node, on which the NetWorker User for Microsoft GUI is started.
- NetWorker User for Microsoft GUI is started from outside the cluster, and the deleted virtual machine is recovered to the cluster owner node.

Before you perform recovery, consider the following limitations:

- NMM does not support redirected recovery of virtual machines over SMB 3.0.
- NMM does not enable you to recover a virtual machine to a node that is not active on the cluster that contains the virtual machine.
If you perform a relocated recovery of a virtual machine to a node on a cluster, and the virtual machine is active on another cluster node, the recovery fails.

Before you perform recovery, ensure that you meet the following requirements:

- Ensure that the NMM client is installed on all the clustered nodes.
- To recover virtual machines to the source Hyper-V server, ensure that the original drive letters or mount points for the virtual machines exist on the source server. The directory paths are automatically created. Recovering virtual machines to the source server overwrites the source virtual machines.
- Because Hyper-V recognizes virtual machines by using an internal GUID, ensure that you do not either move or rename the virtual machines during the recovery operation.

Procedure

1. Open the NetWorker User for Microsoft GUI.
2. Click the icon that is beside the **NetWorker server** field, and specify the NetWorker server.
3. From the **Client** list, select the Hyper-V standalone server that contains the virtual machines that you want to recover.

If the Hyper-V server does not appear in the **Client** list, add it to the list:

- a. On the menu bar, click **Options > Configure Options**.

- b. In the **Configuration Options** dialog box, click the icon beside the **Client name** field.
 - c. In the **Select Viewable Clients** dialog box:
 - a. From the **Available clients on <NetWorker_server_name>** list, select the Hyper-V clustered server, and then click **Add**.
The selected Hyper-V clustered server appears in the **Clients to list on menu bar** list.
 - b. Click **OK**.
 - d. In the **Configuration Options** dialog box, click **OK**.
4. From the left panel, click **Recover > Hyper-V Recover Session > Image Recovery**.
5. In the middle panel, on the **Browse** tab:
 - a. Expand **Microsoft Hyper-V**.
 - b. To select all the virtual machines to recover, select at root-level (**Microsoft Hyper-V**). Otherwise, select individual virtual machines to recover.
6. To view required volumes of a selected virtual machine either at the root-level or the individual-level to recover, right-click the virtual machine, and then select **Required volumes**.
In the **Required NetWorker Volumes** dialog box, review the list of volumes, and then click **OK**.
7. To select a particular version or backup time of a virtual machine:
 - a. Right-click the virtual machine, and then select **Versions**.
 - b. In the **NetWorker Versions** dialog box:
 - a. Select the backup time.
 - b. Select **Use selected item backup time as new browse time**.
 - c. Click **OK**.
8. To search for either a virtual machine or a VHD or VHDx that is attached to a virtual machine, in the middle panel:
 - a. Perform one of the following steps:
 - To search for a virtual machine, perform one of the following steps:
 - On the **Browse** tab, select **Microsoft Hyper-V**. Click the **Search** tab.
 - On the **Browse** tab, right-click **Microsoft Hyper-V**, and then select **Search for**.
The **Path** field displays the Microsoft Hyper-V path.
 - To search for a VHD or VHDx that is attached to a virtual machine, perform one of the following steps:
 - On the **Browse** tab, expand **Microsoft Hyper-V**, and then select the virtual machine. Click the **Search** tab.
 - On the **Browse** tab, expand **Microsoft Hyper-V**, right-click the virtual machine, and then select **Search for**.
The **Path** field displays the virtual machine path.

- b. (Optional) In the **Name** field, type the name of the search item. You can refine the search by using any of the following types of search:
- Literal match (case-insensitive): Type `abc` to return `abc`, `ABC`, or `AbC` but not `abcd` or `ABCD`.
 - Literal match (case-sensitive): Type `"abc"` to return `abc`, but not `AbC` or `abcd`.
 - Name contains (case-insensitive): Type `%abc%` to return `abc`, `abcd`, `ABCD`, or `xyzABCde`.
 - Name starts with (case-insensitive): Type `abc%` to return `abcd` or `ABCde`, but not `xyzABCde`.
 - Name ends with (case-insensitive): Type `%abc` to return `xyzAbc`, but not `ABCde`.
 - Single-character match search by using the `?` wildcard:
 - Type `?` to return single character entries and drive volumes, such as `C` or `D`.
 - Type `<writer_name>?Writer` to return the `<writer_name> Writer`, for example, `abc Writer`.
 - Multiple-character match search by using the `*` wildcard:
 - Type `*.txt` to return all entries with a `.txt` extension.
 - Type `*` to return all items within the selected container.
 - Type `*writer*` to return all writers.
 - Search by using the `*` and `?` wildcards: Type `*???*writer*` to return the `abc Writer`.

c. Click **Search**.

The **Result** panel displays the search results.

9. From the **Hyper-V Recover Session** toolbar, click **Recover**.

The **Hyper-V Virtual Machine Recovery** wizard appears with the **Hyper-V Virtual Machine Recovery Options** page.

Figure 14 Clustered - Hyper-V virtual machine recovery options page

Hyper-V Virtual Machine Recovery

Hyper-V Virtual Machine Recovery Options
Specify Hyper-V virtual machine recovery options.

Recovery Options

- ☒ Recover the virtual machines to the respective source (original) locations
(This option overwrites the source virtual machines.)
- ☐ Recover the virtual machine to an alternate location
(This option is enabled only when you have selected one virtual machine to recover.)

Diagnostic Level

Select the diagnostic output level: Normal

Host to Run the Recovery Process

- ☐ Run the recovery process on the respective hosts
- ☒ Run the recovery process on an alternate host

CSVNODE3.MSAPPS.COM

< Back Next > Cancel Help

10. Perform one of the following types of recovery:

- To recover the selected virtual machines to their source locations:
 - a. On the **Hyper-V Virtual Machine Recovery Options** page:
 - a. Select **Recover the virtual machines to the respective source (original) locations**.

Note

This option overwrites the source virtual machines.

- b. From the **Select the diagnostic output level** list, select a debug level, **Normal** (or 1) through **9** according to the amount of recovery debug information that you want the GUI to log. The debug level **9** logs more information.
- c. Select one of the following host options to run the recovery process:
 - **Run the recovery process on the respective hosts:** This option runs the recovery process on the hosts, on which the virtual machines are present. If a selected virtual machine is not present (deleted) in the cluster, the recovery process for the deleted virtual machine is run on the host, on which the NetWorker User for Microsoft GUI runs.

Note

This option is disabled for Hyper-V Server 2016.

- **Run the recovery process on an alternate host:** This option enables you to select a host in the cluster to run the recovery process. From the list, select the Hyper-V Server. By default, the local host is selected.

- d. Click **Next**.
- b. On the **Hyper-V Virtual Machine Recovery Options Summary** page:
 - a. Review the settings that you have configured.
To modify the settings, click **Back**.
 - b. To start the recovery operation, click **Start**.

Note

When the recovery is in progress, you cannot perform other tasks in the **Hyper-V Virtual Machine Recovery** wizard.

- To recover the selected virtual machine to an alternate location:
 - a. On the **Hyper-V Virtual Machine Recovery Options** page:
 - a. Select **Recover the virtual machine to an alternate location**.

Note

This option is enabled when you have selected only one virtual machine to recover.

- b. From the **Select the diagnostic output level** list, select a debug level, **Normal** (or 1) through **9** according to the amount of recovery debug information that you want the GUI to log. The debug level **9** logs more information.
- c. Click **Next**.
The **Destination Cluster Node or Hyper-V Server, and Path** page appears.

Figure 15 Clustered - Destination cluster node or Hyper-V Server, and path page

Hyper-V Virtual Machine Recovery

Destination Cluster Node or Hyper-V Server, and Path
Specify the destination cluster node or Hyper-V Server, and path, to which you want to recover the virtual machine.

Original Hyper-V Server: PradeepMittalCSV

Destination Cluster Node or Hyper-V Server

☒ Recover the virtual machine to the active cluster node
☐ Recover the virtual machine to an alternate cluster node
☐ Recover the virtual machine to an alternate Hyper-V Server

Destination Path

Specify the destination path for the Hyper-V client configuration files:

Select the destination path for the virtual machine:

VM	VHD	Destination
VM3_Node3	C:\ClusterStorage\Volume1\V...	

- b. On the **Destination Cluster Node or Hyper-V Server, and Path** page:

- a. Select one of the following alternate location options:
 - **Recover the virtual machine to the active cluster node:** This option recovers the selected virtual machine to the active cluster node.
This option is available only when the selected virtual machine to recover is active in the cluster.
 - **Recover the virtual machine to an alternate cluster node:** This option enables you to select a cluster node to recover the selected virtual machine.
From the list, select the destination cluster node.

This option is available only when the selected virtual machine to recover is deleted from the cluster.
 - **Recover the virtual machine to an alternate Hyper-V Server:** This option enables you to select a Hyper-V Server that is outside the cluster to recover the selected virtual machine.
From the list, select the destination Hyper-V Server.
- b. In the **Specify the destination path for the Hyper-V client configuration files** field, click **Browse**, and then select the destination path to recover the client configuration files.
- c. From the **Select the destination path for the virtual machine** table or list, select the destination path to recover the selected virtual machine.
To change the destination path, click **Change Destination**, and then select the path.
- d. Click **Next**.
- c. On the **Hyper-V Virtual Machine Recovery Options Summary** page:
 - a. Review the settings that you have configured.
To modify the settings, click **Back**.
 - b. Click **Start** to start the recovery operation.

Note

When the recovery is in progress, you cannot perform other tasks in the **Hyper-V Virtual Machine Recovery** wizard.

11. To view the status of the recovery, in the left panel of the NetWorker User for Microsoft GUI, click **Monitor**.

After you finish

Verify whether the recovery operation registered the virtual machine as a cluster resource. Otherwise, manually register the virtual machine as a cluster resource by using Failover Cluster Manager.

Recovering Hyper-V Server virtual machines at granular level

Granular level recovery (GLR) enables you to recover specific files from a single backup without recovering the full virtual machine. GLR reduces the recovery time. You can perform GLR by using the NMM client software.

You can perform GLR of NMM backups of Hyper-V virtual machine that has a Windows operating system. NMM does not support GLR of non-Windows virtual machines on Hyper-V, for example Linux virtual machines. The NMM GLR feature mounts the virtual machine that contains the items to recover.

NMM GLR supports the following configurations:

- Windows virtual machines that are hosted on Hyper-V 2012, 2012R2, and 2016
- VHD and VHDX data
- Dynamic and fixed disk type data
- FAT32, NTFS, and ReFS volume data
- Hyper-V VMs over SMB

Before you perform GLR, consider the following limitations:

- You can browse and recover one virtual machine at a time.
- You cannot simultaneously recover the same backup to multiple clients.
- While you perform GLR, you must not mount another virtual machine. Otherwise, NMM unmounts the first virtual machine, with which you lose access to the contents till you remount the virtual machine.
- You cannot recover data from a Windows Server 2012 Storage Spaces disk on a virtual machine.
- You cannot recover deduplicated data.
- You cannot recover differencing disk with parent and child hard disk on different hard drives.

While you perform the steps in the following procedure, you can view the progress of the recovery operation, and check for error messages on the **Monitor** page. The error messages pertain to attaching hard disks, recognizing virtual machines, and expanding virtual machines.

Procedure

1. Open the NetWorker User for Microsoft GUI.
2. Click the icon that is beside the **NetWorker server** field, and specify the NetWorker server.
3. From the **Client** list, select the Hyper-V server (standalone or clustered) that contains the virtual machines that you want to recover.

If the Hyper-V server does not appear in the **Client** list, add it to the list:

- a. On the menu bar, click **Options > Configure Options**.
- b. In the **Configuration Options** dialog box, click the icon beside the **Client name** field.
- c. In the **Select Viewable Clients** dialog box:
 - a. From the **Available clients on <NetWorker_server_name>** list, select the Hyper-V Server, and then click **Add**.
The selected Hyper-V Server appears in the **Clients to list on menu bar** list.
 - b. Click **OK**.
- d. In the **Configuration Options** dialog box, click **OK**.
4. From the left panel, click **Recover > Hyper-V Recover Session > Granular Level Recovery**.
5. In the middle panel, on the **Browse** tab, expand **Microsoft Hyper-V**, and then select the Hyper-V virtual machine that contains the folders or items that you want to recover.

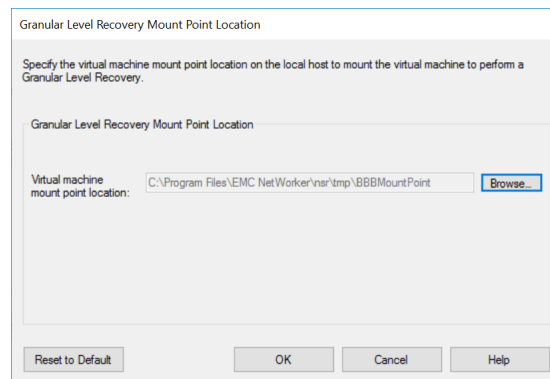
Though all the virtual machines (Windows, Linux, and so on) that are backed up for Hyper-V client appear, you can perform GLR of only Windows virtual machines.

6. To select a particular version or backup time of the selected virtual machine:
 - a. Right-click the virtual machine, and then select **Versions**.
 - b. In the **NetWorker Versions** dialog box:
 - a. Select the backup time.
 - b. Select **Use selected item backup time as new browse time**.
 - c. Click **OK**.

7. Right-click the virtual machine, and then select **Mount**.

The **Granular Level Recovery Mount Point Location** dialog box appears.

Figure 16 GLR mount point location



8. In the **Granular Level Recovery Mount Point Location** dialog box, the **Virtual machine mount point location** field contains the default mount point location for the items that you want to recover, on the local host. The default path is typically `<NetWorker_installation_path>\nsr\ntp\BBBMountPoint`. For example, `C:\Program Files\EMC NetWorker\nsr\ntp\BBBMountPoint`.

To specify a different path, click **Browse**, and then select the path on the local host.

In the **Granular Level Recovery Mount Point Location** dialog box, click **OK**.

NMM mounts the virtual machine. If another virtual machine is mounted for GLR, the GUI notifies you that it unmounts the first virtual machine. To continue to mount the second virtual machine, click **OK**. Otherwise, to leave the first virtual machine mounted, click **Cancel**.

9. After the mount operation succeeds, in the middle panel, on the **Browse** tab, select the mounted virtual machine.

The list of the VHDs that the mounted virtual machine contains appear in the right panel.

10. To view required volumes of a virtual machine to recover, right-click the virtual machine, and then select **Required volumes**.

In the **Required NetWorker Volumes** dialog box, review the list of volumes, and then click **OK**.

11. To search for a specific item within a virtual machine:

a. Perform one of the following steps:

- On the **Browse** tab, expand **Microsoft Hyper-V**, and then select the virtual machine. Click the **Search** tab.
- On the **Browse** tab, expand **Microsoft Hyper-V**, right-click the virtual machine, and then select **Search for**.

The **Path** field displays the virtual machine path.

b. (Optional) In the **Name** field, type the name of the search item. You can refine the search by using any of the following types of search:

- Literal match (case-insensitive): Type `abc` to return `abc`, `ABC`, or `AbC` but not `abcd` or `ABCD`.
- Literal match (case-sensitive): Type `"abc"` to return `abc`, but not `AbC` or `abcd`.
- Name contains (case-insensitive): Type `%abc%` to return `abc`, `abcd`, `ABCD`, or `xyzABCde`.
- Name starts with (case-insensitive): Type `abc%` to return `abcd` or `ABCde`, but not `xyzABCde`.
- Name ends with (case-insensitive): Type `%abc` to return `xyzAbc`, but not `ABCde`.
- Single-character match search by using the `?` wildcard:
 - Type `?` to return single character entries and drive volumes, such as `C` or `D`.
 - Type `<writer_name>?Writer` to return the `<writer_name> Writer`, for example, `abc Writer`.
- Multiple-character match search by using the `*` wildcard:
 - Type `*.txt` to return all entries with a `.txt` extension.
 - Type `*` to return all items within the selected container.
 - Type `*writer*` to return all writers.
- Search by using the `*` and `?` wildcards: Type `*???*writer*` to return the `abc Writer`.

c. Click **Search**.

The **Result** panel displays the search results.

Note

The **Search for** feature is enabled only when the source virtual machine is mounted.

12. Select the items that you want to recover.

13. On the **Hyper-V Recover Session** toolbar, click **Recover**.

The **Hyper-V Virtual Machine Recovery** wizard appears with the **Hyper-V Virtual Machine Granular Level Recovery Options** page.

Figure 17 Hyper-V virtual machine GLR options page

14. On the **Hyper-V Virtual Machine Granular Level Recovery Options** page:
 - a. Specify the following fields:
 - **Select the diagnostic output level:** From this list, select a debug level, **Normal** (or 1) through **9** according to the amount of GLR debug information that you want the GUI to log. The debug level **9** logs more information.
 - **Specify the destination path:** This field contains the default destination path, to which you want to recover the items that you have selected. The default path is typically <NetWorker_installation_path>\nsr\tmp\HyperVGlrRestore. **For example**, C:\Program Files\EMC NetWorker\nsr\tmp\HyperVGlrRestore. **To specify a different path**, click **Browse**, and then select the path.
 - b. Click **Next**.
15. On the **Hyper-V Virtual Machine Recovery Options Summary** page:
 - a. Review the settings that you have configured.
To modify the settings, click **Back**.
 - b. To start the recovery operation, click **Start**.
NMM recovers the items to the destination path by creating the original folder hierarchy.

Note

When the recovery is in progress, you cannot perform other tasks in the **Hyper-V Virtual Machine Recovery** wizard.

16. To view the status of the recovery, in the left panel of the NetWorker User for Microsoft GUI, click **Monitor**.

After you finish

After you recover the files or folders that are local to the proxy server, you must manually move the files according to your requirement.

You can perform browse and recovery actions on the mounted virtual machine.

You can either unmount the virtual machine or perform another Hyper-V GLR operation.

Using NMM 9.1 or later to recover the backups that were performed by using NMM 8.2.x

To use NMM 9.1 or later to recover the backups that were performed by using NMM 8.2.x, upgrade NMM to 9.1 or later by ensuring that you select the **Restore of NMM 8.2.x and Earlier Backups (VSS workflows)** option in the NMM 9.1 or later installer. In the case of federated backups, you must perform this upgrade operation on all the cluster nodes.

The *NetWorker Module for Microsoft Installation Guide* provides information about how to upgrade, and the **Restore of NMM 8.2.x and Earlier Backups (VSS workflows)** option in the installer.

CHAPTER 5

File Level Recoveries

This chapter includes the following sections:

- [Introduction.....](#) 96
- [Performing a browser-based file level restore.....](#) 97
- [Performing a directed file level restore.....](#) 98
- [Monitoring file level restores.....](#) 99
- [Hyper-V FLR web UI log files.....](#) 100

Introduction

The NMM Hyper-V File Level Restore (FLR) user interface allows NetWorker administrators to restore files that are stored on any Hyper-V virtual machine that is configured for NetWorker protection. As a NetWorker administrator, you can select Hyper-V backups to restore from, browse and search for files, select files and folders, and perform either browser-based or directed recoveries. The Hyper-V FLR user interface is fully web-based and runs in a web browser.

The Hyper-V FLR web UI requires the following software and privileges:

- NetWorker 9.0.1 or later with NetWorker Authentication Service configured.
- The `lgtoadapt.rpm` package installed on the NetWorker server if you use Linux.
- To access the Hyper-V FLR web UI, you can log in as a NetWorker administrator, who is part of the NetWorker Security Administrator and Application Administrator user groups.
The Hyper-V FLR web UI provides access to any Hyper-V data that has been backed up on the NetWorker server.
- To perform Hyper-V FLR recoveries without logging in as a NetWorker administrator:
 - Add the remote client to the Backup Operators group on the Hyper-V Server host.
 - Install Hyper-V Integrated Services on all virtual machines on the remote host.
 - In PowerShell, run the `enable-psremoting` cmdlet. During redirected recovery, NMM requires this cmdlet so that it can fetch the list of running VMs and open live VMs from the remote Hyper-V server.
- NMM Recovery Agent (NMM RA), which is installed as part of the standard NMM installation package. To perform recoveries to a Hyper-V virtual machine, install NMM Proxy RA. The *NetWorker Module for Microsoft Installation Guide* provides information.
The proxy server communicates with NW Adapter, mounts virtual machine backups, and then serves data for browser-based and directed restores. The proxy server also queries Hyper-V servers to get a list of running VMs on each Hyper-V server. The proxy server can be any host that has required permissions on the NetWorker server to access virtual machine backups from different Hyper-V servers that it manages. The proxy server requires Failover Clustering to mount virtual machines on a CSV.

The Hyper-V FLR web UI supports the following restore workflows:

- Browser-based restore: Restores the backup to a local folder or network location.
- Directed restore: Restores items to a specific virtual machine and location.

Required ports for Hyper-V File Level Restore GUI

The Hyper-V FLR service uses the following ports:

- 10000 HTTP
- 11000 Secure HTTPS
- 10099 Cache Service
- 10024 Persistence Service

- 9090 NetWorker Authentication Service

Performing a browser-based file level restore

The Hyper-V FLR web UI browser-based file level restore process follows a step-by-step, wizard-like workflow similar to the directed file level restore process. You can use the browser-based restore option to select items to recover and then download the recovered files.

Procedure

1. From a supported web browser, type the URL of the Hyper-V FLR web UI:

`http://server_name:http_service_port`
where:

- *server_name* is the name of the Hyper-V FLR web UI.
- *http_service_port* is the port for the embedded HTTP server. The default HTTP port is 10000.

For example: `http://houston: 10000`

2. Log in using NetWorker administrator credentials.

The **Proxy URL** dialog box appears. This dialog box appears the first time you log in.

3. In **Proxy URL** dialog box, provide the cluster name that was used for Hyper-V cluster level backup. Provide Hyper-V server name for standalone.
4. In the **Select the Hyper-V Server** area, select the Hyper-V server that contains the virtual machine you want to restore from by using one of the following options:
 - Select a server from the **Select the Hyper-V Server** list.
 - Click the sort button to sort the list of servers alphabetically.
 - Type a search term in the search field. The results display in the **Select the Hyper-V Server** list. Select a server from this results list. To clear the search term, click the "x" in the search field.
5. In the **Select the Virtual Machine** area, select the virtual machine that you want to restore from by using one of the following options:
 - Select a virtual machine from the list that is displayed in the **Select the Virtual Machine** list.
 - Click the sort button to sort the list of VMs alphabetically.
 - Type a search term in the search field. The results display in the **Select the Virtual Machine** list. Select a virtual machine from this results list. To clear the search term, click the "x" in the search field.
6. In the **Select a backup containing the items for restore** area, select the backup that you want to restore from by using one of the following options:
 - Select a backup that is displayed in the list.
 - To sort the list of backups alphabetically, click the sort button.
 - To filter the list of backups by date, Click the filter button, select the **Date** filtered view, select **Before** or **After**, and then select a backup date. To clear the filter. Click the filter button and select **None**. Dates are displayed in GMT time.

7. On the **Restore Items** page, browse or search for items to restore.
8. Select an item to restore by doing one of the following:
 - Double-click the item. The item displays in the **Items to Restore** area.
 - Drag and drop the item into the **Items to Restore** area.

To remove items from the **Items to Restore** area, click the "x" next to the item.
9. On the **Restore Options** page, verify that the **Restore to a browser download location** setting is **Yes** and then click **Finish**.
10. Check the restore status in the **Restore Monitor** pane.

To cancel a pending restore, click or tap **Cancel**.
11. After the download completes, click the **Download** button on the restored items status pane.
12. In the **Save As** window that displays, browse to a location to save the restored items and click **Save**.

Results

The restored items are downloaded to the location you specified.

Performing a directed file level restore

The Hyper-V FLR web UI directed file level restore process follows a step-by-step, wizard-like workflow similar to the browser-based file level restore process. Use the Hyper-V FLR web UI to select items to restore and then specify a destination file path for the restored items.

Procedure

1. Log in using NetWorker administrator credentials.

The **Proxy URL** dialog box appears. This dialog box appears the first time you log in.
2. In **Proxy URL** dialog box, provide the cluster name that was used for Hyper-V cluster level backup. Provide Hyper-V server name for standalone.
3. In the **Select the Hyper-V Server** area, select the Hyper-V server that contains the virtual machine you want to restore from by using one of the following options:
 - Select a server from the **Select the Hyper-V Server** list.
 - Click the sort button to sort the list of servers alphabetically.
 - Type a search term in the search field. The results display in the **Select the Hyper-V Server** list. Select a server from this results list. To clear the search term, click the "x" in the search field.
4. In the **Select the Virtual Machine** area, select the virtual machine that you want to restore from by using one of the following options:
 - Select a virtual machine from the list that is displayed in the **Select the Virtual Machine** list.
 - Click the sort button to sort the list of VMs alphabetically.
 - Type a search term in the search field. The results display in the **Select the Virtual Machine** list. Select a virtual machine from this results list. To clear the search term, click the "x" in the search field.

5. In the **Select a backup containing the items for restore** area, select the backup that you want to restore from by using one of the following options:
 - Select a backup that is displayed in the list.
 - To sort the list of backups alphabetically, click the sort button.
 - To filter the list of backups by date, Click the filter button, select the **Date** filtered view, select **Before** or **After**, and then select a backup date. To clear the filter. Click the filter button and select **None**. Dates are displayed in GMT time.
 6. On the **Restore Items** page, browse or search for items to restore.
 7. Select an item to restore by doing one of the following:
 - Double-click the item. The item displays in the **Items to Restore** area.
 - Drag and drop the item into the **Items to Restore** area.

To remove items from the **Items to Restore** area, click the "x" next to the item.
 8. On the **Restore Options** page, verify that the **Restore to a browser download location** setting is **No** and then click **Finish**.
 9. In the **Select the Virtual Machine** area, select the destination virtual machine by doing one of the following:
 - Select a virtual machine from the list that is displayed in the **Select the Virtual Machine** list.
 - Click the sort button to sort the list of virtual machines alphabetically.
 - Type a search term in the search field. The results display in the **Select the Virtual Machine** list. Select a virtual machine from this results list. To clear the search term, click the "x" in the search field.
 10. In the **Restore to location** area, browse to the wanted destination file path for the restored items by doing one of the following:
 - Select a location in the list that is displayed.
 - Select a location in the **Restore to location** list, then type a search term in the search field and click **Go**. The results display in the **Restore to location** list. To clear the search term, click the "x" in the search field.
 11. Click **Finish**.
 12. Check the restore status in the **Restore monitor** pane.
- To cancel a pending restore, click **Cancel**.

Results

The items are restored to the specified destination file path.

Monitoring file level restores

You can monitor Hyper-V FLR restores from any page in the web UI by using the Status bar. Collapse or expand the Status bar to hide or display the **Restore Monitor** toolbar. The **Restore Monitor** toolbar allows you to view restore details in tile or list view and display an expanded view of a restore status.

Status bar

When collapsed, the Status bar displays condensed information about running, pending, successful, and failed restores. Clicking or tapping the Status bar displays the **Restore Monitor** toolbar.

Restore Monitor toolbar

The **Restore Monitor** toolbar displays restore statuses in either tile view or list view. To change the view, select **Change View** and then select a view. To refresh the **Restore Monitor** with the latest restore information, click or tap **Refresh**. You can also filter the restore statuses to display only successful, failed, or running restores, and you can sort the restore statuses by start time, end time, progress, or status.

Clicking or tapping a restore status displays an expanded view, which provides additional details and available actions. Only one restore status can be expanded at a time. To close the expanded view, double-click or tap the expanded item, or click or tap another item.

Hyper-V FLR web UI log files

If you encounter errors while using the Hyper-V FLR web UI, you can check various Hyper-V FLR web UI, NMM, and NetWorker log files.

Check the log files in the following order:

1. Hyper-V FLR UI

- **Linux:** /nsr/logs/hyperv-flr-ui/hyperv-flr-ui.log
- **Windows:** C:\Program Files\EMC NetWorker\nsr\logs\hyperv-flr-ui\hyperv-flr-ui.log

2. NMM RA (Windows only)

- C:\Program Files\EMC NetWorker\nsr\logs\nsrnmra.log
- C:\Program Files\EMC NetWorker\nsr\logs\nsrnmproxyra.log

3. NetWorker adapter

- **Linux:** /opt/nsr/nsrmq/logs/nsrmq.log
- **Windows:** C:\Program Files\EMC NetWorker\nsr\logs\nsrmq\logs\nsrmq.log

4. RabbitMQ message bus

- **Linux:** /opt/nsr/rabbitmq-server-3.2.4/var/log/rabbitmq/rabbitmq.log
- **Windows:** C:\Windows\System32\config\systemprofile\AppData\Roaming\RabbitMQ\log\rabbit.log

5. NetWorker Server

- **UNIX:** /nsr/logs/daemon.raw
- **Windows:** C:\Program Files\EMC NetWorker\nsr\logs\daemon.raw

6. NetWorker authentication service

- **Linux:** /opt/emc/authc/tomcat/logs/catalina.out
- **Windows:** C:\Program Files\EMC\Authc\tomcat\logs\catalina.out

CHAPTER 6

Data Protection Add-in for SCVMM

This chapter includes the following sections:

• Overview of the Data Protection Add-in for SCVMM	102
• How the Data Protection Add-in works with SCVMM	104
• Installation and uninstallation	107
• Preferences	111
• Data Protection Add-in overview data	114
• SCVMM Recoveries	120
• Monitoring	126
• Troubleshooting	127

Overview of the Data Protection Add-in for SCVMM

The Data Protection Add-in for SCVMM leverages the System Center Virtual Machine Manager (SCVMM) Add-in extension support to enable NetWorker client Hyper-V virtual machine recoveries within the SCVMM console.

The Data Protection Add-in enables you to perform NMM Hyper-V recoveries within the SCVMM console. You can view and recover all current SCVMM-managed virtual machines that have NMM conventional backups. The Data Protection Add-in supports recoveries of Hyper-V virtual machines in cloud, cluster, host, host group, and virtual machine contexts.

You can perform recoveries of Hyper-V virtual machines to the original location or to an alternate host location.

When you create a virtual machine by using SCVMM, the virtual machine is registered with the naming convention, `SCVMM <virtual_machine_name> Resource` in the failover cluster manager. However, the virtual machine is registered with its actual name, that is, `<virtual_machine_name>` in the Hyper-V manager.

When you back up such a virtual machine, it is backed up with its actual name. When you recover the virtual machine, it is re-registered with its actual name in the failover cluster manager.

Recoveries

The Data Protection Add-in feature set supports recovery of Hyper-V virtual machines protected by NetWorker servers. The Data Protection Add-in supports recoveries of conventional backups to the original Hyper-V server on which the virtual machine was backed up or to an alternate Hyper-V server.

The Data Protection Add-in can be used in the following SCVMM configurations:

- SCVMM console on the same host as the SCVMM server
- SCVMM console on a different host from the SCVMM server

To perform recoveries by using the Data Protection Add-in, you must have the required privileges for the client to which you recover the virtual machine. [Required privileges](#) on page 103 provides details about SCVMM privileges. The *NetWorker Administration Guide* provides details about the required NetWorker server privileges.

Backups

The procedure to perform a scheduled backup of a virtual machine that is managed by SCVMM is the same as the procedure to perform a scheduled backup of a standard physical host. Create and configure a NetWorker client for the Hyper-V server to protect the SCVMM virtual machine.

The [Backups](#) on page 33 chapter in this guide, and the *NetWorker Administration Guide* provide information about how to schedule and manage backups.

Supported versions

The Data Protection Add-in supports System Center 2016 and System Center 2012 R2 Virtual Machine Manager.

Note

The *NetWorker E-LAB Navigator*, which is available at <https://elabnavigator.emc.com/eln/elhome>, provides the most up-to-date information about the Windows Server versions that NMM supports. The Data Protection Add-in version must match the NetWorker and NMM client versions.

The Data Protection Add-in is compatible with the following operating systems when imported into the System Center 2016 Virtual Machine Manager Console:

- Windows Server 2016 (64 bit) Standard and Datacenter
- Windows 7 SP1 or later (64 bit or 32 bit) Professional, Enterprise, and Ultimate
- Windows 8 (64 bit or 32 bit) Professional and Enterprise
- Windows 8.1 (64 bit or 32 bit) Professional and Enterprise

The Data Protection Add-in is compatible with the following operating systems when imported into the System Center 2012 R2 Virtual Machine Manager Console:

- Windows Server 2012 (64 bit) Standard and Datacenter
 - Windows Server 2012 R2 (64 bit) Standard and Datacenter
 - Windows 7 SP1 or later (64 bit or 32 bit) Professional, Enterprise, and Ultimate
 - Windows 8 (64 bit or 32 bit) Professional and Enterprise
 - Windows 8.1 (64 bit or 32 bit) Professional and Enterprise
-

Note

Hyper-V Server with Window Server 2016 can not be imported to System Center 2012 R2 Virtual Machine Manager Console.

Software dependencies

The Data Protection Add-in requires the following software:

- The Data Protection Add-in and the NetWorker base and extended client software must be installed on the SCVMM console host. The Data Protection Add-in must match the NetWorker and NMM client versions.
- The NetWorker and NMM 9.0.1 or later client software must be installed on the Hyper-V server to which the virtual machine is recovered.
- The Data Protection Add-in requires access to a NetWorker 9.0.1 or later server.

Required privileges

To perform recoveries, you must be a member of certain SCVMM roles and have certain privileges.

To perform recoveries, you must:

- Be a member of the SCVMM Administrator or Fabric Administrators SCVMM roles.
- Have write access to the folder where the cached datafiles are stored. For example: `C:\Users\%current user%\AppData\Local\EMC\NetWorker\SCVMM`.
- Have NetWorker directed recovery privileges, which requires the following:

- The Data Protection Add-in host is a client of the NetWorker server that contains the backup information. This administering client can be a different platform from the source and destination clients.
- Use the local root or Administrator account to start the recovery. Ensure that the user account is a member of one of the following:
 - The Operators, Application Administrators, Database Administrators, or Database Operators User Group.
 - A customized User Group with the following privileges on the NetWorker server:
 - Remote Access All Clients
 - Operate NetWorker
 - Monitor NetWorker
 - Operate Devices and Jukeboxes
 - Backup Local Data
 - Recover Local Data
 - Recover Remote Data

Installation and configuration overview

To install the Data Protection Add-in, an SCVMM administrator and each user must perform required steps.

Procedure

1. An SCVMM administrator must perform the following steps:
 - a. [Installing SCVMM and the SCVMM console](#) on page 108
 - b. [Installing the Data Protection Add-in](#) on page 108
2. Each user must perform the following steps:
 - a. [Importing the Data Protection Add-in](#) on page 109
 - b. [Activating the Data Protection Add-in](#) on page 109

After you finish

To uninstall the Add-in, an SCVMM administrator, and each user must perform the steps that are described in [Uninstalling the Data Protection Add-in](#) on page 110.

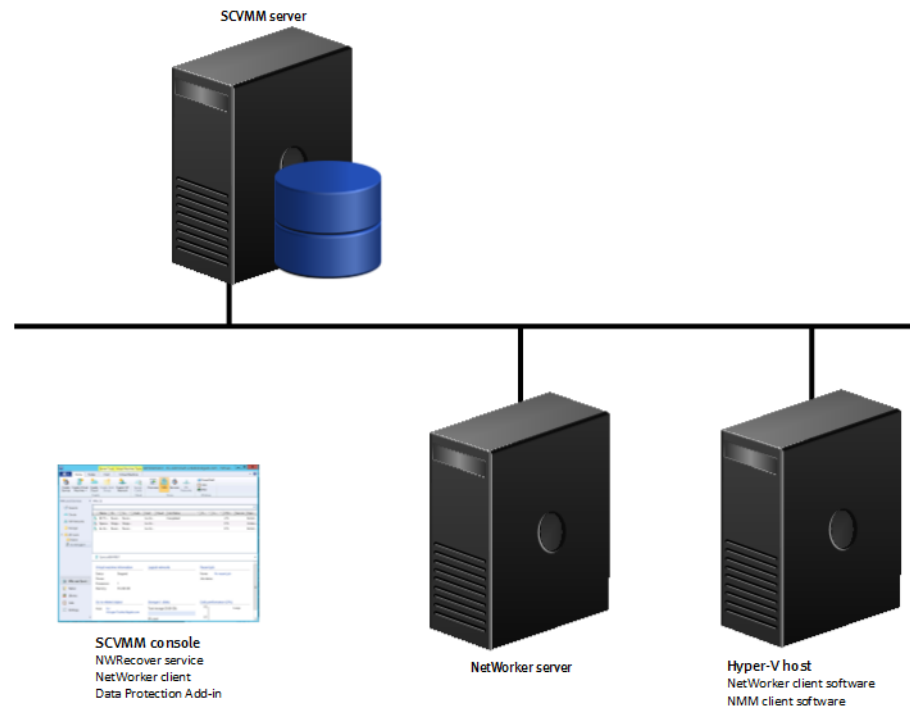
How the Data Protection Add-in works with SCVMM

The NetWorker client and NMM client software must be installed on each Hyper-V physical host. The SCVMM console can be installed on a separate server or on the SCVMM server. However, the NetWorker client must be installed on the SCVMM console host.

Note

The EMC Data Protection Add-in uses a hostname, as displayed on the SCVMM console, when it tries to identify corresponding NetWorker clients on the NetWorker servers. If the hostname in SCVMM does not match its actual server name, then the Add-in is unable to find the corresponding NetWorker Client. Ensure that the hostnames in SCVMM match their actual server names.

The following figure illustrates the Data Protection Add-in architecture.

Figure 18 Data Protection Add-in architecture

Workflows overview

The following sections describe common workflows for the Data Protection Add-in.

Initialize the SCVMM console or change context

When you launch the SCVMM console or change context within the console, the Data Protection Add-in does the following:

Procedure

1. Accesses the SCVMM server to obtain a list of virtual machines for the context you selected.
2. Displays virtual machines for the selected context that have been backed up on servers in the **Preferred NetWorker servers** list.

Refresh the Data Protection Add-in display

When you click the **Refresh** button on any page in the Data Protection Add-in, the Data Protection Add-in does the following:

Procedure

1. Accesses the SCVMM server to obtain a list of all hosts.
2. Accesses the NetWorker server to obtain a list of all clients and save sets.
3. Accesses the SCVMM server to obtain a list of all virtual machines in the current context.
4. Displays the updated protection information on the **Overview** page and virtual machines available for recovery on the **Recover** page.

Perform a recovery

When you perform a recovery, the following occurs:

Procedure

1. The Data Protection Add-in passes the virtual machine, backup time, and destination options you selected to the NWRecover service. The NWRecover service starts the recovery process.
2. The NWRecover service runs a remote agent on the Hyper-V server and passes the required information.

The NWRecover service posts recover messages to the **Data Protection Add-in Monitor** page

3. The remote agent performs the requested recovery.

During the recovery process, the NWRecover service updates the log shown in the **Monitor** page as well as the Windows event log under **Applications and Services > NetWorker Recovery Service**.

Results

The NWRecover service posts the recover success message in the monitor log and the Windows event log.

GUI overview

The Data Protection Add-in consists of the **Overview**, **Preferences**, **Recover**, and **Monitoring** pages.

- **Overview**—Displays the protection status for all virtual machines in the current SCVMM context.
- **Preferences**—Allows you to specify NetWorker servers, set the refresh rate, and set the debug level.
- **Recover**—Allows you to perform recoveries and view virtual machines available for recovery.
- **Monitoring**—Allows you to view in-progress and completed operations.

After you import the Data Protection Add-in, when you select the **All Hosts** or **Cloud** scope in the SCVMM console, the EMC Data Protection Add-in button displays in the **SCVMM** ribbon within the **VMs and Services** context.

If you select a non-supported scope (within the **VMs and Services** context), the **Data Protection Add-in** button is disabled.

SCVMM user roles and allowed actions

The Data Protection Add-in is cloud and tenant-aware, so you can only recover virtual machines to which you have access. You cannot direct a recovery to a Hyper-V server to which you do not have access.

The following table lists the supported SCVMM User Roles and the actions that the Data Protection Add-in allows for each supported role.

Table 24 SCVMM user roles and actions allowed by the Data Protection Add-in

Role	Actions allowed
Fabric Administrator (Delegated Administrator)	Can see all virtual machines, hosts, and clouds. Can recover all virtual machines managed by SCVMM to original and alternate locations.
Tenant Administrator	Can see and recover virtual machines within the private cloud they manage. Only recovery to original location is supported. On the Recover page, unable to see the Hyper-V Host and Recover Destination columns.
Read-Only Administrator	Can see the virtual machines and hosts within the private cloud they manage. No recovery operations are allowed.
Application Administrator (Self-Service Administrator)	Can see and recover virtual machines within the private cloud they manage. Only recovery to original location is supported. On the Recover page, unable to see the Hyper-V Host and Recover Destination columns.

Supported scopes and contexts

The Data Protection Add-in supports the following SCVMM scopes:

- Cloud
- Cluster
- Host (clustered and stand-alone)
- HostGroup
- Virtual machine

The Microsoft website provides more information about SCVMM scopes.

Installation and uninstallation

This section describes the Data Protection Add-in required components and the order in which they must be installed and configured.

The SCVMM administrator must install the components in the following order:

1. SCVMM and SCVMM console
2. SCVMM update rollups
3. NetWorker base client kit
4. NetWorker extended client kit
5. Data Protection Add-in

Note

The NetWorker base and extended client kits must be installed on the SCVMM console host and must be the same version as the Data Protection Add-in.

After the SCVMM administrator has installed these components, each user imports and activates the Data Protection Add-in.

To uninstall the Data Protection Add-in, each user removes the Data Protection Add-in from the SCVMM console, and an SCVMM administrator uninstalls the Data Protection Add-in.

Installing SCVMM and the SCVMM console

Download and install SCVMM from the Microsoft website. Install the SCVMM console so that it is available for all users. This installation requires system administrator privileges.

Installing the Data Protection Add-in

To install the Data Protection Add-in, you access the installation files from a DVD disk or EMC Online Support. To install the Data Protection Add-in on the SCVMM server, you must have local administrator privileges.

NetWorker client packages and the Data Protection Add-in must be installed on the SCVMM console host. The Data Protection Add-in requires the NetWorker base client and extended client packages to be installed on the SCVMM console host before installing the Add-in itself. The Data Protection Add-in must match the NetWorker and NMM client versions. The *NetWorker Installation Guide* provides details about the NetWorker client package installations.

Because the Data Protection Add-in does not have built-in foreign language support, only install the English language pack on the NetWorker client for use with the SCVMM add-in.

Procedure

1. To access the Data Protection Add-in software from a local DVD disk:
 - a. Log in as an administrator or equivalent on the NetWorker client.
 - b. Insert the Data Protection Add-in DVD disk into the DVD drive.
 - c. Run `EMC_Data_Protection_UI_Addin_for_SCVMM.msi` directly from the DVD.
 - d. Accept the default values during the installation.
2. To access the Data Protection Add-in software from EMC Online Support:
 - a. Log in as administrator or equivalent on the NetWorker client.
 - b. Browse to EMC Online Support (<http://support.emc.com>).
 - c. Browse to the Downloads for NetWorker Module for Microsoft page.
 - d. Download the 32-bit or 64-bit Data Protection Add-in software Zip file to a temporary folder that you create.
 - e. Extract the Zip file to the temporary folder.
 - f. Run `EMC_Data_Protection_UI_Addin_for_SCVMM.msi`.
 - g. Accept the default values during the installation.

Results

The installer places a Data Protection Add-in Zip file in the public user documentation folder and installs the required NWRecover Service. The NWRecover Service automatically starts during the installation process.

The default installation path for the Data Protection Add-in Zip file is:

`C:\Users\Public\Documents\EMC NetWorker\nsr\addins
\VMM_DataProtection\`. If you encounter any issues while installing or importing

the Data Protection Add-in, then ensure that you have read and write permission for all folders in this path.

Make note of the Data Protection Add-in .zip file installation path as it is used in [Importing the Data Protection Add-in](#) on page 109. The default installation path for the NWRecover Service is:

C:\Program Files\EMC NetWorker\nsr\addins\VMM_DataProtection.

Importing the Data Protection Add-in

Each Data Protection Add-in user must import the Data Protection Add-in. The users must have write access to the folder where the cached datafiles are stored. For example: C:\Users\%current user%\AppData\Local\EMC\NetWorker\SCVMM.

Procedure

1. Launch the SCVMM console and connect to a Virtual Machine Manager server. The console opens.
2. In the workspaces pane, click **Settings**.
3. In the navigation pane, click **Console Add-ins**.
4. If a previous version of the Data Protection Add-in exists, select it and click **Remove** on the SCVMM ribbon.
5. On the SCVMM ribbon, click **Import Console Add-in**.
6. In the **Import Console Add-in** wizard, browse to the folder in which you installed the Data Protection Add-in Zip file.
7. Select **EMC.DP.ScvmAddIn.zip** and click **Open**. For example: C:\Users\Public\Documents\EMC NetWorker\nsr\addins\VMM_DataProtection.
8. To continue installing, select the checkbox and click **Next**.
9. Click **Finish** and then click **Close** to close the **Jobs** window that displays.

Results

If an error message displays, delete the pre-existing add-in folder. For example: C:\Program Files\Microsoft System Center 2016\Virtual Machine Manager\bin\AddInPipeline\AddIns\<domain_username>.

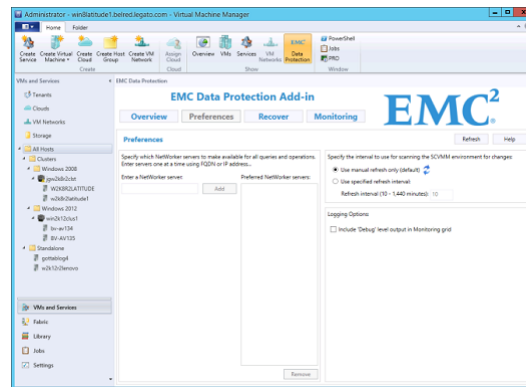
Activating the Data Protection Add-in

After you install SCVMM, the SCVMM console, the SCVMM update rollups, and the Data Protection Add-in, you must activate the Data Protection Add-in.

Procedure

1. In the workspace pane of the SCVMM console, click **VMs and Services**.
2. In the navigation pane, select a host or cluster.
3. On the SCVMM ribbon, click **EMC Data Protection**.

After about 5-10 seconds, the main content area of the console will be replaced by the Data Protection Add-in.

Figure 19 Data Protection Add-in Preferences page

When a user launches the Data Protection Add-in for the first time, the **Preferences** page displays. After initial configuration and refresh, subsequent launches of the Add-in display the **Overview** page first.

Uninstalling the Data Protection Add-in

To uninstall the Data Protection Add-in, each user must remove the Data Protection Add-in from the SCVMM console, and an SCVMM administrator must uninstall the Data Protection Add-in. If no users perform other NetWorker operations on this computer, you can also uninstall the NetWorker software. These tasks can be performed in any order.

Removing the Data Protection Add-in from the SCVMM console

Each user must remove the Data Protection Add-in from the SCVMM console. Removing the Data Protection Add-in from the SCVMM console removes all components that are copied to the SCVMM AddIn folder during the import process, but not the originally downloaded Data Protection Add-in .zip file itself.

Note

Removing the Data Protection Add-in only affects individual users. Other users who imported the Add-in are not affected.

Procedure

1. In the SCVMM console, click the **Settings** workspace.
2. Click the **Console Add-ins** setting.
3. In the list of installed Add-ins, select **EMC Data Protection Add-in**.
4. On the top ribbon, click **Remove**.
5. On the confirmation window that displays, click **Yes**.

After you finish

The Data Protection Add-in creates persistent data cache files during the refresh operation. These files are created for each user. If a user removes the Add-in and is not expected to upgrade or otherwise re-import the add-in in the future the files can be manually removed from the following folder: `C:\User\<user name>\AppData\Local\EMC\NetWorker\SCVMM`.

Uninstalling the Data Protection Add-in by using Windows Program and Features

An SCVMM administrator must uninstall the Data Protection Add-in from the SCVMM server. Uninstalling the Data Protection Add-in ensures the Data Protection Add-in

(.zip file) is removed from the SCVMM console host and ensures that the Data Protection service is stopped and uninstalled.

Note

This step affects all users who imported the Data Protection Add-in. If the Data Protection Add-in is uninstalled, no users can perform a recovery by using the Data Protection Add-in. Verify that each SCVMM console user has removed the Data Protection Add-in before uninstalling.

Procedure

1. For Windows Server 2012 or Windows 8 or later: Click **Control Panel** and then click **Programs and Features**.
2. For Windows 7 or earlier: Click **Control Panel** and then click **Uninstall a program**.
3. Select **EMC Data Protection UI Addin for SCVMM**.
4. Click **Uninstall**.

Upgrading the Data Protection Add-in

To upgrade the Data Protection Add-in, you must complete the uninstallation procedures to uninstall the current version and then complete the installation procedures to install the new version.

Before you begin

Before upgrading the Data Protection Add-in, ensure that the NetWorker and NMM client software and the NetWorker Server software are compatible with the Data Protection Add-in. The Data Protection Add-in version must match the NetWorker and NMM client versions. The *NetWorker E-LAB Navigator*, which is available at <https://elabnavigator.emc.com/elab/elhome>, provides the most up-to-date information about supported versions.

Procedure

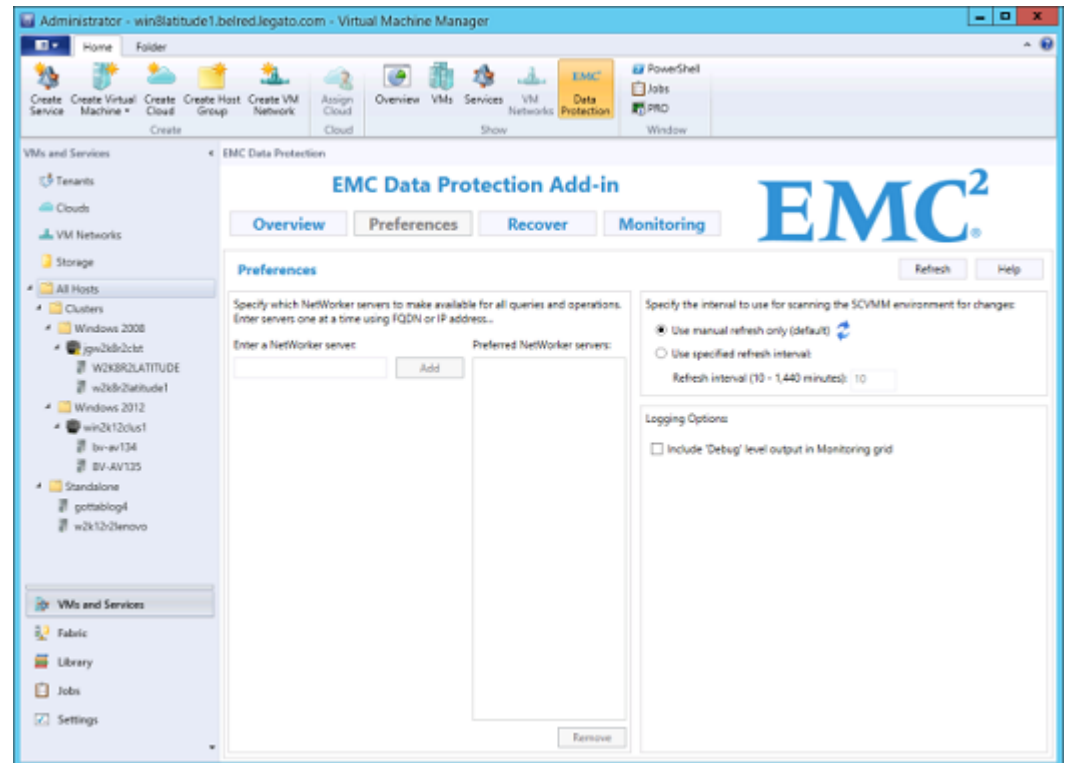
1. Obtain the new Data Protection Add-in installer MSI file from the EMC Support site.
2. For all users, remove the existing Data Protection Add-in from the SCVMM console.
3. Follow the steps that are described in [Uninstalling the Data Protection Add-in](#) on page 110.
4. Follow the steps that are described in [Installing the Data Protection Add-in](#) on page 108.
5. Follow the steps that are described in [Importing the Data Protection Add-in](#) on page 109.
6. Follow the steps that are described in [Activating the Data Protection Add-in](#) on page 109.
7. In the **Upgrade Successful** window, click **OK**.
8. Click **Refresh** to repopulate the Data Protection Add-in.

Preferences

After completing the installation process, you must configure the Data Protection Add-in to access the NetWorker servers that contain virtual machine backups for

recovery. You can also set the refresh frequency and specify the logging debug level. After making any configuration changes to the SCVMM environment, perform a Refresh operation in the Data Protection Add-in to ensure that the Add-in is displaying current information.

Figure 20 Data Protection Add-in



Adding NetWorker servers

You can search for virtual machine backups on multiple NetWorker servers. Contact the NetWorker administrator to learn which NetWorker servers protect the virtual machines you manage, and then add them to the Data Protection Add-in.

Procedure

1. In the workspaces pane of the SCVMM console, click **VMs and Services**.
2. In the left navigation pane, select the host or cloud you want to manage.
3. In the SCVMM ribbon, click **EMC Data Protection**.
4. In the Data Protection Add-in, click **Preferences**.
5. In the text box next to the **Preferred NetWorker servers** list, type the FQDN or IP address of a NetWorker Server and click **Add**.

The preferred NetWorker servers list is displayed.

Note

Do not use IPv6 addresses, localhost, or 127.0.0.1 as the NetWorker server in the Preferences page.

6. In the notification that displays, click **OK**.
7. Follow the directions in the notification.

8. Click **Refresh** to view the newly added NetWorker server virtual machine protection status on the **Overview** page and the available virtual machine backups on the **Recover** page.

Note

If adding more than one NetWorker server at a time, it is recommended to add all servers before starting the Refresh operation.

Removing NetWorker servers

Procedure

1. In the workspaces pane of the SCVMM console, click **VMs and Services**.
2. In the navigation pane, select the host or cloud you want to manage.
3. In the SCVMM ribbon, click **EMC Data Protection**.
4. In the Data Protection Add-in, click **Preferences**.
5. In the **Preferred NetWorker servers** list, select a server and click **Remove**.

The Data Protection Add-in automatically performs a Refresh operation to display virtual machine data for the remaining NetWorker servers.

Setting the refresh interval

On the **Preferences** page, the Data Protection Add-in provides two options for scanning the SCVMM environment for changes:

- Use manual refresh only—This is the default setting. When you select this option, you must manually scan for changes by clicking the **Refresh** button on any Data Protection Add-in page. With this setting, the Data Protection Add-in does not scan for changes automatically.
- Use specified refresh interval—You can specify the interval at which the Data Protection Add-in automatically refreshes the data. When you select this option, type a refresh interval and click anywhere in the SCVMM console to apply the change. The refresh rate should correspond to how often a virtual machine is backed up in the environment and the amount of time a refresh process takes to complete. If the refresh process does not complete within the interval you specify, lengthen the interval accordingly.

Including debug output for logging purposes

You can choose to include debug output in log files. This can be especially helpful for troubleshooting purposes. To include debug level output, on the **Preferences** page, select the **Include debug level output** checkbox.

Using multiple NetWorker servers that define the same clients and virtual machine save sets

The Data Protection Add-in learns about protected virtual machines by querying the NetWorker servers that are specified on the **Preferences** page. Because the Hyper-V Server is a client of multiple NetWorker servers, if there is conflicting data that pertains to a Hyper-V Server and its virtual machine protection, the Data Protection Add-in can display inconsistent data.

If you use multiple NetWorker servers that define the same Hyper-V clients and virtual machine save sets, it is recommended to change the **Preference** page to one

NetWorker server at a time. This task reduces NMM data protection metric inconsistency on the **Overview** page and protected virtual machine listings on the **Recover** page.

In scenarios, where the **Preferences** page does include NetWorker servers that define the same Hyper-V clients and virtual machine save sets, the Data Protection Add-in arbitrarily chooses information from one NetWorker server if conflicts exist. This behavior prevents scenarios, where a virtual machine is mis-counted for protection metrics or shows twice on the **Recover** page.

Data Protection Add-in overview data

The **Overview** page summarizes the current NMM data protection metrics for the managed virtual machines in the currently selected SCVMM context. For Administrator, Fabric Administrator, and Read-Only Administrator user roles, the Data Protection Add-in displays virtual machine protection status. For Tenant Administrator and Application Administrator (Self-Service Administrator) user roles, the Data Protection Add-in displays virtual machine backup status.

Overview page for Administrator, Fabric Administrator, and Read-Only Administrator user roles

For Administrator, Fabric Administrator, and Read-Only Administrator user roles, the **Overview** page displays multiple sub-panes:

- Clouds, Clusters, Hosts, and Virtual machines sub-panes—These sub-panes list the number of clouds, clusters, hosts, and virtual machines the user role manages within the currently selected SCVMM context.
- Configured for protection—This sub-pane provides protection characteristics for the virtual machines that are protected on the NetWorker servers that are listed on the **Preferences** page.

The pie chart provides the following data about virtual machines:

- VMs excluded from protection—These virtual machines are currently listed in the NSR_EXCLUDE_COMPONENTS attribute for a NetWorker client resource and not protected by another client resource.
- VMs not protected—These virtual machines are not configured for a scheduled backup as part of a NetWorker client resource and not explicitly excluded for backup. Virtual machines with missing status are shown under 'VMs not protected' and can be viewed in the pie chart under VMs not protected section of Overview page.
The tool tip on the "VMs not protected" slice of the pie chart lists the virtual machines names that are not configured in the NetWorker server and have a status of 'Missing' in the SCVMM.
- VMs protected—These virtual machines are configured for scheduled backup as part of a NetWorker client resource.

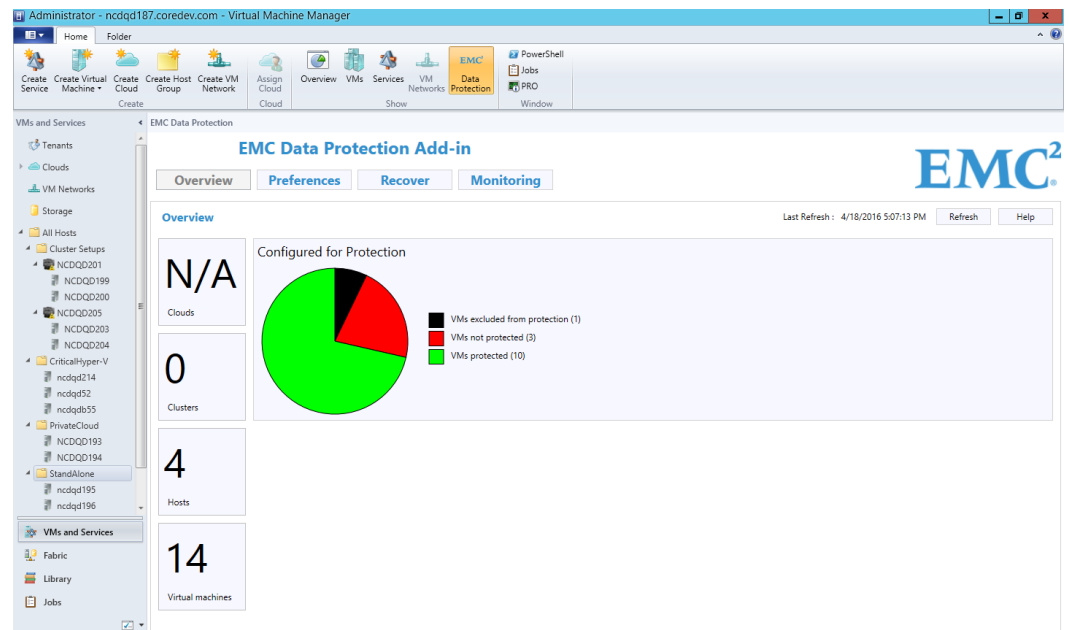
A virtual machine is protected when it is configured for scheduled backups as part of a NetWorker client resource. A virtual machine that is configured for scheduled backups but does not have existing backups is considered protected. Conversely, a virtual machine that is not configured for scheduled backups but has existing backups is not considered protected.

Note

The Data Protection Add-in is unable to distinguish between multiple virtual machines with the same name on the same host. If a host has multiple virtual machines with the same name, and any of these virtual machines are backed up, the Data Protection Add-in shows all the virtual machines as backed up.

The following figure shows the Data Protection Add-in **Overview** page for Administrator, Fabric Administrator, and Read-Only Administrator user roles.


Figure 21 Data Protection Add-in Overview page for Administrator, Fabric Administrator, and Read-Only Administrator user roles



When you position the mouse over a protection category in the pie chart, a tooltip lists the first 10 virtual machines for that protection category. If there are more than 10 virtual machines in that category, the list is truncated with an ellipsis. To view the full list, click the wanted section of the pie chart. If the virtual machine name is more than 40 characters in length, the tooltip truncates the virtual machine name with an ellipsis.

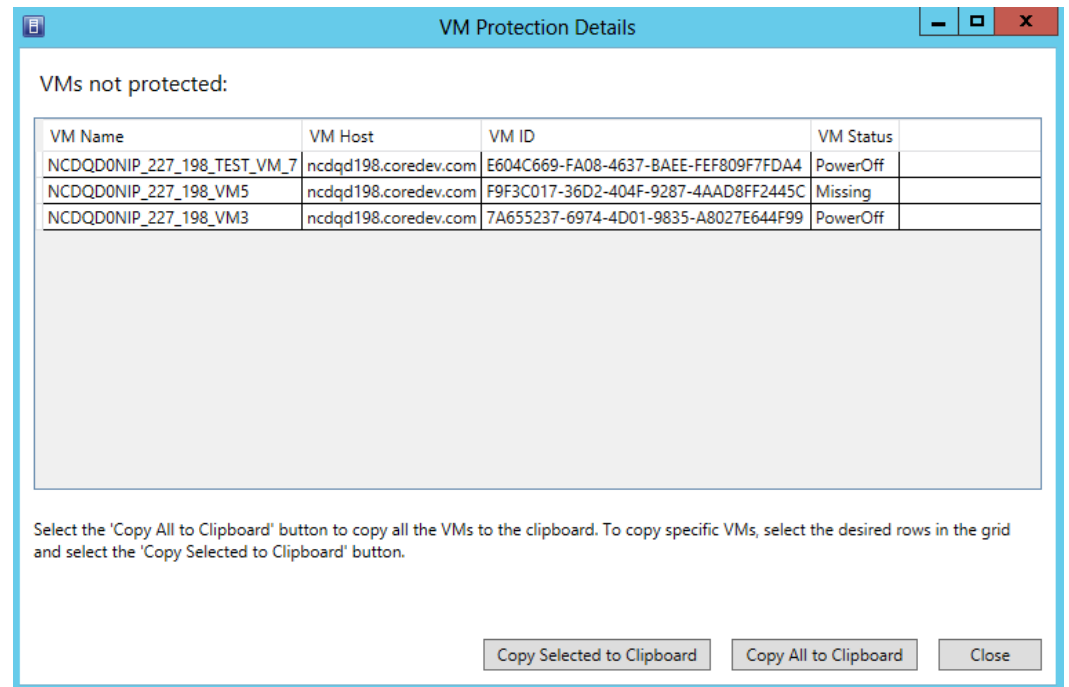
The following figure shows the pie chart and tooltip for Administrator, Fabric Administrator, and Read-Only Administrator user roles.

Figure 22 Virtual machine Protection Details tooltip for Administrator, Fabric Administrator, and Read-Only Administrator user roles



When you click a protection category in the pie chart, the virtual machine **Protection Details** window displays. This window contains a table that lists the name, host, and ID for each virtual machine in the selected protection category. To copy data for all the virtual machines to the clipboard, click the **Copy All to Clipboard** button. To copy data for specific virtual machines, select the wanted rows in the table and click the **Copy Selected to Clipboard** button. You can press Ctrl or Shift to select multiple rows, similar to other Windows applications.

Figure 23 Virtual machine Protection Details window for Administrator, Fabric Administrator, and Read-Only Administrator user roles



Overview page for Tenant Administrator and Application Administrator user roles

For Tenant Administrator and Application Administrator user roles, the **Overview** page displays multiple sub-panes:

- Clouds, Clusters, and Hosts sub-panes—These sub-panes display "NA", since Tenant Administrator and Application Administrator user roles do not have access to other clouds, clusters, or hosts.
- Virtual machines sub-panes—This sub-pane lists the number of virtual machines that the Tenant Administrator and Application Administrator user roles can access.

The pie chart provides the following data about virtual machines:

VMs not backed up

These virtual machines are not currently backed up as part of a NetWorker client resource.

VMs backed up

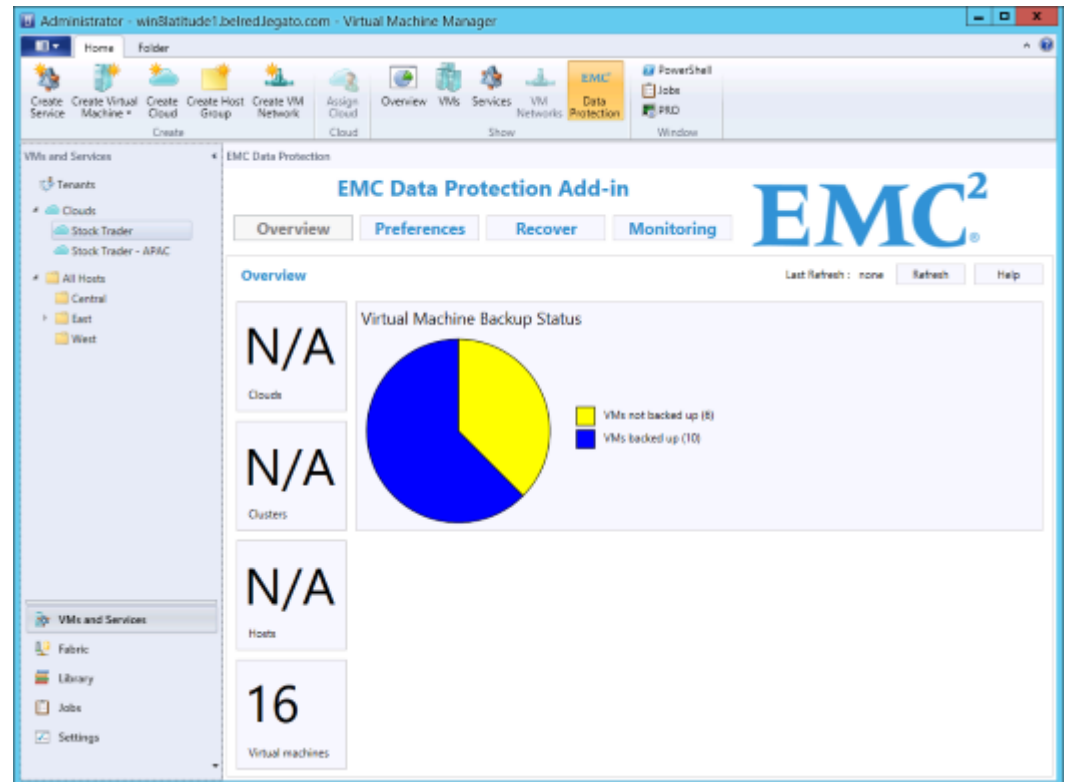
These virtual machines are currently backed up as part of a NetWorker client resource.

Note

The Data Protection Add-in is unable to distinguish between multiple virtual machines with the same name on the same host. If a host has multiple virtual machines with the same name, and any of these virtual machines are backed up, the Data Protection Add-in shows all the virtual machines as backed up.

The following figure shows the Data Protection Add-in **Overview** page for Tenant Administrator and Application Administrator user roles.


Figure 24 Data Protection Add-in Overview page for Tenant Administrator and Application Administrator user roles



When you position the mouse over a backup status category in the pie chart, a tooltip lists the first 10 virtual machines for that backup status category. If there are more than 10 virtual machines in that category, the list is truncated with an ellipsis. To view the full list, click the wanted section of the pie chart. If the virtual machine name is more than 40 characters in length, the tooltip truncates the virtual machine name with an ellipsis.

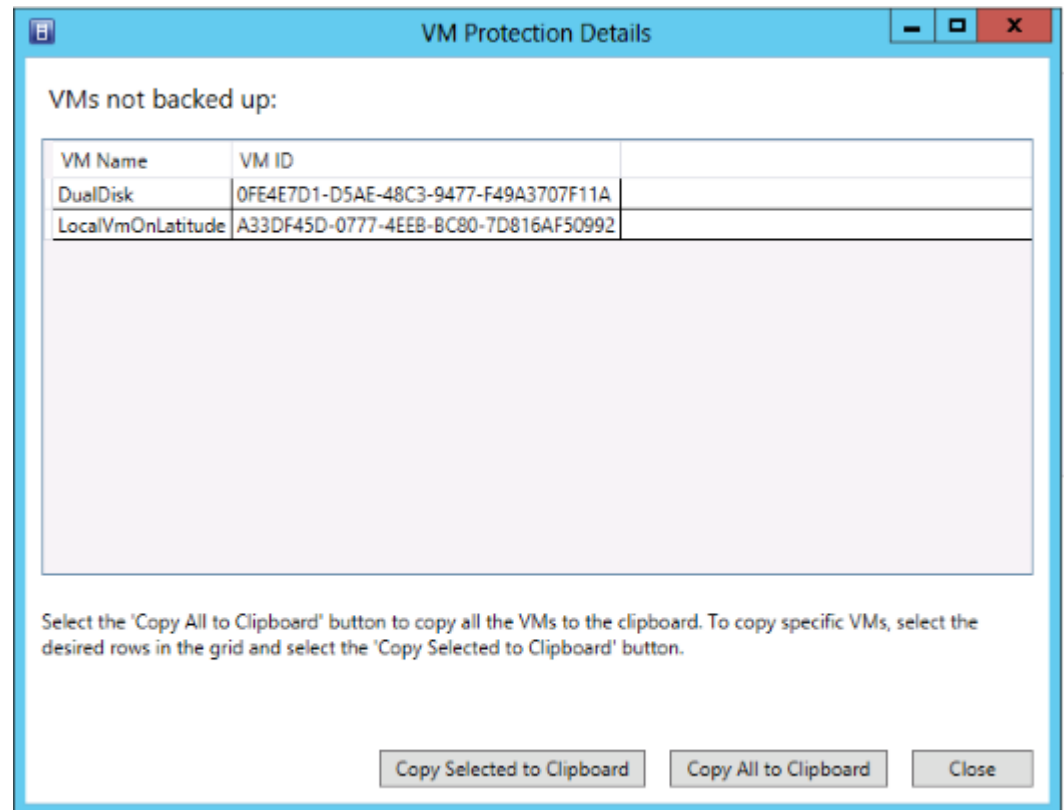
The following figure shows the pie chart and tooltip for Tenant Administrator and Application Administrator user roles.

Figure 25 Virtual Machine Backup Status tooltip for Tenant Administrator and Application Administrator user roles



When you click a backup status category in the pie chart, the virtual machine **Protection Details** window displays. This window contains a table that lists the virtual machine name and virtual machine ID for each virtual machine in the selected backup status category. To copy data for all the virtual machines to the clipboard, click the **Copy All to Clipboard** button. To copy data for specific virtual machines, select the wanted rows in the table and click the **Copy Selected to Clipboard** button. You can press Ctrl or Shift to select multiple rows, similar to other Windows applications.

Figure 26 Virtual machine Protection Details window for Tenant Administrator and Application Administrator user roles



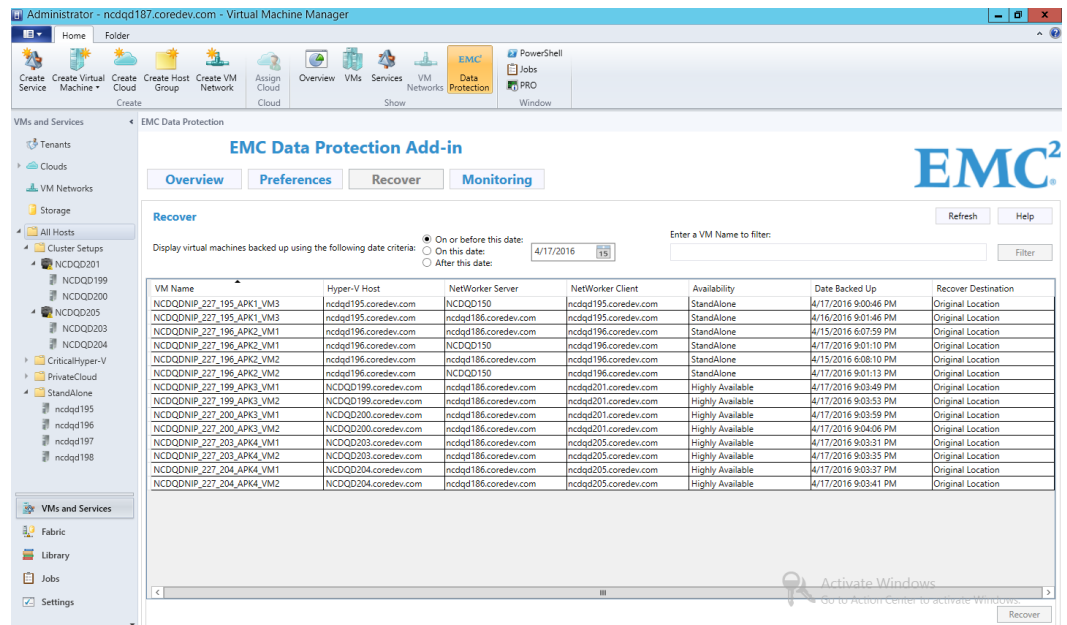
SCVMM Recoveries

When you access the SCVMM **Recover** page for the first time, click **Refresh** to populate the grid with the backups performed on the virtual machines.

Data Protection Add-in for SCVMM Recover page

The backups displayed on the Data Protection Add-in for SCVMM Recover page include those backed up by the NetWorker server according to the **Preferred Servers** list and those matching the data filtering criteria.

If you change the SCVMM environment, such as adding a virtual machine, adding a NetWorker server on the **Preferences** page, or performing a redirected recovery, click **Refresh** to update the list of virtual machines on the **Recover**.

Figure 27 Data Protection Add-in for SCVMM Recover page**NOTICE**

To recover backups that were created using an NMM release earlier than 9.0, click **Start > EMC NetWorker > NetWorker Tools > Restore previous NMM release backups** to start the NetWorker Module for Microsoft GUI. Browse the backups and perform the recovery from the GUI that appears.

Considerations

When performing virtual machine recoveries by using the Data Protection Add-in for SCVMM, consider the following:

- The Data Protection Add-in supports recoveries only from conventional backups. You cannot use the Data Protection Add-in to recover virtual machines from NMM Hyper-V persistent snapshots.
- The Data Protection Add-in is unable to distinguish between multiple VMs with the same name on the same host. If a host has multiple VMs with the same name, the Data Protection Add-in shows incorrect recovery options.
- The Data Protection Add-in does not support recoveries of VMs that have differencing disks.
- The Data Protection Add-in does not perform multiple operations simultaneously, such as recovering multiple VMs or refreshing the list of VMs during a recovery. The **Recover** and **Refresh** buttons are disabled while a recovery or refresh operation is in progress.
- The recovery progress log messages are reported in the following locations:
 - On the **Monitoring** page in the Data Protection Add-in.
 - On the Hyper-V server where the actual recovery is performed.
 - Open the Windows Event Viewer on the server that is hosting the SCVMM console. To access the event logs, browse to **Application and Services Logs > Networker Recovery Service**.
- When recovering to a Windows Server 2016 destination host, the Data Protection Add-in for SCVMM performs import based recovery or proxy based recovery.

- For highly available VMs in cluster configurations, recovery to the original location is always to the active node of the cluster, regardless of the existing virtual machine physical host location. Before starting the recovery, confirm that the cluster active node is the same as the virtual machine physical host. After the recovery is complete, you might need to use Microsoft Cluster Manager to make the virtual machine highly available again.

Note

If this practice is not followed, the resulting conflict of the same virtual machine on different nodes can be very difficult to repair and might require a cluster restart.

- For highly available virtual machine recoveries, when you recover to a cluster physical node rather than to the cluster virtual server, you must use Microsoft Cluster Manager to make the virtual machine highly available after the recovery completes.
- For recoveries of VMs on Hyper-V servers over SMB 3.0 configurations, the Data Protection Add-in supports recovery of stand-alone and clustered configurations.

Note

Because the Data Protection Add-in performs Hyper-V recoveries by using NMM, the NMM Hyper-V considerations that are described in the documentation also apply to performing Hyper-V recoveries by using the Data Protection Add-in.

Viewing available virtual machines

The **Recover** page displays a list of all virtual machines that match the following criteria:

- Reside within the currently selected context of the SCVMM navigation pane
- Have been backed up by a NetWorker server in the **Preferred Servers** list
- Have at least one backup date that matches the current date filter

You can sort the list by **VM Name**, **Hyper-V Host**, or **Availability**.

You can further filter virtual machine names by typing a search string in the **Enter a VM name to filter** field and clicking **Filter**.

Note

The search string can be the full virtual machine name or a sub string of the virtual machine name; The search is case in-sensitive; The search filter is based on the virtual machine name only; Regular expressions are not allowed as part of the search string.

By default, the **Recover** page shows all virtual machines that were backed up on or before the current date. You can filter the virtual machines by selecting one of the date criteria options and choosing a date on the calendar. Only virtual machines with backup times that match the specified date filtering criteria are displayed. If you select criteria that results in no matching backup dates for a particular virtual machine, then that virtual machine does not display in the table.

Recovering virtual machines to the original location

The recovery operation runs on the Hyper-V server that is hosting the virtual machine or, if the virtual machine is highly available, on the active node of the cluster. The **Monitoring** page displays the status of the recovery.

Procedure

1. In the SCVMM console, ensure the **Home** tab is selected.
2. In the **workspaces** pane of the SCVMM console, click **VMs and Services**.
3. In the navigation pane, select the host or cloud that contains the virtual machine you want to recover.
4. On the SCVMM ribbon, click **EMC Data Protection**.
5. In the Data Protection Add-in, click the **Recover** tab.
6. On the **Recover** page of the Data Protection Add-in, select multiple virtual machines at a time in the table and perform recovery.

To include multiple virtual machines in a recovery, select the corresponding save sets by using CTRL + Select and then perform recovery. The selected virtual machines are recovered in a sequential manner.

Note

In NMM releases earlier than 9.2, you can select only a single save set or a virtual machine at a time for recovery by using the Data Protection Add-in for SCVMM.

-
7. Select the **Date Backed Up** cell, click again to activate the drop-down list, and select the backup date and time.
 8. Click the **Recover** button.

Redirected recoveries

The Data Protection Add-in supports redirected recovery of virtual machines to an alternate host to which you have access in the SCVMM console, provided the host is protected with NetWorker Server.

In the SCVMM host, the virtual machine placement path properties contain one or more paths. The redirected recovery location is the first location in this list.

The Data Protection Add-in recovers to the default SCVMM placement path that the Hyper-V administrator configured during the Hyper-V role installation.

The Data Protection Add-in does not support redirected recoveries of Hyper-V backups that were taken before an NMM 8.2 upgrade.

The Data Protection Add-in supports redirected recoveries to a host running the same or later operating system version. For example: The Data Protection Add-in supports redirected recovery from a Windows Server 2016 source host to a Windows Server 2016 destination host, but the Data Protection Add-in does not support redirected recovery from a Windows Server 2016 source host to a Windows Server 2012 R2 destination host.

The Data Protection Add-in does not support virtual machine redirected recovery to an SMB path location. If a virtual machine placement path property specifies a path to an SMB location as the first item in the path list, then a redirected virtual machine recovery to this Hyper-V server is not supported.

Virtual machine IDs after redirected recovery

NMM assigns a new virtual machine ID in certain redirected recovery scenarios. The redirected recovery continues normally, regardless of whether NMM assigns a new ID or uses the existing ID. If NMM assigns a new ID during redirected recovery, then the virtual machine appears in both the source and destination hosts.

The following table provides details about whether NMM assigns the existing virtual machine ID or a new virtual machine ID during a redirected recovery:

Table 25 Virtual machine IDs after redirected recovery

Source operating system	Destination host		Destination virtual machine ID assigned
	Operating system	Configuration type	
Windows Server 2008 R2	Windows Server 2008 R2	Not applicable	Existing
Windows Server 2008 R2	Windows Server 2012 or 2012 R2	CSV	New
Windows Server 2008 R2	Windows Server 2012 or 2012 R2	Stand-alone	Existing
Windows Server 2012	Windows Server 2012 or 2012 R2	CSV	New
Windows Server 2012	Windows Server 2012 or 2012 R2	Stand-alone	Existing
Windows Server 2012 R2	Windows Server 2012 or 2012 R2	CSV	New
Windows Server 2012 R2	Windows Server 2012 or 2012 R2	Stand-alone	Existing

File paths for redirected recovery virtual machines and VHDs

For a redirected recovery, the Data Protection Add-in uses the SCVMM placement path property as the default location for recoveries. The Data Protection Add-in extends the default SCVMM placement path property value by appending the virtual machine name and the recovery time (`vmname_timestamp`) to create a unique subfolder.

If you recover multiple virtual machines with the same name on different source hosts to the same destination host, the Data Protection Add-in recovers these virtual machines to two different folders with unique subfolders by appending `vmname_timestamp` to the folder names. For example, if two virtual machines that are both named `Virtual_Machine` are recovered to the default SCVMM placement path property "C:\ProgramData\Microsoft\Windows\Hyper-V", the virtual machines are recovered to the following unique subfolders:

- C:\ProgramData\Microsoft\Windows\Hyper-V
 \Virtual_Machine_20140917143500\
- C:\ProgramData\Microsoft\Windows\Hyper-V
 \Virtual_Machine_20140917152205\

If the virtual machine has multiple disks with the same name, the Data Protection Add-in recovers these disks to separate folders. For example, if a virtual machine with two VHDs that are both named `DualDisk.vhd` are recovered to the default SCVMM placement path property "C:\ProgramData\Microsoft\Windows\Hyper-V", the virtual machines are recovered to the following unique subfolders:

- C:\ProgramData\Microsoft\Windows\Hyper-V
 DualDisk_20140625133500\1\DualDisk.vhd
- C:\ProgramData\Microsoft\Windows\Hyper-V
 DualDisk_20140625133500\2\DualDisk.vhd

Note

Microsoft limits virtual machine file paths to 260 characters. If the appended file path exceeds 260 characters, the recovery fails.

Performing a redirected recovery

Perform the following steps for redirected recovery of NMM 18.2. You cannot perform a redirected recovery of NMM 8.2 and later save sets to a NMM 18.2 Hyper-V Server and vice-versa.

Procedure

1. Take the original virtual machine offline to avoid conflicts during the recovery operation.
2. In the SCVMM console, ensure the **Home** tab is selected.
3. In the **workspaces** pane of the SCVMM console, click **VMs and Services**.
4. In the left navigation pane, select the host or cloud that contains the virtual machine you want to recover.
5. On the SCVMM ribbon, click **EMC Data Protection**.
6. In the Data Protection Add-in, click the **Recover** tab.
7. On the **Recover** page of the Data Protection Add-in, select multiple virtual machines at a time in the table and perform recovery.

To include multiple virtual machines in a recovery, select the corresponding save sets by using CTRL + Select and then perform recovery. The selected virtual machines are recovered in a sequential manner.

Note

In NMM releases earlier than 9.2, you can select only a single save set or a virtual machine at a time for recovery by using the Data Protection Add-in for SCVMM.

8. Select the **Date Backed Up** cell, click again to activate the drop-down list, and select the relevant backup date and time.
9. Select the **Recover Destination** cell, click again to activate the drop-down list, and select the destination host.

The **Recover Destination** drop-down list shows physical Hyper-V hosts that are NetWorker clients and that are visible in SCVMM for the current user. The **Recover Destination** drop-down list does not list the NetWorker virtual server clients representing the clusters.

10. Click the **Recover** button.

11. If an **Action Needed** message displays, click **OK** to clear the message.
12. Confirm that the virtual machine is successfully recovered by verifying that the virtual machine appears in the hypervisor on the Hyper-V server where you recovered the virtual machine.
13. If the **Action Needed** message displayed, delete the original virtual machine from its original host by using the SCVMM console. Alternatively, delete the virtual machine by using Hyper-V Manager or PowerShell, and then refresh the SCVMM console.
14. In the navigation pane of the SCVMM console, right-click the destination host and click **Refresh Virtual Machines**.
15. If a new ID was assigned to the virtual machine as described in [Virtual machine IDs after redirected recovery](#) on page 124, in the navigation pane of the SCVMM console, right-click the source host and click **Refresh Virtual Machines**.
16. Ask the NetWorker administrator to perform a backup of the virtual machine from its new Hyper-V host.

Viewing virtual machines after a redirected recovery

If you perform a redirected recovery of a virtual machine to a Hyper-V host, then the virtual machine will not meet the criteria that are listed in [Viewing available virtual machines](#) on page 122 until after a new backup of the virtual machine is completed. Therefore, the **Recover** page does not immediately display the redirected virtual machines. After you perform a redirected recovery for a virtual machine, ask the NetWorker administrator to perform a backup of the Hyper-V host where the virtual machine currently resides.

Because the Data Protection Add-in displays only backups for the current Hyper-V host of the virtual machine, if you want to recover a virtual machine from a backup that was taken before a redirected recovery, you must use NMM.

Recovering a deleted virtual machine

The Data Protection Add-in does not support recovering virtual machines that have been deleted from SCVMM. The NetWorker administrator must perform the recovery by using NMM.

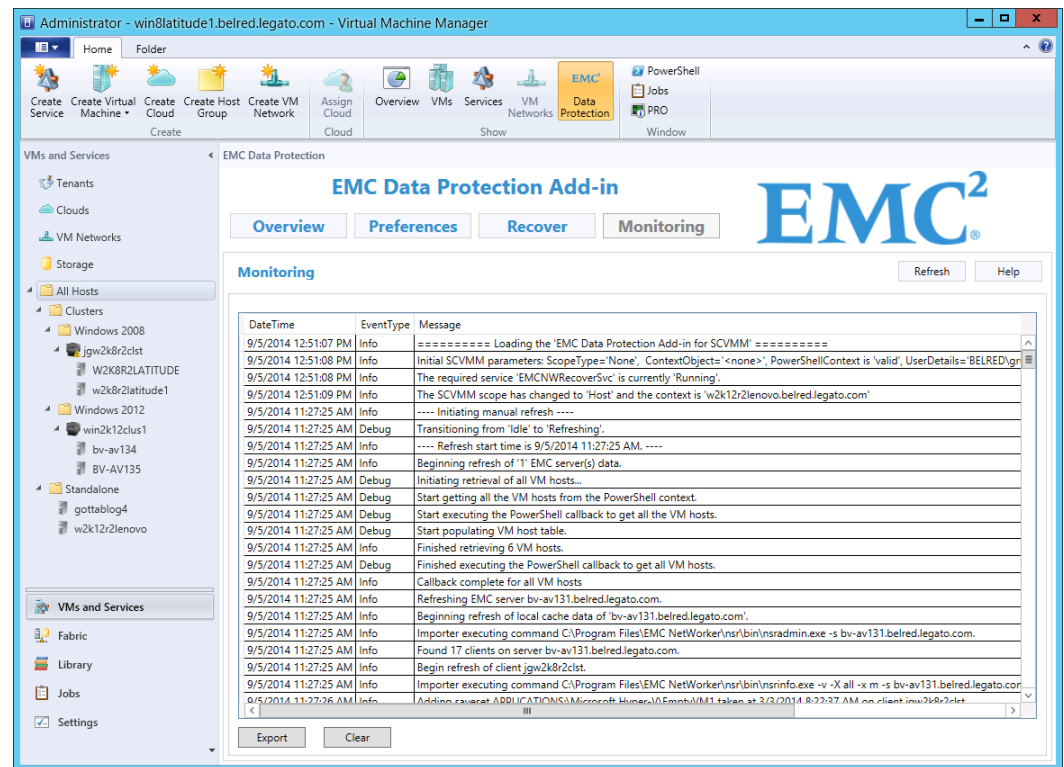
Monitoring

The **Monitoring** page provides information about Data Protection Add-in events and operations.

The **Monitoring** page displays:

- Status of recovery operations in progress
- Details of queries to the NetWorker servers and the SCVMM server
- All logging entries from previous uses of the Data Protection Add-in (if any)

Figure 28 Data Protection Add-in for SCVMM Monitoring page



The **Monitoring** page shows three columns, all of which can be sorted:

- **DateTime**
- **EventType**
- **Message**

The monitor log information is updated in real-time as operations occur. To manually scan for updated protection information, click **Refresh**.

You can export the log file by clicking **Export** at the bottom of the **Monitoring** page. NMM logs are stored on the destination host, where the virtual machine is restored. The exported log file name is **MonitorExportFile** and is at **C:\Users\<current user>\AppData\Local\EMC\NetWorker\SCVMM**.

Troubleshooting

This section includes information about how to resolve general issues you might encounter while using the Data Protection Add-in. The *NetWorker Administration Guide* and the *NetWorker Module for Microsoft Administration Guide* provide additional troubleshooting details.

Recovered virtual machine does not start

If a recovered virtual machine does not start, perform the following steps:

Procedure

1. Select the recovered virtual machine, then right-click the virtual machine and select **Discard Saved State**.
2. Right-click the recovered virtual machine and then select **Properties**.

3. In the **Properties** dialog box, click **Hardware Configuration** and verify the Network Adapter settings of the virtual machine.

Installation fails due to access issue

When you install the Data Protection Add-in, you need access to the following path:

C:\Users\Public\Documents\EMC NetWorker\nsr\addins
\VMM_DataProtection

Note

This path applies to environments in which the system drive is C:.

Solution

Before you install the Data Protection Add-in, verify that you have read/write access permissions to the paths previously noted.

The Data Protection Add-in for SCVMM displays an incorrect NetWorker Server version

If the NetWorker Server software is updated, the Data Protection Add-in for SCVMM displays the previous version number until you remove and re-add that NetWorker server in the Add-in.

Importing fails due to access issue

To import the Data Protection Add-in, you need access to the following paths:

- C:\Program Files\Microsoft System Center 2016\Virtual Machine Manager\bin\AddInPipeline\AddInViews
- C:\Program Files\Microsoft System Center 2016\Virtual Machine Manager\bin\AddInPipeline\AddIns

Note

These paths apply to environments in which SCVMM was installed in the default location of C:\Program Files:

If you do not have access to the required paths, you receive the following error:

The assembly

Microsoft.SystemCenter.VirtualMachineManager.UIAddIns.dll
referenced by the add-in assembly EMC.BRS.ScvmAddIn.AddInView could not be found in the add-in package. Ensure that this assembly was included with the add-in package.

Solution

Before you import the Data Protection Add-in, verify that you have read/write access permissions to the previously noted paths.

Virtual machine attributes might display incorrect values

On the **Monitoring** page of the Data Protection Add-in, the **VM Availability** attribute might occasionally show an incorrect value.

To show the correct information:

1. In the SCVMM navigation pane, refresh the virtual machine.
2. In the Data Protection Add-in, click **Refresh**.

Redirected recovery appears to succeed but no virtual machine appears in Hyper-V Manager

If a redirected recovery appears to succeed but no virtual machine appears in Hyper-V Manager, the network of the target host might be incompatible. For example, if the target host is in a non-trusted domain, redirected recovery to this target host fails.

If the network of the target host is incompatible, then the virtual machine is disconnected from the network. The recovery succeeds according to the Data Protection Add-in monitor log and the `nsrnmmsv.raw` log, and the virtual machine files are stored on the target host and volume, but Hyper-V Manager does not display or recognize the virtual machine.

Solution

Reconnect the existing switch of the host by using the SCVMM GUI or by using the following PowerShell command:

```
$sw=Get-virtual machineswitch;get-vm -Id <vmID> |  
Get-VMNetworkAdapter | Connect-VMNetworkAdapter -SwitchName  
$sw.Name
```

After reconnecting the existing switch, re-attempt the redirected recovery.

Checks for redirected recovery failures

Redirected recovery of a virtual machine might fail due to virtual machine network or saved state incompatibility between the original Hyper-V host and the target Hyper-V host. The Hyper-V writer cannot register the virtual machine because of errors in virtual machine configuration files which the writer cannot resolve.

If you suspect this is the problem for a failed redirected recovery, then examine the target host destination location for the virtual machine files. Look in the Monitor logs for the `redirected restore cmd line options`: output. If the virtual machine files are there, then try to register the virtual machine manually by using the SCVMM UI.

Avoid virtual machine names with the same name within an SCVMM context

The Data Protection Add-in primarily uses the virtual machine name, as displayed in Hyper-V Manager or Failover Cluster Manager, as an identifier for the virtual machines. If multiple virtual machines have the same name in the same SCVMM context, then the Add-in is unable to distinguish between the virtual machines. Although not required, it is considered best practice for virtual machine names to be unique.

Cluster virtual machine backups do not display on the Recover page

If a cluster virtual machine backup does not display on the **Recover** page, check that the cluster is configured as highly available in Microsoft Cluster Manager.

If a virtual machine is removed from Microsoft Cluster Manager and is no longer shown by PowerShell as highly available, the backups for that virtual machine do not display on the **Recover** page.

Redirected recovery is not supported when the virtual machine name or virtual machine configuration path contains special characters

NMM Hyper-V restricts the use of special characters in virtual machine names and virtual machine configuration paths.

When you try to recover a Hyper-V save set, virtual machine name, or file path that contains a character that is not listed above, the Data Protection Add-in checks the name and path of the virtual machine objects and displays a warning message stating that the virtual machine contains unsupported characters and cannot be recovered to an alternate location.

You can use SCVMM to perform the following workaround:

Procedure

1. Recover the virtual machine to the original location of the backup.
2. Use SCVMM to export the virtual machine to a temporary location.
3. Copy the virtual machine files to an appropriate location on the target host.
4. Use SCVMM to import the virtual machine.

CHAPTER 7

Windows Bare Metal Recovery Solution

This chapter includes the following sections:

- [Microsoft Hyper-V Backup and BMR.....](#) 132
- [Microsoft System Center Virtual Machine Manager backup and BMR.....](#) 133

Microsoft Hyper-V Backup and BMR

Ensure that the Hyper-V Server is backed up before performing BMR.

Backing up Hyper-V for BMR

Configure a Hyper-V client resource to back up Hyper-V for BMR.

Review the "Configuring a client resource by using the Client Backup Configuration Wizard" and "Configuring a client resource manually by using the NetWorker Management Console" sections in the Backups chapter for detailed information about how to create a client resource.

Note

Separately create a client resource to back up the application data and the file system data which includes the `DISASTER_RECOVERY : \ save set`.

Performing BMR of Hyper-V

Ensure that the cluster service account is enabled. Otherwise, the explorer stops responding when you browse CSVs and the recovery fails.

To enable the cluster service computer account, log in to the Domain Controller with the domain administrator credentials, and use the Active Directory Users and Computer snap-in.

Procedure

1. Perform the procedures that the "Performing a Windows BMR recovery to physical or virtual computers" section in the *NetWorker Administration Guide* describes.

After the recovery completes, the status of Hyper-V VMs in the Failover Cluster Manager window are either Pending or Failed.
2. Reconfigure quorum and cluster storage (CSV and non-CSV volumes):
 - a. Create or configure the quorum, CSV, and non-CSV shared volumes.
 - b. Ensure that you have documented the drive letters for quorum and non-CSV (shared volume) and used here.
 - c. Ensure that you have configured the same CSV volume name as the original.
3. Recover all Hyper-V cluster nodes by using the NMM client.
4. Reconfigure the VMs with either the **Pending** or the **Failed** status.
5. Recover the Hyper-V VMs that the CSV volumes host with the virtual node as the client resource, by using the NMM client.
6. Recover the Hyper-V VMs that the non-CSV volumes host with the physical node as the client resource, by using the NMM client.

Microsoft System Center Virtual Machine Manager backup and BMR

Ensure that the System Center Virtual Machine Manager is backed up before performing BMR.

Backing up System Center Virtual Machine Manager for BMR

Back up and recover the following key components of VMM to ensure a protected VMM environment:

- **VMM server:** The VMM server is the central component of a VMM deployment. The server contains the core Windows service that includes the VMM engine.
- **VMM database:** The VMM database resides on either the VMM server or on a remote database server. You must regularly save the VMM database to quickly recovery the VMM environment.
- **VMM library:** The VMM library is either one file server or multiple file servers with file shares that index specific file types that the VMM uses. The VMM library shares include VHD, VMDK, ISO, PS1, INF, VFD, FLP, and XML files besides stored virtual machines.

You must frequently back up the data such as, the data on the virtual machines, which VMM manages, that often changes. NMM supports the Microsoft Hyper-V VSS writer to protect Hyper-V virtual machines.

You must perform BMR backups for the Hyper-V servers also.

If you have installed any component of VMM on a Hyper-V virtual machine, you can protect it by using the NMM client's Hyper-V. Use NMM to protect the VMM server, the VMM library, or the VMM database on a virtual machine. NMM takes image backups of Hyper-V virtual machines. You do not need to install an NMM client on the virtual machine. Instead, install an NMM client on the Hyper-V server and configure NMM client resources to back up the virtual machines. Use the NMM GUI Hyper-V plug-in to recover a virtual machine that contains a VMM component.

Backing up the VMM library for BMR

The VMM library is a catalog of resources that enables you to store objects that are not running or associated with a host. The library contains the files that are present on library shares, templates, operating system profiles, and hardware profiles.

Procedure

- Use the NMM backup and recovery techniques for Windows files servers to back up and recover a VMM library server.
- Use the NetWorker client to back up the VMM library by specifying the `ALL` save set to perform a Windows BMR backup of the VMM library. The `ALL` save set enables you to back up the local file systems and the `DISASTER_RECOVERY:\` save set of the VMM library.

[Creating a client resource for a VSS-based backup by using the Client Backup Configuration wizard](#) on page 47 and [Manually creating a client resource for a VSS-based backup by using the Client Properties dialog box](#) on page 52 provide information about how to create client resources to back up the application data and the file system data respectively.

Backing up the VMM server for BMR

The VMM server relies on the VMM database to store data. The VMM database uses the server account to encrypt that data.

You must perform a full backup of the VMM server host to ensure that you meet the following requirements:

- The same host account is available to interact with the VMM database after recovery.
- All virtual machine hosts associates and communicates with the VMM server.
- You can recover any communication certificates so that the existing agents continue to work as you expected.

You can protect the VMM server by performing a NetWorker Windows BMR backup along with an NMM backup of local file systems and application-specific data, such as the data associated with Hyper-V Server, SQL Server, or SharePoint Server. The advantage of a Windows BMR backup and recovery capability for a VMM server is that the backup includes the system state. When you perform recovery, the system state recovers the computer security identifier (SID) that uniquely identifies the computer to the network. The computer-recovered SID enables the VMM server to quickly and easily interact with any existing VMM database, VMM library, and managed servers without additional post-recovery operations. The Windows BMR backup captures the VMM authorization datafile which is typically `%ProgramData%\Microsoft\Virtual Machine Manager\HyperVAuthStore.xml` when protecting the critical system drive.

Procedure

1. Use the NetWorker client to back up the VMM server non-application files by specifying the `ALL` save set. This setting enables the backup of the local file systems and the `DISASTER_RECOVERY:\` save set of the VMM server.

[Creating a client resource for a VSS-based backup by using the Client Backup Configuration wizard](#) on page 47 and [Manually creating a client resource for a VSS-based backup by using the Client Properties dialog box](#) on page 52 provide information about how to create client resources to back up the application data and the file system data respectively.

2. Use an NMM client resource specification to perform a full backup of any application data on the VMM server.

Backing up the VMM database for BMR

The VMM database is a SQL Server database that contains all the configuration information for VMM. Give the highest priority to protect the VMM database.

The default installation of VMM uses Windows Internal Database (WID), a version of Microsoft SQL Express, on the same host as the VMM server. The database name is `MICROSOFTVMM\Virtual Manager DB`. However, the VMM database can also reside on either an existing Microsoft SQL Server Standard or SQL Server Enterprise server. To find out the location of the database for VMM, inspect the Registry values entries under the `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft System Center Virtual Machine Manager Server\Settings\Sql` key.

NMM clients support the backup and restore of SQL databases.

Back up the VMM database that is part of a SQL Server installation.

The *NetWorker Module for Microsoft for SQL VDI User Guide* provides information about how to protect a SQL server.

Procedure

1. Configure an NMM SQL Server client resource for the system with the VMM database.
2. Include the VMM databases when you schedule a backup.
3. In the **Client Properties** dialog box, in the **Save set** field, specify the VMM database. You can save either the entire SQL Server or named instances and individual databases.

Performing BMR of a System Center Virtual Machine Manager

The VMM components have interdependencies. Recover the entire VMM environment for a disaster recovery.

Procedure

1. Recover the library file server.
 - a. Recover the physical VMM library file server with a Windows BMR save set.
 - b. Use the NetWorker client file system plug-in to recover the file system that you shared as the VMM library.
2. Recover the VMM database.
 - a. Recover the physical SQL server with a Windows BMR save set.
 - b. Stop the VMM services.
 - c. Use the NMM GUI SQL plug-in to restore the VMM database.
 - d. Start the VMM services.
3. Recover the VMM host.
4. After the operating system and VMM server are stable, perform the required post-recovery VMM steps that the Microsoft documentation describes.

Performing BMR of the VMM server

Perform the procedures described in the "Performing a Windows BMR recovery to physical or virtual computers" section in the *NetWorker Administration Guide*.

Performing BMR of the VMM library

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. After you recover the file share directories, either adjust the VMM Library Refresher interval or manually trigger the Library Refresher to enable VMM to refresh the contents of the file server share.

The VMM Administrator Console embedded help describes how to set the Library Refresher interval and perform a manual refresh.

Performing BMR of the VMM database

Perform the procedures that the "Performing a Windows BMR recovery to physical or virtual computers" section in the *NetWorker Administration Guide* describes.

If you only back up and recover the VMM database, then you must use the same instance of the VMM server component for the restore. This is important to reduce VMM post -restore steps. Microsoft VMM documentation provides more information.

You must manually stop the VMM Windows service before you restore the VMM database. After you restore the database, start the service. After you restore the VMM database and start the VMM Windows service, VMM synchronizes the database with the VMM server and VMM library computers.

Note

You must reapply any changes that you made to the VMM environment between the time of the backup and the restore. For example, you must re-create the templates that you created in VMM after performing the backup that you use to perform the restore.

Before you restore a SQL database, verify whether all the SQL services have started. To perform a SQL Server recovery as part of BMR, perform the steps that the Bare Metal Recovery chapter in the *NetWorker Module for Microsoft for SQL VDI User Guide* describes.

CHAPTER 8

Troubleshooting

This chapter includes the following sections:

- [Troubleshooting generic issues](#)..... 138
- [Troubleshooting backups issues](#)..... 138
- [Troubleshooting recovery issues](#)..... 141

Troubleshooting generic issues

The following topics explain generic issues that might occur in a Hyper-V environment, and provide solution or workaround.

Sometimes Hyper-V virtual machine does not start, and an error message appears

Problem

Sometimes Hyper-V virtual machine does not start, and one of the following similar error messages appears:

- The application encountered an error while attempting to change the state of <virtual_machine_name>
- General access denied error' (0x80070005)

Solution

<https://support.microsoft.com/en-gb/help/2249906/hyper-v-virtual-machine-may-not-start--and-you-receive-a-general-access-denied-error> provides information about the issue and its resolution.

Troubleshooting backups issues

The following topics explain issues that might occur during the backup process for a Hyper-V environment, and provide solution or work around.

Redirected I/O status does not update after CSV backup

Problem

During a CSV backup, the CSV is in redirected I/O status. Other nodes cannot directly write to disks. Instead, the I/O is redirected over the LAN to the owner node that is performing the backup.

Solution

If the redirected I/O status does not update correctly after the CSV backup is complete, clear the status by performing one of the following procedures:

1. Type the following commands at the Windows PowerShell command prompt to delete the stale shadows:

```
diskshadow
DISKSHADOW> list shadows all
DISKSHADOW> delete shadows all
```

Test-ClusterResourceFailure "volume name"

Note

This command might clear the "backup in progress" status.

2. If the "redirected access" status is not cleared after performing step 1, change the coordinator node by moving the volume to another node in the cluster and verifying that the volume is online.

Type the following command at the Windows PowerShell command prompt to clear the backup state for the affected volume:

```
nsr_csvutil -c <csv_volume_path>
```

For example: `nsr_csvutil -c "c:\ClusterStorage\Volume1"`

Hyper-V pass-through disks are not backed up in a child partition backup

Problem

For Hyper-V backups, the child partition pass-through disks are skipped in the Hyper-V parent partition backup, and child partition pass-through disks are supported by backups within the child partition. However, sometimes the Hyper-V parent partition backup of a child partition with a pass-through disk might fail completely.

Solution

If this failure occurs, contact Microsoft support for assistance because the problem might be with the hardware configuration or the Microsoft Hyper-V writer.

Hyper-V configuration requirements for backing up a virtual machine contains multiple volumes

Problem

When the guest contains multiple virtual hard disks, the backup of the associated virtual machine from the Hyper-V server might fail because of a Microsoft limitation. When multiple volumes exist on the guest, VSS determines the shadowstorage area for the snapshots, based on the volume that has more space. This can cause snapshots of two volumes to reside on the volume that has more space. For example, the snapshots of volume C and volume D may reside on volume D because VSS has determined that volume D contains more space. During the snapshot revert stage, PostSnapshot, the snapshot of volume C may be lost if the snapshot of volume D is reverted first.

Solution

To prepare a multiple volume guest for backup:

1. Use the `vssadmin` command to force the shadowstorage of each volume to occur on the same volume. Run the following commands from inside each guest, not the parent physical Hyper-V Server.

```
vssadmin Add ShadowStorage /For=C: /On=C:
vssadmin Add ShadowStorage /For=D: /On=D:
```

2. Repeat as needed for each volume in the virtual machine.

Hyper-V CSV virtual machine backups fail when the VHDX is stored in the root of the CSV volume

Problem

Due to a Microsoft limitation, Hyper-V CSV virtual machine backups fail when the VHDX is stored in the root of the CSV volume. For example: C:

```
\ClusterStorage\Volume1\Test Lab.Vhdx
```

Solution

Move the virtual machine to a subfolder. For example: C:\ClusterStorage
 \Volume1\Test Lab\Test Lab.Vhdx.

Backup of an online virtual machine might reset the uptime value of the virtual machine

Problem

After an Application-Aware backup of an online virtual machine, Hyper-V Manager might display an inaccurate system uptime. System uptime inside the guest operating system is accurate.

Solution

Contact Microsoft for help with this issue.

Hyper-V backups are backed up in a crash-consistent state

Problem

When the guest contains multiple virtual hard disks, the backup of the associated virtual machine from the Hyper-V server might fail because of a Microsoft limitation. When multiple volumes exist on the guest, VSS determines the shadowstorage area for the snapshots, based on the volume that has more space. This can cause snapshots of two volumes to reside on the volume that has more space. For example, the snapshots of volume C and volume D may reside on volume D because VSS has determined that volume D contains more space. During the snapshot revert stage, PostSnapshot, the snapshot of volume C may be lost if the snapshot of volume D is reverted first.

Solution

To prepare a multiple volume guest for backup:

1. Use the `vssadmin` command to force the shadowstorage of each volume to occur on the same volume. Run the following commands from inside each guest, not the parent physical Hyper-V Server.

```
vssadmin Add ShadowStorage /For=C: /On=C:
```

```
vssadmin Add ShadowStorage /For=D: /On=D:
```

2. Repeat as needed for each volume in the virtual machine.

Unable to create a checkpoint of a virtual machine for an RCT-based backup

Problem

Sometimes you cannot create a check point of a virtual machine when you perform an RCT-backup. This is because one of the previous backups of the virtual machine would have failed, and checkpoints would remain with the virtual machine.

Solution

1. Review the NMM and VMMS logs to find the reason for the failure of the checkpoint creation.
2. Create a user checkpoint by using the Hyper-V Server Manager.
3. Merge the stale checkpoints of the virtual machine by running the following PowerShell command:

```
Get-VMSnapshot -VMName <virtual_machine_name> -ComputerName
<Hyper-V_Server_name> | Remove-VMSnapshot
```

[Removing and merging stale recovery checkpoints](#) on page 44 provides more information about how to remove and merge stale recovery checkpoints.

Backup of Hyper-V clustered server virtual machines fail

Problem

In a Hyper-V clustered server environment, a backup fails to start on one or multiple remote Hyper-V Server nodes if either the remote node is not reachable or NetWorker and NMM are not installed on the remote node.

The following similar message appears in the log file:

```
nsrrpcinfo: Remote system error No connection could be made
because the target machine actively refused it. Removed
node1.domain.com from PSOL as the server is not accessible.
```

Solution

1. Validate the node on the cluster.
2. Ensure that NetWorker and NMM are installed on the node.
3. Ensure that all the NetWorker and NMM services are started on the node.

Troubleshooting recovery issues

The following topics explain issues that might occur while performing a Hyper-V recovery, and provide solution or work around.

Failure to establish a Client Direct session during GLR

The Client Direct feature must be enabled to perform GLR. To verify that the environment has Client Direct enabled, perform 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 *EMC 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 the client has name resolution for the Data Domain device. If an AFTD storage node is being used, ensure the client has name resolution for the storage node.
3. Run the following `save` command from the command prompt:

```
PS C:\Program Files\EMC NetWorker\nsr\bin> save -D1 -a
DIRECT_ACCESS=yes -b networker_pool 'C:\Windows\System32\drivers
\etc\hosts'
```

Where *networker_pool* is the NetWorker pool containing the volumes where the save sets for recovery reside.
4. Check the output for messages indicating the Client Direct session is established:

```
10/16/16 23:59:27.094472 Default DFA handling by client is
'Fallback'
10/16/16 23:59:27.094472 DIRECT_ACCESS=yes: Client direct set to
'Yes'
```

```
10/16/16 23:59:27.129477 Device attribute block size is 262144
10/16/16 23:59:29.185589 libDDBoost version: major: 3, minor: 3,
patch: 0, engineering: 2, build: 545054
10/16/16 23:59:29.197590 load ddp_get_file_segment_type
129292:save: Successfully established Client direct save session
for save-set ID '889485007' (mb-vm-sql-2.dpsg-sea.emc.c
om:C:\Windows\System32\drivers\etc\hosts) with Data Domain volume
'ddveselssemccom.002'.
10/16/16 23:59:29.299596 using DFA save for ssid = 889485007
10/16/16 23:59:29.299596 ssid 889485007 using DFA save to 'mb-vm-
nw-2'
10/16/16 23:59:29.299596 Successfully setup direct saves
```

5. (Optional) If the save command fails:

a. Run the save command again after replacing -D1 to -D3:

```
PS C:\Program Files\EMC NetWorker\nsr\bin> save -D3 -a
DIRECT_ACCESS=yes -b networker_pool 'C:\Windows\System32\drivers
\etc\hosts'
```

Where *networker_pool* is the NetWorker pool containing the volumes where the savesets for recovery reside.

b. Check for output messages indicating the Client Direct session is established.

c. If a Client Direct session is not established, find the messages indicating the cause of the failure, and fix the problem as required.

Cannot enable a Client Direct session and GLR failing as a result

Client Direct is required for Hyper-V GLR, including GLR of backups taken with a previous version of NMM. 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. 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.

When recovering multiple Hyper-V CSV virtual machine through proxy, all the virtual machines are recovered but all the virtual machines are not registered

Problem

In a Hyper-V CSV setup, when you recover multiple Hyper-V CSV virtual machines through proxy, all the virtual machines are recovered although only one virtual machine is registered.

Solution

After recovery of multiple Hyper-V CSV virtual machines through proxy is complete, NMM recovers the .vhd and .xml files. Manually run the following PowerShell command to register the virtual machines that are not registered: `ps C:\Users\administrator.CONTOSO> Import-VM -path "C:\ClusterStorage\Volume3\CSV-VM-013\CSV-VM-013\Virtual Machines\E45E8DBB-FAEF-4A79-B891-5386AB20F66B.xml"`

Name	State	CPUUsage(%)	MemoryAssigned(M)	Uptime	Status
CSV-VM-013	Off	0 0	00:00:00	Operating	normally

After Hyper-V CSV disaster recovery, application data recovery fails and CSV mount point is not browsable

Problem

After disaster recovery, if NMM is used to recover Hyper-V data, the following issues are observed:

- Hyper-V recovery of virtual machines that are located in a shared disk (but non-CSV volume) fails.
- The CSV volumes are not browsable, and recovery of virtual machines that are on the CSV volume fails.

Solution

Perform the following steps:

1. Remove stale entries from the cluster resource.
2. In the Domain Controller, start the Active Directory Users and Computers Snap-In, and ensure that the failover cluster virtual network name account of Hyper-V Virtual Server is enabled.

Through Advanced Recovery option, recovery of online virtual machine to other node in same cluster setup completes

Problem

In a Hyper-V CSV setup, when a child partition is running, the same child partition can be recovered to another node using the Advanced Recovery option. This creates multiple virtual machines in different CSV nodes.

Solution

If the virtual machine is online or active, recover the virtual machine to the same node.

NMM registers corrupted Hyper-V child partition to Hyper-V Server

Problem

Although a recovery operation for a Hyper-V child partition fails, NMM registers the corrupted Hyper-V child partition to the Hyper-V Server.

Solution

After receiving a confirmation about a failed recovery operation, the Hyper-V system administrator must delete the following:

1. The corrupted Hyper-V child by using the Hyper-V Manager.
2. The corresponding child partition .vhd files.

Sometimes recovery to a particular node fails, but the same recovery to an alternate node succeeds

Problem

Sometimes recovery to a particular node fails. In the case of a failure, data is recovered, but import of virtual machines fails because of presence of stale

entries. So, the logs report the entire recovery as failed. However, the same recovery to an alternate node succeeds.

Solution

On the node, to which the recovery failed, delete stale entries from the following folders:

- C:\ProgramData\Microsoft\Windows\Hyper-V\Planned Virtual Machines
- C:\ProgramData\Microsoft\Windows\Hyper-V\Planned Virtual Machines Cache

APPENDIX A

Recovering SQL Server, Exchange Server, and SharePoint Server Items from a Hyper-V Virtual Machine

This appendix includes the following sections:

- [Overview](#) 146
- [Recovering items that are stored on a Hyper-V virtual machine](#) 146

Overview

This appendix describes how to perform granular-level recovery (GLR) of Microsoft SQL Server, Exchange Server, and SharePoint Server items that are stored in Hyper-V virtual machines.

The Hyper-V Writer of Microsoft Hyper-V Server supports only full backups (VSS_BT_FULL). The Hyper-V requestor performs a full backup of virtual machines that run a Microsoft application (SQL Server, Exchange Server, or SharePoint Server). If a requestor specifies VSS_BT_COPY, the Hyper-V Writer performs a full backup, according to the VSS MSDN documentation.

The following table shows the backup types that the requestor can set by using the SetBackupState on the host and the backup type set by the Hyper-V requestor inside the guest.

Table 26 Backup types

Backup type set by requestor through SetBackupState on the host	Backup type set by Hyper-V's requestor inside the guest
VSS_BT_FULL	VSS_BT_FULL
VSS_BT_COPY	VSS_BT_FULL

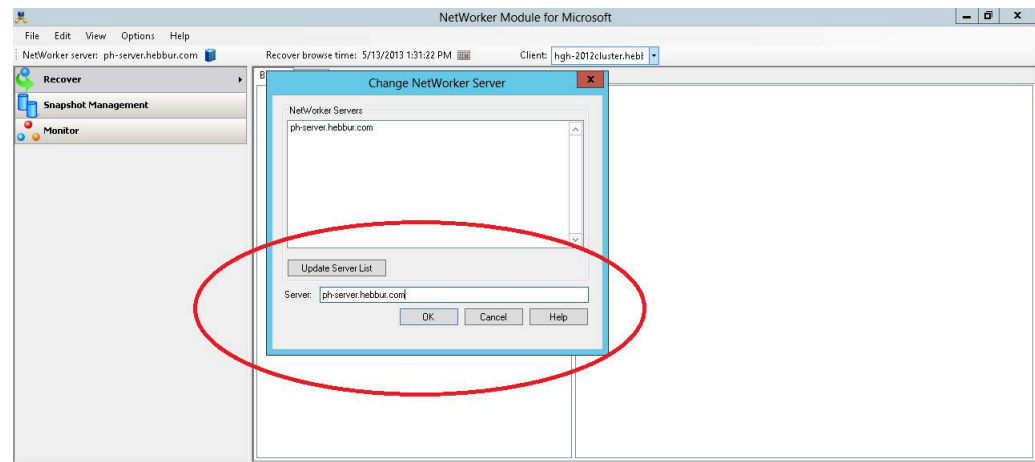
Virtual machine image backups are copy-type backups in-guest for applications. Log grooming requires separate in-guest application backups. The Microsoft documentation provides information about the VSS_BT_FULL backup type.

Recovering items that are stored on a Hyper-V virtual machine

Procedure

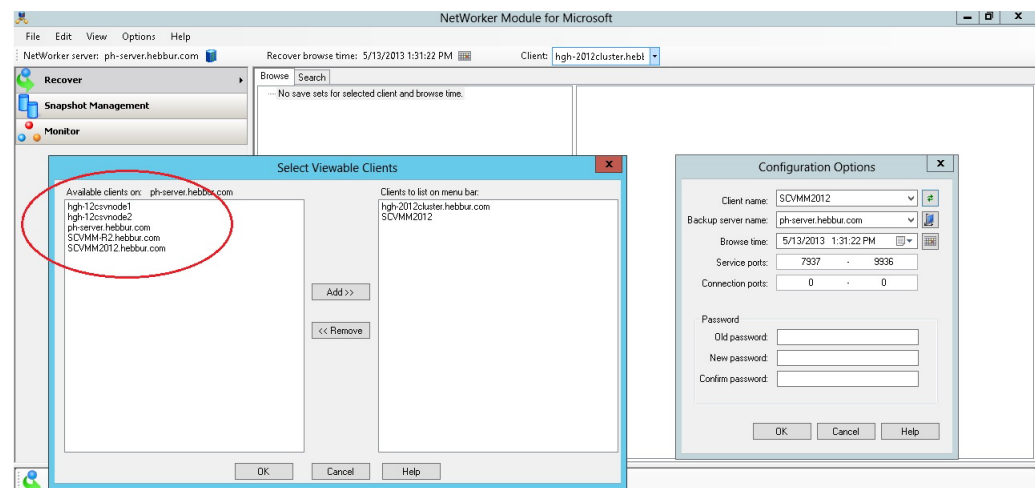
1. Configure the Hyper-V client resources on the NetWorker server, and then select the save set at root-level (**Microsoft Hyper-V**) for backup.
2. Perform a full backup.
3. Open the NetWorker User for Microsoft GUI on the FLR proxy server that you configured for GLR.
4. Select the NetWorker server where you performed the Hyper-V Server backup. Click the icon next to **NetWorker server** to view the list of NetWorker servers in the **Change NetWorker Server** window.

Figure 29 Change NetWorker Server window

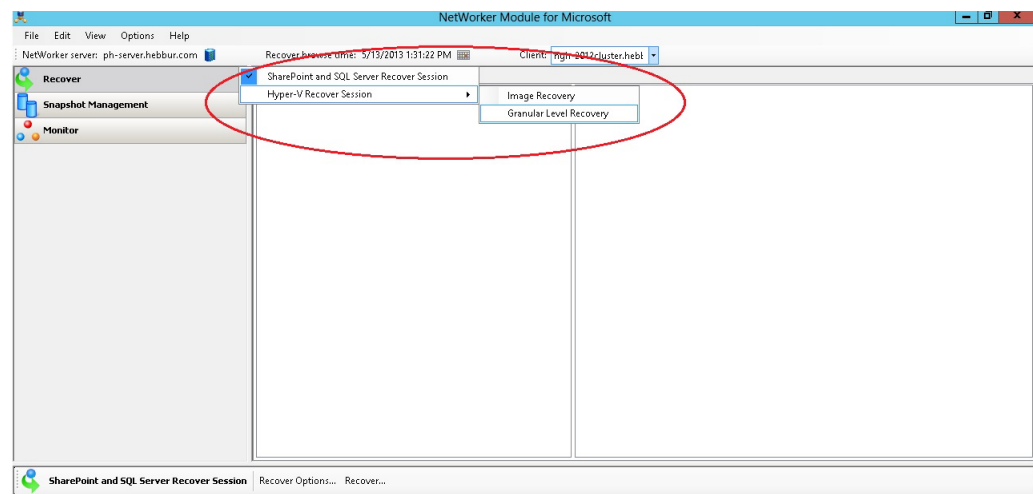


5. Under **Options**, select the **Configure Option** and click the icon next to the **Client name** field in the **Select Viewable Clients** window to view the Hyper-V Server client resources.

Figure 30 Select Viewable Clients window



6. Click **Recover > Hyper-V Recover Session > Granular Level Recovery**.

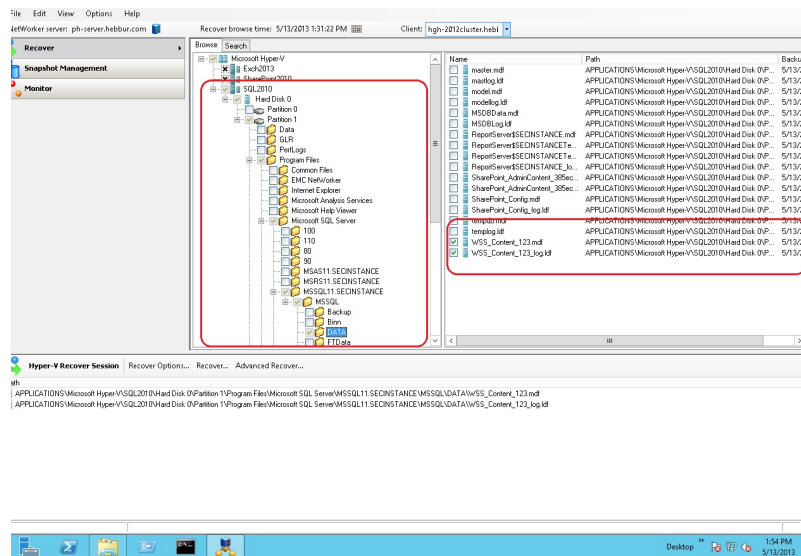
Figure 31 Selecting granular level recovery

Recovering SQL Server items from a Hyper-V virtual machine

Procedure

1. Mount the virtual machine that hosts the SQL Server, attach the hard disk, and then browse to the folder that contains the database and logs from which you will recover the items.

The following figure provides an example:

Figure 32 Selecting SQL Server items for recovery from a Hyper-V virtual machine

2. Select the database (mdf) and logs (ldf) files.
3. Perform the recovery to the folder of your choice.

If the database is offline in SQL Server Management Studio, perform the following steps:

- a. Copy the recovered database and logs files to the actual path.
- b. Bring the database online.

- c. Check that the recovered data is intact.

If the database is online in SQL Server Management Studio with some data corruption or loss, perform the following steps:

- a. Bring the database offline.
- b. Replace the existing database and logs with the recovered database and logs files.
- c. Bring the database online.
- d. Check that the recovered data is intact.

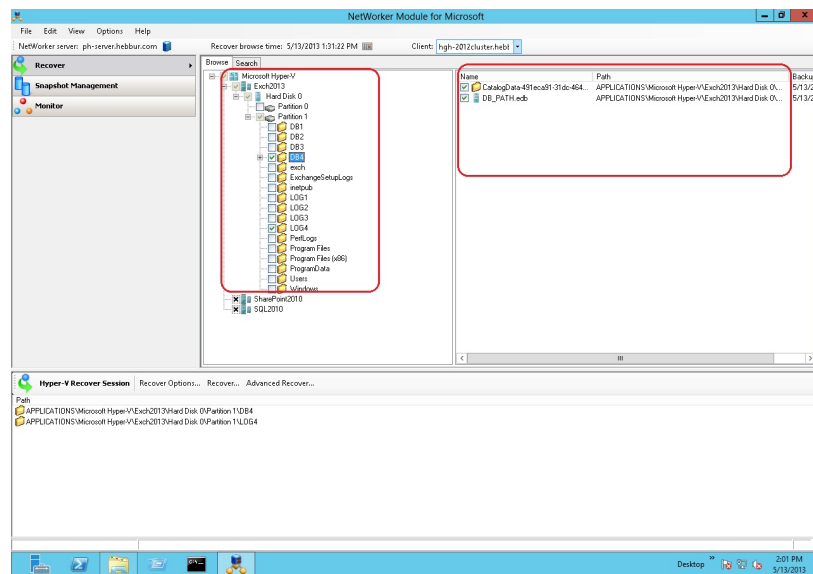
Recovering Exchange Server items from a Hyper-V virtual machine

Procedure

1. Mount the virtual machine that hosts the Exchange Server, attach the hard disk, and browse to the folder that contains the database and logs from which you will recover the items.

The following figure provides an example:

Figure 33 Selecting Exchange Server items for recovery from a Hyper-V virtual machine



2. Select the database and log files.
3. Perform the recovery to the folder of your choice.

If the database is online in the Exchange Management Console with some data corruption or loss, perform the following steps:

- a. Bring the database offline.
- b. Replace the existing database and logs folder with the recovered database and logs folder in the actual path.
- c. Bring the database online.
- d. Check that the recovered data is intact.

If the database is offline in the Exchange Management Console, perform the following steps:

- a. Replace the existing database and logs folder with recovered database and logs folder in the actual path.
- b. Bring the database online.
- c. Check that the recovered data is intact.

Recovering SharePoint Server items

Use EMC ItemPoint for SharePoint Server to perform SharePoint GLR.

Procedure

1. Mount the virtual machine that hosts the SharePoint database, attach the hard disk, and browse to the folder that contains the database and logs from which you will recover the items.
2. Select the database and log files.

NMM mounts the Hyper-V VHD file in a location that you define during GLR recovery. The default location is `C:\Program Files\EMC NetWorker\nsr\tmp\`.

3. Use EMC ItemPoint for SharePoint Server to perform the SharePoint GLR.

Install EMC ItemPoint for SharePoint Server on the SharePoint Server and on the FLR proxy server where you mount the Hyper-V virtual machine. These steps are similar to the procedure described in the *NetWorker Module for Microsoft for SQL and SharePoint VSS User Guide*. In this document, you directly mount the database under SharePoint and SQL Server Recover Session.

However, to recover items from a virtual machine that hosts the SharePoint Server, you must configure EMC ItemPoint for SharePoint Server differently. In **Add the Source Path** for the database, select the path where the Hyper-V VHD is mounted and then browse through the folder to select the database.

For example: `C:\Program Files\EMC NetWorker\nsr\tmp\HyperVMountPoints\SQL2010\Hard Disk 0\Partition1\sqlfirstins\MSSQL11.FIRSTINSTANCE\MSSQL\DATA.`

4. Provide the target SharePoint Server with credentials.

EMC ItemPoint for SharePoint Server configures itself with the SharePoint Server and FLR proxy server by scanning the logs, pre-scanning the logs, hashing the logs, and retrieving the content database.

5. After the EMC ItemPoint for SharePoint Server configuration completes, copy the content to be recovered from the source to the target location.