

# EMC® NetWorker® Module for Microsoft for Hyper-V VSS

Version 8.2 Service Pack 1

## User Guide

302-001-234

REV 01



**EMC<sup>2</sup>**

Copyright © 2007-2015 EMC Corporation. All rights reserved. Published in USA.

Published January, 2015

EMC 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. EMC Corporation 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 EMC software described in this publication requires an applicable software license.

EMC<sup>2</sup>, EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>).

EMC Corporation  
Hopkinton, Massachusetts 01748-9103  
1-508-435-1000 In North America 1-866-464-7381  
[www.EMC.com](http://www.EMC.com)

# CONTENTS

<b>Figures</b>	<b>7</b>	
<b>Tables</b>	<b>9</b>	
<b>Preface</b>	<b>11</b>	
<b>Chapter 1</b>	<b>Introduction</b>	<b>15</b>
	Overview.....	16
	Microsoft Hyper-V environments.....	16
	Hyper-V CSV in a failover cluster.....	16
	How NMM works with Hyper-V.....	17
	Using NMM with Hyper-V.....	18
	Using NMM with Hyper-V in a stand-alone server environment.....	18
	Using NMM in a CSV environment.....	20
	Using NMM with Hyper-V VMs over SMB 3.0.....	21
	Granular level recovery.....	21
	Using NMM with Client Direct to AFTD or Data Domain Boost storage devices.....	22
	Required privileges.....	22
	Required SMB privileges.....	23
	Required Hyper-V CSV privileges.....	23
	Required RDZ privileges.....	24
	Example Hyper-V configurations.....	24
	Hyper-V on physical server configurations.....	24
	Hyper-V storage configurations.....	25
	Hyper-V configuration requirements for backing up a VM that contains multiple volumes.....	27
	Backup overview.....	27
	Backup types.....	28
	Types of supported backup.....	28
	Files included in backups.....	28
	Backup process workflow.....	29
	Recovery overview.....	34
	Types of supported recovery.....	34
	Recovering to the original Hyper-V Server.....	35
	Redirected recovery of a VM to an alternate Hyper-V Server.....	35
	Special character considerations.....	35
<b>Chapter 2</b>	<b>Backups</b>	<b>37</b>
	Planning backups.....	38
	NMM capabilities for Window Server 2012 and 2012 R2 environments .....	38
	Integration services components version.....	39
	Viewing valid application data save sets.....	39
	Configuring backups.....	40
	Configuring a Hyper-V client resource.....	43
	Creating the client by using the Client Configuration Wizard.....	43

	Creating the client manually by using the NetWorker Management Console.....	46
	Performing cluster-level and CSV VM backups.....	47
	Configuring multi-proxy backups.....	48
<b>Chapter 3</b>	<b>Recoveries</b>	<b>51</b>
	Overview.....	52
	Turning the VM offline for recovery.....	52
	Hosting the recovered virtual systems.....	52
	Specifying the destinations for the Hyper-V configuration files and virtual system.....	52
	Selecting Hyper-V recovery options.....	53
	Selecting the Hyper-V recovery destination.....	53
	Performing Hyper-V recovery to the original machine and location....	53
	Performing a directed Hyper-V recovery to a different machine or location.....	54
	Recovering Hyper-V CSV VMs.....	56
	SMB 3.0 VM recovery.....	56
	CSV VM recovery.....	56
	Recovering multiple CSV VMs to the original location .....	57
	Recovering an individual CSV Hyper-V VM to a different location.....	57
	Recovering with a Windows Server 2012 and 2012 R2 proxy CSV server.....	58
	Troubleshooting RPC service failure messages.....	59
<b>Chapter 4</b>	<b>Granular Level Recoveries</b>	<b>61</b>
	Overview.....	62
	Windows Server 2012 Hyper-V GLR features.....	62
	Recovering Hyper-V files and folders.....	62
	Recovering AES-encrypted data from a Hyper-V VM.....	64
<b>Chapter 5</b>	<b>EMC Data Protection Add-in for SCVMM</b>	<b>65</b>
	Overview.....	66
	Recoveries.....	66
	Backups.....	66
	Supported versions.....	66
	Software dependencies.....	67
	Required privileges.....	67
	Installation and configuration overview.....	68
	How the Data Protection Add-in works with SCVMM.....	68
	Workflows overview.....	69
	GUI overview.....	69
	SCVMM user roles and allowed actions.....	70
	Supported scopes and contexts.....	70
	Installation and uninstallation.....	70
	Installing SCVMM and the SCVMM console.....	71
	Installing SCVMM Update Rollups.....	71
	Installing the Data Protection Add-in.....	71
	Importing the Data Protection Add-in.....	72
	Activating the Data Protection Add-in.....	72
	Uninstalling the Data Protection Add-in.....	73
	Upgrading the Data Protection Add-in.....	74
	Configuration.....	75

Adding NetWorker servers.....	75
Removing NetWorker servers.....	76
Setting the refresh interval.....	76
Including debug output for logging purposes.....	76
Using multiple NetWorker Servers that define the same clients and VM savesets.....	76
<b>Data Protection Add-in overview data.....</b>	<b>77</b>
<b>Recoveries.....</b>	<b>81</b>
Viewing available VMs.....	83
VM Encrypted Recoveries.....	83
Recovering a VM to the original location.....	84
Redirected recoveries.....	85
Recovering a deleted VM.....	88
<b>Monitoring.....</b>	<b>88</b>
<b>Troubleshooting.....</b>	<b>89</b>
Recovered VM doesn't start.....	89
Installation fails due to access issue.....	89
Importing fails due to access issue.....	89
VM attributes might display incorrect values.....	90
Redirected recovery appears to succeed but no VM appears in Hyper-V Manager.....	90
Checks for redirected recovery failures .....	90
Avoid VM names with the same name within an SCVMM context.....	90
Cluster VM backups do not display on the Recover page.....	91
Unable to recover if 'localhost' used as NetWorker server name.....	91
Redirected recovery fails when the VM name or VM configuration path contains special characters.....	91
<b>Chapter 6 Best Practices and Recommendations</b>	<b>93</b>
<b>Overview.....</b>	<b>94</b>
Hyper-V Server backup and recovery best practices.....	94
Hyper-V VM applications backup and recovery best practices.....	95
Improving backup performance in Windows Server 2012 and 2012 R2 clusters with CSV.....	95
Data mining using Hyper-V granular level recovery.....	96
Restrictions and requirements for relocating and recovering data to a different location.....	96
Restrictions for backup and recovery of Hyper-V VMs in a Windows Server Failover Cluster.....	97
Restrictions for Hyper-V VM granular level recovery.....	97
<b>Chapter 7 Troubleshooting</b>	<b>99</b>
<b>Troubleshooting backups.....</b>	<b>100</b>
T	100
<b>Appendix A Recovering SQL Server, Exchange Server, and SharePoint Server Items from a Hyper-V VM</b>	<b>105</b>
<b>Overview.....</b>	<b>106</b>
<b>Recovering items.....</b>	<b>106</b>
<b>Recovering SQL Server items.....</b>	<b>107</b>
<b>Recovering Exchange Server items.....</b>	<b>108</b>
<b>Recovering SharePoint Server items.....</b>	<b>109</b>

Glossary	113
----------	-----

# FIGURES

1	Two-node Hyper-V failover cluster.....	17
2	Guest backup and recovery environment.....	18
3	Image-level backup and recovery environment.....	19
4	Granular level recovery environment.....	22
5	Windows Server 2008 host with Hyper-V VMs.....	24
6	NMM backup of Hyper-V components.....	25
7	Hyper-V storage options.....	26
8	Image-level VSS backup workflow.....	30
9	Image-level saved state backup workflow.....	32
10	Image-level backup workflow for federated backups.....	33
11	Data Protection Add-in architecture.....	68
12	Data Protection Add-in for SCVMM Preferences page.....	75
13	Data Protection Add-in Overview page for Administrator, Fabric Administrator, and Read-Only Administrator user roles.....	78
14	VM Protection Details tooltip for Administrator , Fabric Administrator, and Read-Only Administrator user roles.....	78
15	VM Protection Details window for Administrator, Fabric Administrator, and Read-Only Administrator user roles.....	79
16	Data Protection Add-in Overview page for Tenant Administrator and Application Administrator user roles .....	80
17	Virtual Machine Backup Status tooltip for Tenant Administrator and Application Administrator user roles.....	80
18	VM Protection Details window for Tenant Administrator and Application Administrator user roles.....	81
19	Data Protection Add-in for SCVMM Recover page.....	82
20	Data Protection Add-in for SCVMM Monitoring page.....	88

## FIGURES

# TABLES

1	Revision history.....	12
2	Comparison of guest and image-level backup and recovery.....	19
3	Access privileges needed for backup and recovery .....	23
4	Hyper-V VM configurations.....	26
5	Types of supported backups.....	28
6	VM files supported by the VSS Hyper-V Writer.....	29
7	Types of supported recoveries.....	34
8	Common special characters and their URL-encoded values.....	36
9	Backup tasks for Hyper-V .....	40
10	Hyper-V save set syntax.....	41
11	Hyper-V application information variable settings.....	42
12	SCVMM user roles and actions allowed by the Data Protection Add-in.....	70
13	VM IDs after redirected recovery.....	85
14	Backup types.....	106

## TABLES

# Preface

As part of an effort to improve its product lines, EMC periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.

Contact your EMC technical support professional if a product does not function properly or does not function as described in this document.

---

## Note

This document was accurate at publication time. Go to EMC Online Support (<https://support.emc.com>) to ensure that you are using the latest version of this document.

---

## Purpose

This guide contains information about using the NetWorker Module for Microsoft (NMM) Release 8.2 SP1 software to back up and recover Hyper-V VMs using the Volume Shadow Copy Service (VSS) technology.

---

## Note

The *NetWorker Module for Microsoft Administration Guide* supplements the backup and recovery procedures described in this guide and must be referred to when performing application-specific tasks. Ensure to download a copy of the *NetWorker Module for Microsoft Administration Guide* from EMC Online Support (<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:

- EMC NetWorker software
- EMC NetWorker snapshot management
- Microsoft Volume Shadow Copy Service (VSS) technology

## Revision history

The following table presents the revision history of this document.

**Table 1** Revision history

Revision	Date	Description
01	January 28, 2015	First release of this document for EMC NetWorker Module for Microsoft release 8.2 SP1.

## 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 VSS User Guide*
- *NetWorker Module for Microsoft for Windows Bare Metal Recovery Solution User Guide*
- *NetWorker Module for Microsoft Advanced Recovery Guide*
- *NetWorker Performing backup and recovery of SharePoint Server by using NetWorker Module for Microsoft SQL VDI solution Technical Notes*
- *NetWorker Performing Exchange Server Granular Recovery by using NetWorker Module for Microsoft with Ontrack PowerControls Technical Notes*
- *NetWorker SharePoint BLOB Backup and Recovery by using NetWorker Module for Microsoft and Metalogix StoragePoint Technical Notes*

## Special notice conventions used in this document

EMC uses the following conventions for special notices:

### NOTICE

Addresses practices not related to personal injury.

### Note

Presents information that is important, but not hazard-related.

## Typographical conventions

EMC uses the following type style conventions in this document:

*Italic* Use for full titles of publications referenced in text

Monospace Use for:

- System code
- System output, such as an error message or script
- Pathnames, file names, prompts, and syntax
- Commands and options

*Monospace italic* Use for variables

<b>Monospace bold</b>	Use for user input
[ ]	Square brackets enclose optional values
	Vertical bar indicates alternate selections - the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate non-essential information omitted from the example

---

## Where to get help

EMC support, product, and licensing information can be obtained as follows:

### Product information

For documentation, release notes, software updates, or information about EMC products, go to EMC Online Support at <https://support.emc.com>.

### Technical support

Go to EMC Online Support and click Service Center. You will see several options for contacting EMC Technical Support. Note that to open a service request, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

### Online communities

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

### Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to [DPAD.Doc.Feedback@emc.com](mailto:DPAD.Doc.Feedback@emc.com)



# CHAPTER 1

## Introduction

This chapter includes the following sections:

• <a href="#">Overview</a> .....	16
• <a href="#">Microsoft Hyper-V environments</a> .....	16
• <a href="#">How NMM works with Hyper-V</a> .....	17
• <a href="#">Using NMM with Hyper-V</a> .....	18
• <a href="#">Required privileges</a> .....	22
• <a href="#">Example Hyper-V configurations</a> .....	24
• <a href="#">Backup overview</a> .....	27
• <a href="#">Recovery overview</a> .....	34
• <a href="#">Special character considerations</a> .....	35

## 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 single physical server to consolidate workloads. EMC NetWorker Module for Microsoft (NMM) provides image level backup and recovery of the Microsoft Hyper-V role installed on Windows Server 2008, 2008 R2, 2012, and 2012 R2 and on Server Core installations for Windows Server 2008, 2008 R2, 2012, and 2012 R2.

NMM utilizes the VSS infrastructure, including writers and providers, to back up and recover each VM and the Hyper-V Initial Store configuration file (or, in Windows Server 2012 and 2012 R2, the host component/parent partition)

NMM supports Hyper-V cluster shared volume (CSV) backup and recovery on Microsoft Windows Server 2008 R2, 2012, or 2012 R2.

For Hyper-V cluster and CSV environments, including proxy environments, you must install the NetWorker client and NMM on all nodes in the cluster.

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

## Microsoft Hyper-V environments

The Microsoft Hyper-V documentation provides a complete and updated list of system requirements and supported guest operating system versions.

The *EMC NetWorker Online Software Compatibility Guide* on EMC Online Support lists the most up-to-date information about the operating systems and versions that NMM supports.

## Hyper-V CSV in a failover cluster

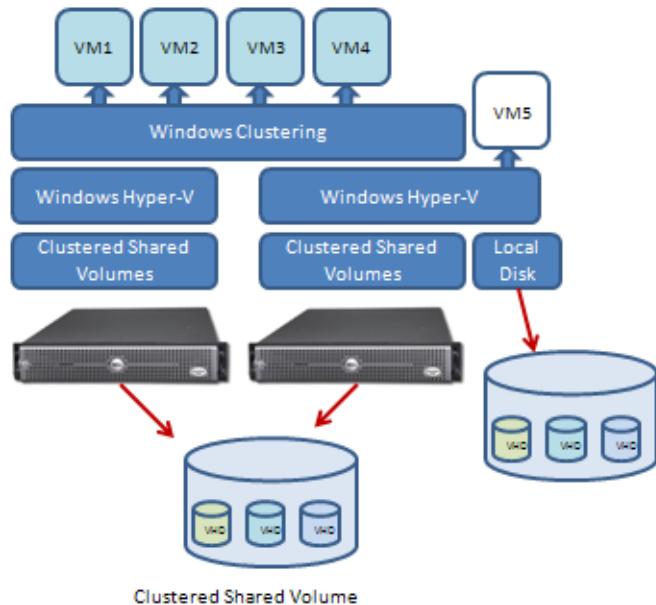
To prevent a Hyper-V Server from becoming a single point of failure, you can create a failover cluster. In a failover cluster, all servers (nodes) run Hyper-V and can host one or more VMs. A VM can run on multiple nodes in the cluster, but can only be active on one node at a time.

A failover cluster usually includes a shared storage device that is physically connected to all servers in the cluster. However, only one server at a time can access each storage volume.

NMM supports failover clustering for Hyper-V through Clustered Shared Volumes (CSV). A CSV is a standard cluster disk with an NTFS volume that is accessible for read and write operations by all cluster nodes. This gives a VM complete mobility through the cluster, as any node can be an owner of the VM.

For Windows Server 2008 R2 environments, throughout the duration of the backup, the CSV is in redirected I/O mode. Other nodes cannot directly write to disks. Instead, the I/O is redirected over the LAN to the owner node performing the backup.

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

**Figure 1** Two-node Hyper-V failover cluster

- NMM protects failover cluster configurations by performing a CSV backup. NMM performs the backup and recovery of the VM on the cluster node where the VM is currently active.
- Take the VM offline before you perform the recovery.
- Perform a CSV backup of the cluster.

[NMM federated architecture for Hyper-V failover clusters on page 31](#) provides additional information.

## How NMM works with Hyper-V

Hyper-V is a configurable feature on Windows Server 2008, 2008 R2, 2012, and 2012 R2 that you can use to host the VMs.

Each VM is usually a server operating system that runs Microsoft applications, such as:

- Exchange Server
- SharePoint Server
- SQL Server
- Data Protection Manager

Hyper-V runs as a role in Windows Server 2008, 2008 R2, 2012, and 2012 R2. NMM uses the Hyper-V VSS writer (for Windows Server 2008 and 2008 R2) and the Hyper-V VSS writer and Clustered Share Volumes VSS writer (for Windows Server 2012 and 2012 R2) on the host to back up and recover Hyper-V data by using APPLICATION save sets. The Hyper-V Writer backs up and recovers Hyper-V configuration and VM files.

VSS is a framework that enables volume backups to be performed while applications on a system continue to write to the volumes. The Hyper-V VSS writer enables the creation of image backups for VMs by quiescing the Windows operating system and applications within the guest for operating system and application consistency.

For Windows Server 2012 and 2012 R2 environments, when you install integration components on a VM, and the VM runs an application such as Exchange or SQL on that VM, the Hyper-V backup takes a copy-type backup of the application data. For Windows

Server 2008 R2 and earlier environments, the Hyper-V backup takes a full backup of the application data.

## Using NMM with Hyper-V

You can use NMM with Hyper-V in stand-alone or clustered environments, over SMB 3.0, and with Client Direct to AFTD or Data Domain Boost devices.

### Using NMM with Hyper-V in a stand-alone server environment

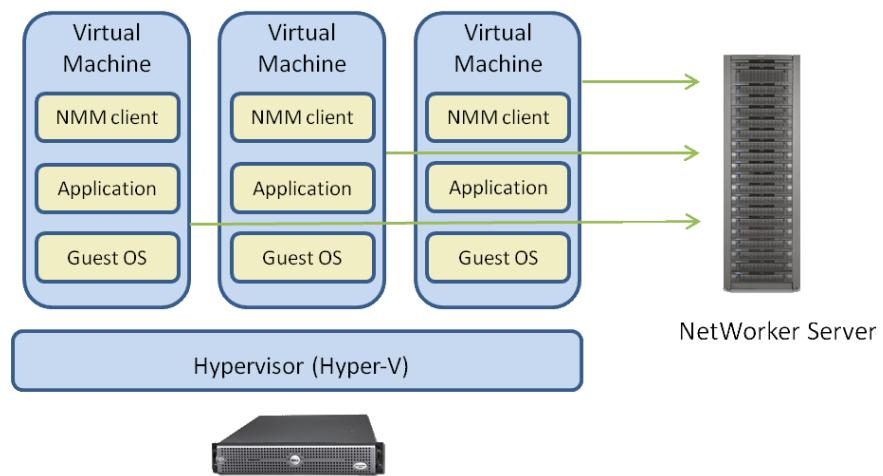
You can use NMM to protect a stand-alone Hyper-V environment at the guest and image level.

#### Guest backup and recovery

With guest backup and recovery, you install an NMM client on each VM that hosts databases or specific applications on the Hyper-V server, for example Microsoft Exchange or Microsoft SharePoint. NMM considers each VM to be a separate client, and you can perform individual backups of each VM and Microsoft application.

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

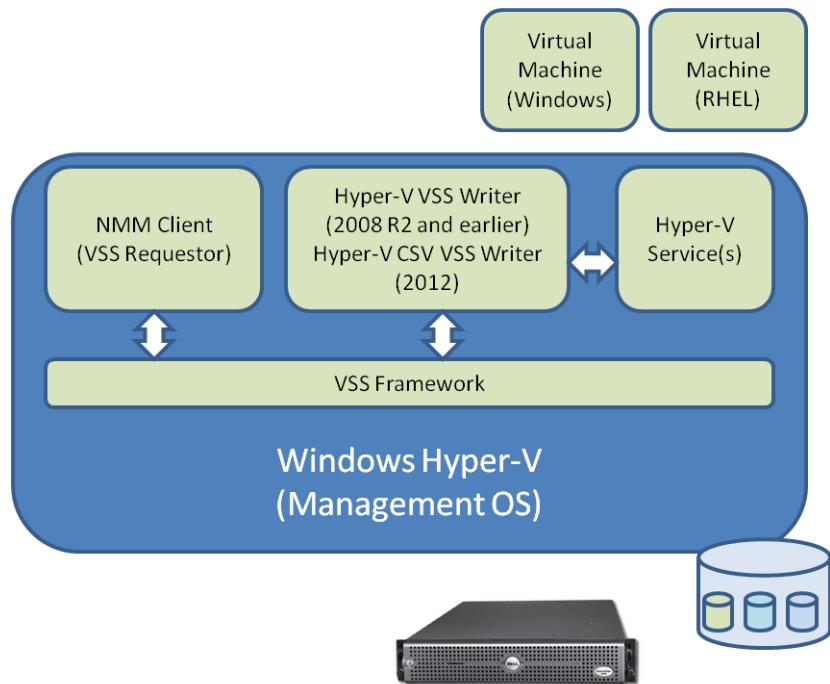
**Figure 2** Guest backup and recovery environment



#### Image-level backup and recovery

With image-level backup and recovery, you install the NMM client on the Hyper-V Management operating system or the parent partition.

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

**Figure 3** Image-level backup and recovery environment

You can perform full image-level backups of individual VMs or the initial store. For Windows Server 2012 and 2012 R2, you can perform a full image-level backup of the individual VMs and the host component. The initial store and host component contain the role-based security configuration for Hyper-V. Image-level backups occur from the Hyper-V Management operating system instead of from the individual VMs.

## Comparing NMM protection methods for Hyper-V

You can choose whether to perform Hyper-V guest or image-level backup and recovery depending on criterion such as user knowledge of Hyper-V, the Windows operating system running on the guest, and where the NMM software is installed.

The following table provides a comparison of the two methods.

**Table 2** Comparison of guest and image-level backup and recovery

Criterion	Guest backup and recovery	Image-level backup and recovery
User knowledge of Hyper-V	No advanced Hyper-V knowledge is required	Requires advanced Hyper-V knowledge
Windows guest operating system	Windows guest operating systems that Hyper-V supports, through the use of NMM clients	All guest operating systems that Hyper-V supports
NMM software installation	NMM and NetWorker on each guest operating system	NMM and NetWorker on the management operating system
NetWorker server network connection	Required for each VM	Required only for the Hyper-V server

**Table 2** Comparison of guest and image-level backup and recovery (continued)

Criterion	Guest backup and recovery	Image-level backup and recovery
Deduplication with the appropriate Data Domain or Avamar device	Data within each VM	Data at image-level
Support for iSCSI/pass-through media	Yes	No
Support for individual backup of each VM	Yes	Yes
Application-aware backup and recovery	Yes, with NMM for applications such as: <ul style="list-style-type: none"> <li>• Microsoft Exchange Server</li> <li>• Microsoft SharePoint Server</li> <li>• Microsoft SQL Server</li> <li>• Microsoft Active Directory</li> </ul>	No
VM status for backup	VM must be running	VM does not need to be running
Backup consumption of CPU, RAM, and disk resources	On the VM	On the Hyper-V server
Backup customization, including exclusion of certain files or file types	Yes	No
Recovery of individual files and folders	Yes	By using GLR
Disaster recovery requirements	Windows bare metal recovery (BMR) uses a two-step recovery: Recover the operating system state critical volumes.  Use NMM to recover applications and non-critical volume data.	One-step recovery of backup data from NMM. However, backups are a “crash-consistent” snapshot of the full VM image, which might or might not reliably support a full system recovery without data loss.

## Using NMM in a CSV environment

CSV is a feature of failover clustering available in Windows Server 2008 R2, 2012, and 2012 R2 for use with the Hyper-V role. CSV is available for Hyper-V VMs created with Windows Server 2008 R2, 2012, or 2012 R2 and is supported by NMM.

A CSV is a clustered disk that contains an NTFS volume. Each node within the cluster can access the volume for read and write operations. This gives the VM complete mobility throughout the cluster as any node can be the VM owner, and changing owners is easy.

A CSV is not owned by any one node in the cluster. Instead, the CSV travels between cluster nodes as the backup and recoveries occur. Microsoft and NMM refer to the node in the cluster where a CSV is locally mounted as the “coordinating node”. NMM must

perform the backup and recovery operations from the coordinating node of the CSV of each VM that participates in the backup or recovery operation.

To perform a backup or recovery operation, NMM locates the VM cluster node that owns the CSV and then makes the node the CSV coordinating node. The same is true for clustered VM recovery: NMM finds the currently configured node for the VM and recovers the VM to that cluster node after making it the CSV coordinating node. When a VM already exists in the cluster, you must perform the recovery operation on the cluster node that owns the VM.

NMM supports physical proxy nodes for Windows Server 2012 and 2012 R2 Hyper-V CSV backups. When you specify a Preferred Server Order List (PSOL) in the Application Information attribute for the client resource of Cluster Server Name, NMM performs shadow copies and backups on each proxy node. The shadow copies are done serially, one proxy node at a time. After all shadow copies are successfully completed, the proxy nodes perform VM data backups in parallel. The recovery process is the same as for a normal Hyper-V VM.

For Hyper-V cluster and CSV environments, including proxy environments, you must install the NetWorker client and NMM on all nodes in the cluster.

## Using NMM with Hyper-V VMs over SMB 3.0

NMM supports Hyper-V VMs residing on Windows Server 2012 and 2012 R2 SMB 3.0. Windows Server 2012 and 2012 R2 allows Hyper-V VMs to store their data on SMB 3.0 shares and provides capabilities to take snapshots and back up the data remotely.

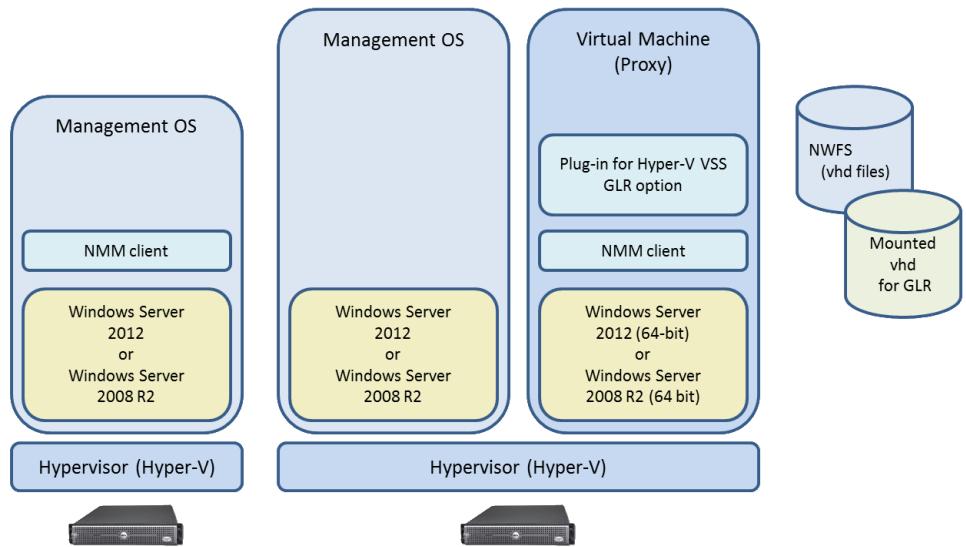
To back up application-consistent data, install NMM on the Hyper-V servers. The storage location presents no difference in configuring backups and in performing backups and recoveries for a VM on a stand-alone server or failover cluster. The same operations that protect local VMs also apply to the VMs on SMB file shares. [Required SMB privileges on page 23](#) describes the required permissions for SMB backup and recovery.

## Granular level recovery

When you perform image-level backups with NMM, you can use granular level recovery (GLR) to recover an image backup to a temporary file system on a different client, and then browse and recover individual files and folders.

To browse and recover individual files and folders, select the GLR option when you install NMM on the virtual or physical machine that you want to use for GLR. This machine is typically a different machine than the management operating system.

The following figure illustrates an environment where NMM is installed on a VM to perform a GLR.

**Figure 4** Granular level recovery environment

## Using NMM with Client Direct to AFTD or Data Domain Boost storage devices

You can store Hyper-V backups on the NetWorker server, on an AFTD device, or on an EMC Data Domain® system. By default, NMM stores backups on devices that are local to the NetWorker server.

The NMM software supports the NetWorker Client Direct feature. The Client Direct feature:

- Enables clients with network access to AFTD or Data Domain Boost storage devices to send their backup data directly to the devices, bypassing the NetWorker storage node. The storage node manages the devices for the NetWorker clients, but does not handle the backup data.
- Reduces bandwidth usage and bottlenecks at the storage node.
- Provides highly efficient backup data transmission.

Destination devices must specify their complete paths in their Device Access Information attribute. If the Client Direct backup is not available, NMM performs a traditional storage node backup instead. When you create an NMM client resource in NMC, NetWorker enables the Client Direct feature by default, but you can disable the Client Direct feature in each client resource.

The nmm.raw backup log will display details about the Client Direct activity for the Hyper-V server.

The *NetWorker Administration Guide* provides details about the Client Direct to AFTD or Data Domain Boost storage devices.

## Required privileges

The required privileges for backing up and recovering Hyper-V VMs are the same as other applications. However, backup and recovery over SMB 3.0 or in a CSV or RDZ environment require additional privileges.

The *NetWorker Module for Microsoft Administration Guide* and the *NetWorker Module for Microsoft Installation Guide* provide additional details.

## Required SMB privileges

SMB backup and recovery requires additional privileges beyond Hyper-V backup and recovery privileges.

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

**Table 3** Access privileges needed for backup and recovery

SMB configuration	Required privileges
All	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• Add backup permissions for the backup user on all file servers in the cluster.</li> <li>• Add the backup user as the cluster administrator (domain administrator).</li> </ul>
File server scale out	In the Local Backup operator group of each SMB node, configure the application server as a member of the Backup Operators group.
Cluster	Add each CSV node to the SMB nodes of the Local Backup Operator group.

Verify that Replication Manager (RM) replication service is running under an account that has backup permissions on all file servers or domain administrator permissions. Verify that the Hyper-V server and the file server are in the same domain. Recoveries require the same permissions as the backup user.

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.

## Required Hyper-V CSV privileges

Hyper-V CSV backup and recovery requires additional privileges beyond Hyper-V backup and recovery privileges. You must create a Domain User for Hyper-V backups and recoveries. During client resource configuration for NMM Hyper-V backups, provide this Domain User account and password (instead of providing a Domain Administrator account and password) for backup and recovery.

### Procedure

1. Create a Domain User for Hyper-V backups and recoveries.
2. Add the following groups to the newly created Domain User:
  - Backup Operators
  - Hyper-V Administrator
  - Windows Authorization Access Group
  - Users
  - Remote Desktop Users
  - Add Group Policy User Control
3. On each cluster node, log in and perform the following steps:
  - a. Provide local administrator privileges to the Domain User.
  - b. Provide access for cluster management to the group. Open PowerShell and type this command:

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

## Required RDZ privileges

NMM supports NetWorker Restricted Data Zones (RDZ). An RDZ adds an additional permission checking layer, which ensures that RDZ administrators accessing areas that have not been specifically coded for this feature, by default, do not have access to those areas.

To perform SMB and CSV backups and recoveries in an RDZ, you must configure 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.

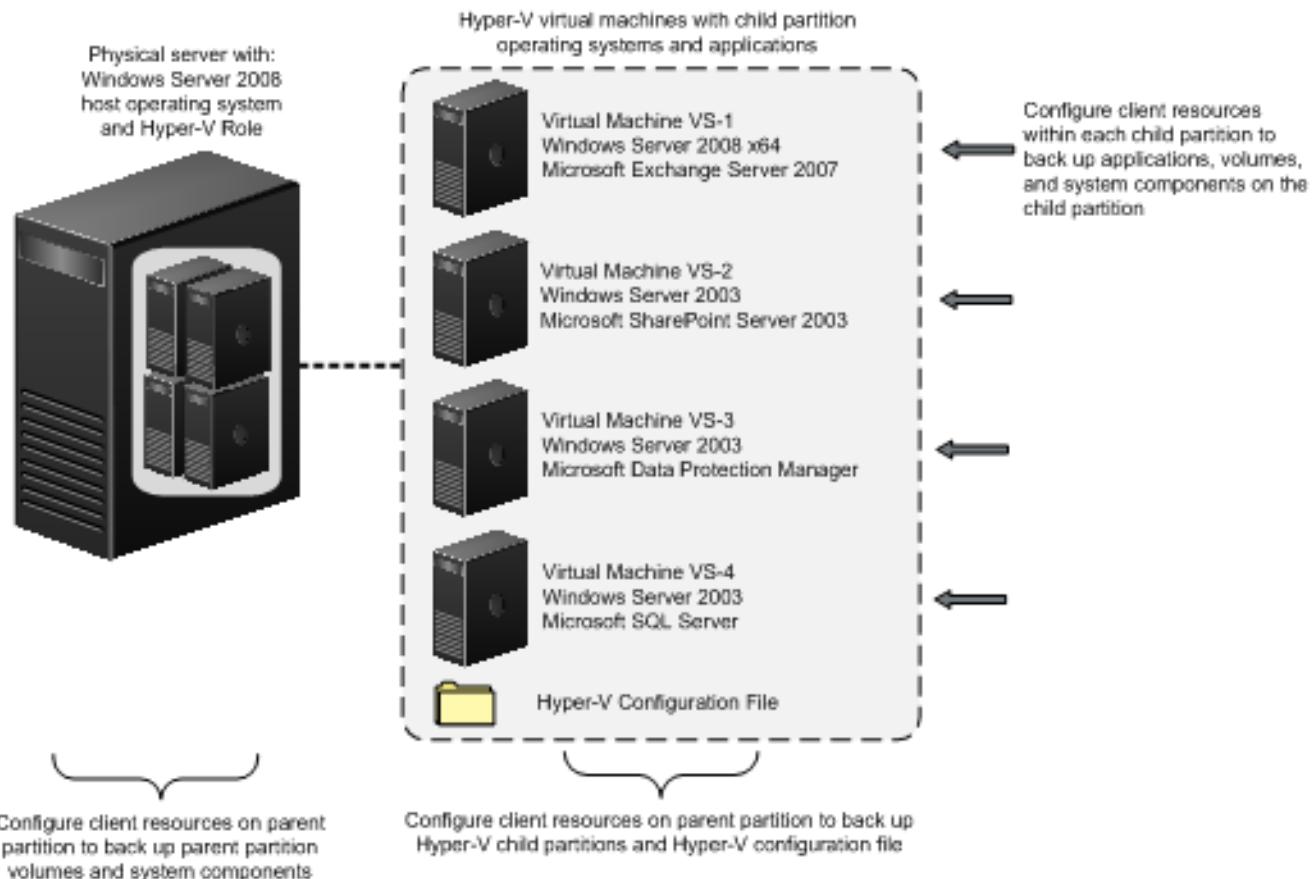
## Example Hyper-V configurations

This section describes some of the possible Hyper-V configurations.

### Hyper-V on physical server configurations

The following figure illustrates a physical server that runs Windows Server 2008. This configuration also applies to Windows Server 2008 R2, 2012, and 2012 R2. The Hyper-V role has been enabled on the physical server, and four VMs have been created, each running a separate operating system and different Microsoft applications.

**Figure 5** Windows Server 2008 host with Hyper-V VMs

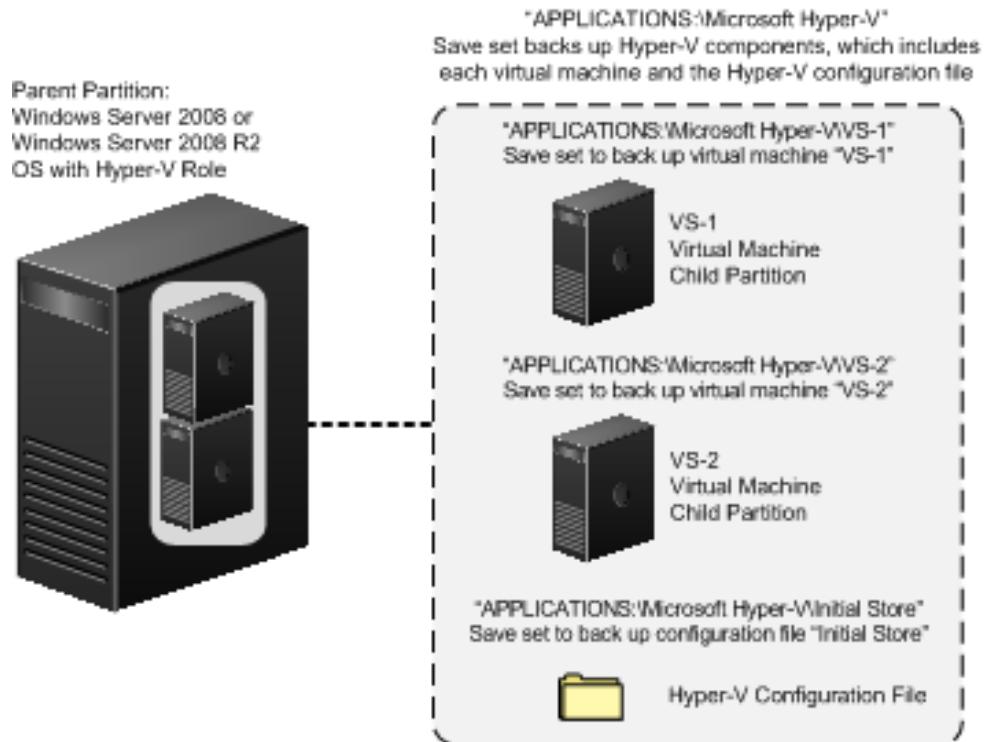


For complete data protection, configure client resources for each of the following:

- Hyper-V VMs on the Hyper-V Server
- The applications within each VM

The following figure describes what the NMM client backs up in Hyper-V, by using the Microsoft Hyper-V VSS Writer and NMM save sets.

**Figure 6** NMM backup of Hyper-V components



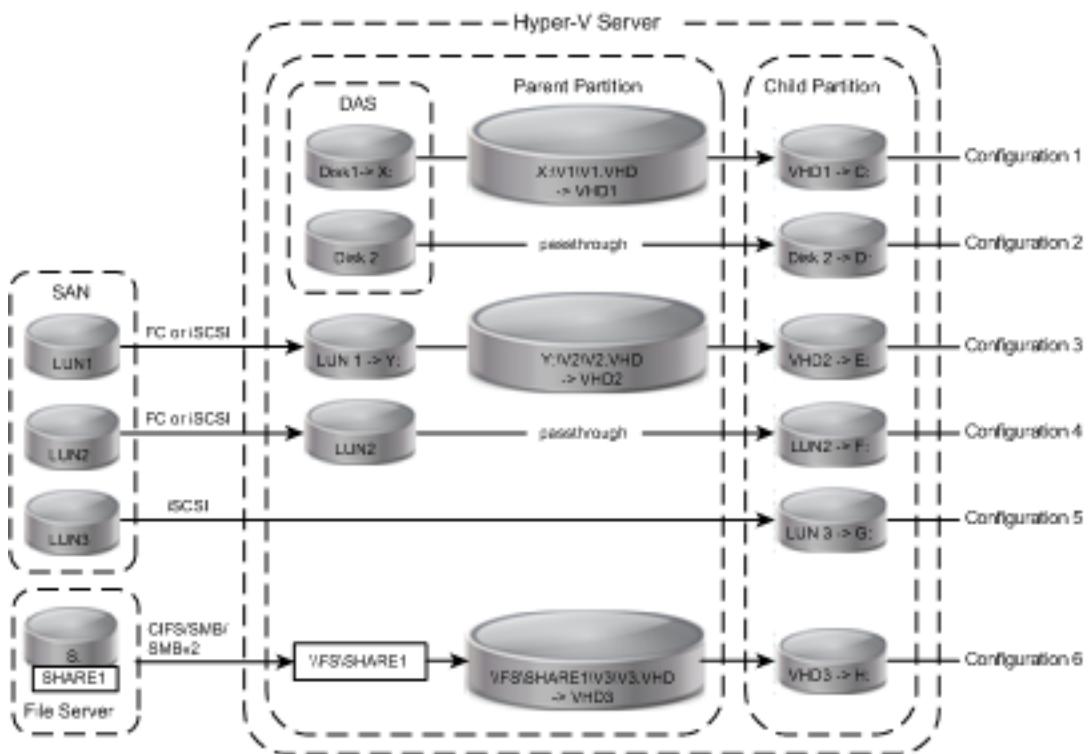
## Hyper-V storage configurations

There are a wide variety of storage configurations available for Hyper-V VMs, such as passthrough disks, direct-attached storage (DAS), storage area networks (SANs), and file servers.

The following documentation provides more details about hardware and software requirements for Hyper-V backup and recovery operations:

- The Microsoft website provides more details and the most up-to-date information about storage hardware supported by Hyper-V.
- The *NetWorker Online Software Compatibility Guide* provides the most up-to-date information about supported software for Hyper-V backup and recovery in NMM.

The following figure illustrates Hyper-V storage options.

**Figure 7** Hyper-V storage options

The following table lists the Hyper-V VM configurations.

**Table 4** Hyper-V VM configurations

Configuration	Type
1	VHD1 on DAS
2	DAS passthrough
3	VHD2 on LUN
4	LUN passthrough
5	iSCSI target attached within VM
6	VHD3 on file server

NMM supports Hyper-V snapshots of virtual and physical machines with the Microsoft VSS provider, depending on the hardware storage type and partition type.

The following list describes how NMM supports the configuration types listed in the table above:

- Snapshot support — NMM supports snapshots for both the Hyper-V server and VMs:
  - **VM** — Install NMM on the VM to perform the backup. If you are using the Microsoft Software VSS provider, NMM supports all configurations.
  - **Hyper-V server** — Install NMM on the parent to perform the backup. If you are using the Windows VSS system provider, NMM supports configuration 1 (VHD1), configuration 3 (VHD2), and configuration 6 (SMB 3.0).
- Cluster support — In addition to the supported configurations listed for snapshots, parent and child cluster scenarios support the following storage configurations:

- **VM clustering** — Install NMM on the VM to perform the backup:  
For failover resource drives, NMM supports configuration 5 (LUN exposed directly to VM). This is the only configuration that Microsoft currently supports for Windows Server 2008 Failover Clustering.  
For operating system drives or local drives for the cluster nodes, the VM support listed under Snapshot support applies.
- **Hyper-V server clustering** — Install NMM on the parent to perform the backup.  
NMM supports all configurations.

When performing VM backups while executing on the Hyper-V server, the Microsoft Hyper-V Writer does not include the passthrough or child-attached iSCSI drives for a VM. Configurations such as 2, 4, and 5 are not supported by the Hyper-V Writer. Configuration 6 is not supported because the VSS framework does not support network shares for Windows Server 2008 and 2008 R2.

## Hyper-V configuration requirements for backing up a VM that contains multiple volumes

When there are multiple virtual hard disks in the guest, the backup of the associated VM from the Hyper-V server might fail because of a Microsoft limitation. When there are multiple volumes on the guest, VSS determines the shadowstorage area for the snapshots based on which volume has more space. This can lead to a condition where the snapshots of volumes C and D both reside on volume D because volume D has more space. During the snapshot revert stage, PostSnapshot, the snapshot of volume C snapshot might be lost if the snapshot of volume D snapshot is reverted first.

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 VM.

## Backup overview

You can perform full image-level backups of individual VMs or the initial store (or, in Windows Server 2012 and 2012 R2, the host component), which contains the role-based security configuration for Hyper-V.

Include the following backups in the backup strategy for a Hyper-V environment:

- Stand-alone Hyper-V servers and Hyper-V images
- Clustered Shared Volumes in a Hyper-V environment

Perform these backups regularly on either an on-demand or scheduled basis.

### NOTICE

---

NMM image-level backups do not back up the management operating system. To protect the Hyper-V management operating system, perform a disaster recovery backup.

## Backup types

You can perform both application and crash consistent image-level backups with NMM.

Application and crash consistent backups differ in the following ways:

- With an application-consistent backup, VSS runs in-guest and freezes the operating system and all application states. Ensure that the VM is online and VSS-capable, and ensure that you have installed Microsoft Integration Components (IC).
- With a crash consistent backup, the VM is offline or does not have the IC installed. In this case, the VM is paused before shadow copy creation and resumed after the shadow is created.

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 the backup type when you perform on-demand or scheduled backups.

With image-level saved state or offline backups, the backup operation puts the VM into a saved state during the processing of the `PrepareForSnapshot` event. The backup process takes snapshots of the appropriate volumes and then returns the VM to the previous state during the processing of the `PostSnapshot` event.

Use image-level saved state backups when you cannot install Integration Components (IC) on the Windows clients or when the guest operating system (for example, Linux) does not support VSS. These backups do not communicate with the Hyper-V VSS writer in the VM. As a result, they ensure crash consistency, not application consistency, of the writers running in the VM.

## Types of supported backup

You can perform disaster recovery backups, federated image-level backups of clusters, image level backups of stand-alone servers, and backups over SMB 3.0.

The following table lists the supported types of backup.

**Table 5** Types of supported backups

Type of backup	Includes
Disaster recovery backup for Hyper-V server	Does not include VMs and Initial Store (or, in Windows Server 2012 and 2012 R2, the host component).
Federated Hyper-V image level backup of Hyper-V clusters	Includes VMs which are stored on clustered shared volumes.
Hyper-V image level backup of stand-alone Hyper-V servers	Includes all VMs and Initial Store (or, in Windows Server 2012 and 2012 R2, the host component).
Hyper-V backup over SMB 3.0	Includes all VMs that are stored on SMB 3.0 file servers.

## Files included in backups

The Hyper-V VSS Writer reports certain files for each VM during image-level backups with NMM.

The following table lists the reported file type and extension.

**Table 6** VM files supported by the VSS Hyper-V Writer

File type	File extension	Description
Virtual Hard Disk files (Windows Server 2008 and earlier)	.VHD	For VMs created with Windows Server 2008 and earlier, Hyper-V uses the Microsoft Virtual Hard Disk (VHD) specification to store virtual hard disks for VMs. A VM can have one or more virtual disks.
Virtual Hard Disk files (Windows Server 2012 and later)	.VHDX	For VMs created with Windows Server 2012 and later, Hyper-V uses the Microsoft Virtual Hard Disk (VHDX) specification to store virtual hard disks for VMs. A VM can have one or more virtual disks.
VM configuration	.XML	Hyper-V uses a VM configuration file in XML format to store VM settings (for example, CPU, memory, VHDs).
VM Running State files	.BIN .VSV	Hyper-V uses a VM configuration file in XML format to store VM running state (memory) files.
Virtual Hard Disk Differencing files	.AVHD	A VM snapshot creates one differencing VHD file per VM VHD.
VM Configuration Snapshot(s)	.XML	A VM snapshot creates a copy of the current VM configuration and saves it to enable rollback.

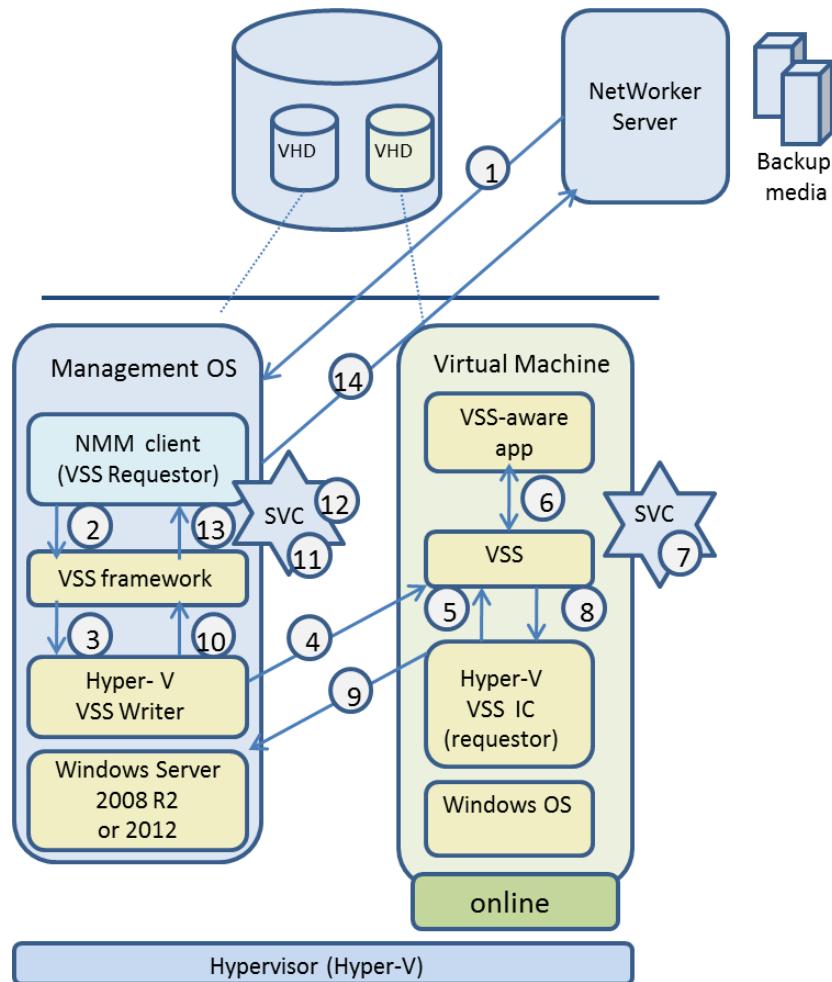
## Backup process workflow

The workflow for Hyper-V backups with NMM differs greatly, depending on the configuration and backup type.

### Image-level VSS backup workflow

The VSS writer enables the creation of image backups for virtual machines by quiescing the Windows operating system and applications within the guest for operating system and application consistency.

The following figure illustrates the workflow for an image-level VSS backup.

**Figure 8** Image-level VSS backup workflow

During an image-level VSS backup, the following events occur:

1. The NetWorker server sends a request to the NMM client on the Hyper-V management operating system to start the backup for the specified VMs.
2. NMM sends a request to the VSS framework to create a point-in-time consistent backup for the VM.
3. The VSS framework contacts the Hyper-V VSS writer and requests that the writer prepare for shadow volume copy (SVC) for the specified VM.
4. The Hyper-V VSS writer establishes a communication path with Hyper-V VSS Integration Components (IC) on the VMs that are being backed up.
5. The Hyper-V VSS IC requests an SVC from VSS inside the VMs.
6. Inside the VM, VSS sends a prepare for SVC request to all applications within the VM. All of the applications are quiesced, and then control is returned to VSS.
7. VSS creates an SVC inside the VM.
8. VSS returns control to the VSS IC requestor.
9. When the SVC completes, the Hyper-V VSS IC requestor notifies the Hyper-V VSS writer and CSV VSS writer (Windows Server 2012 and 2012 R2 only) on the management operating system.
10. The Hyper-V VSS writer returns control to VSS.

11. VSS creates an SVC on the management operating system for the required volumes.
12. The post SVC process occurs to synchronize changes between the SVC created on the VM and the shadow copy created on the physical machine.
13. VSS returns control to NMM.
14. NMM performs a backup from the SVC.  
Ensure that the IC versions for the backup Hyper-V Server and the guest VM match. If you plan to recover the guest VM to an alternate Hyper-V Server, then ensure that the alternate Hyper-V Server uses the same version of IC.

## NMM federated architecture for Hyper-V failover clusters

NMM supports failover clusters through a federated architecture that manages backup and recovery across the Hyper-V cluster. The federated architecture complies with Microsoft CSV and Hyper-V cluster rules. The federated architecture provides failover resiliency for Hyper-V highly available VMs by determining which physical cluster node is running a VM at the time of a backup or recovery operation.

The federated architecture includes two roles:

- Primary role—Provides the point of communication for the NetWorker browse, backup, and recovery jobs. The primary role also controls the high-level VM image backup and recovery workflow across the cluster nodes.
- Secondary role—Provides low-level VM image backup and recovery workflow on specific cluster nodes.

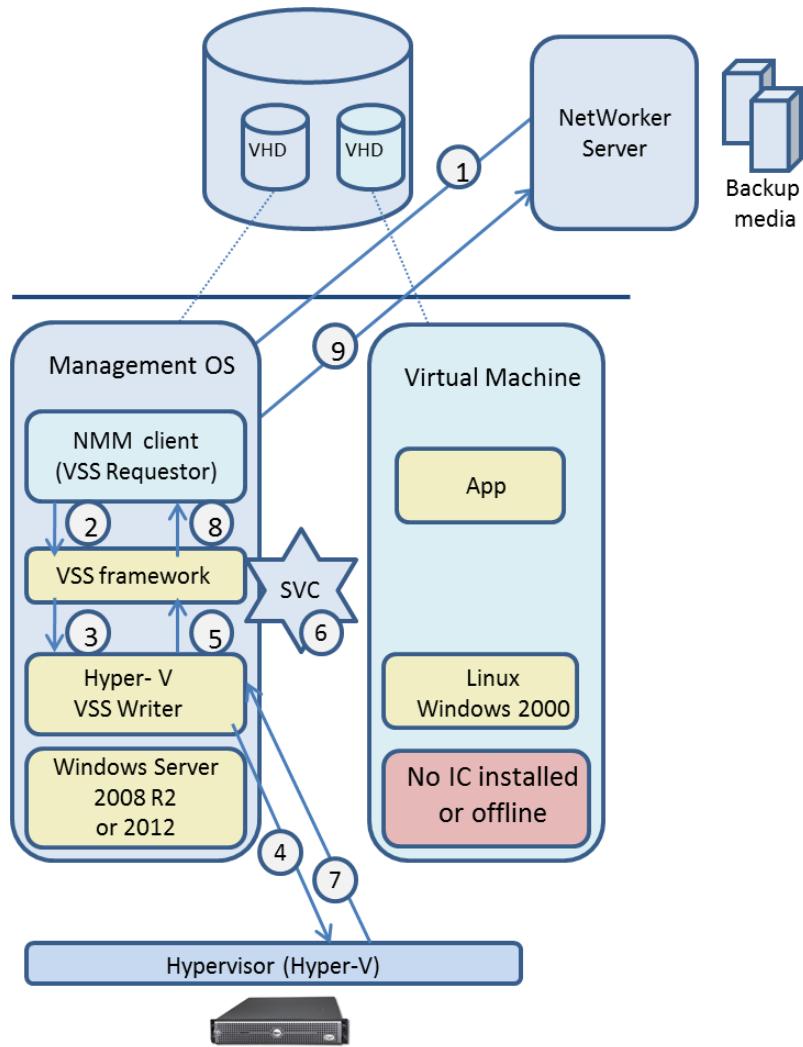
You use the NMC client configuration wizard to configure the NMM Windows CSV client. After you configure the CSV client, NMM issues browse, backup, and recovery jobs against the NMM CSV client name.

The NMM process starts when NetWorker issues a job (workorder) to the NMM Windows CSV client that is operating in the primary role. For VM image backup and recovery operations, the primary role determines which cluster nodes run the VMs specified in the job (workorder), and then dispatches sub-jobs to the appropriate cluster nodes. A dispatched sub-job results in an NMM process starting on the target cluster node, and that process operates in the secondary role. The secondary role manages the CSVs and interacts with the Microsoft Hyper-V VSS writer for backup and recovery operations.

## Image-level saved state backup workflow

If a virtual machine is paused when a backup occurs, then the state changes to a saved state after the backup, which is also called an offline backup.

The following figure illustrates the workflow for an image-level saved state backup.

**Figure 9** Image-level saved state backup workflow

During an image-level saved state backup, the following events occur:

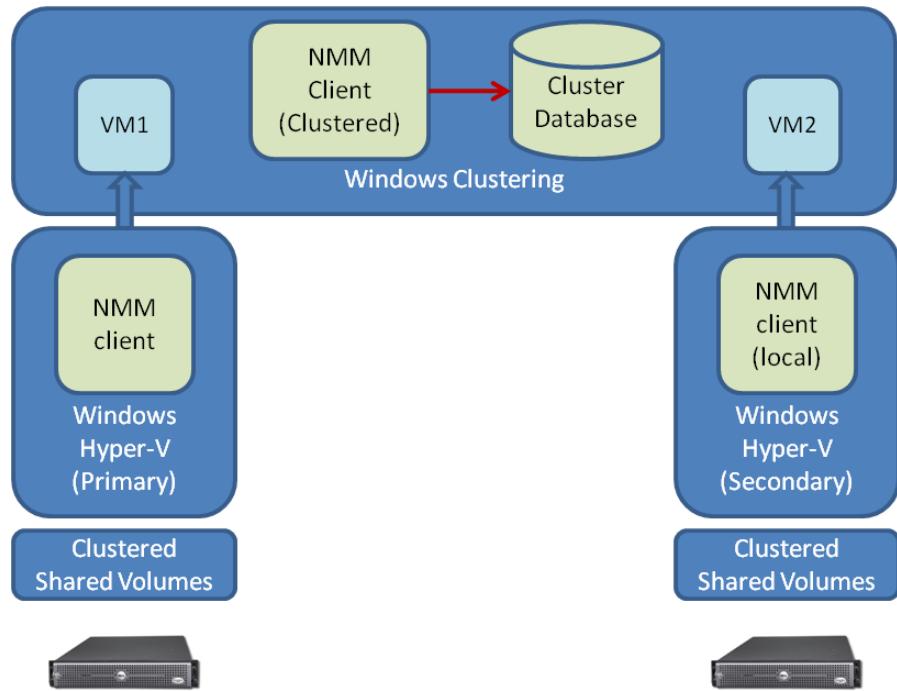
1. The NetWorker server sends a request to the NMM client on the Hyper-V management operating system to start the backup for the specified VMs.
2. NMM sends a request to VSS to create a point-in-time consistent backup for the VM.
3. VSS contacts the Hyper-V VSS writer and requests a prepare for SVC for the specified VM.
4. The Hyper-V VSS writer sends a request to the hypervisor to put the specified VMs into a pause state to freeze the I/O.
5. The Hyper-V VSS writer returns control to VSS.
6. VSS creates an SVC on the management operating system for the required volumes.
7. The Hyper-V writer returns the specified VM to running state.
8. VSS returns control to NMM.
9. NMM performs a backup from the shadow copy volume copy.

## Image-level backup workflow in Hyper-V federated backups

Federated Hyper-V image-level backups of Hyper-V clusters include VMs that are stored on cluster shared volumes (CSVs).

The following figure illustrates the primary and secondary CSV backup workflows in Hyper-V federated backups.

**Figure 10** Image-level backup workflow for federated backups



## Windows Server 2008 R2 and earlier

During a Windows Server 2008 R2 and earlier Hyper-V federated backup, the following events occur:

1. The primary NMM client receives a work order from the NetWorker server. The work order lists the VMs to back up.
2. NMM queries the cluster database to retrieve the active node for each VM.
3. NMM builds the sub-work orders for each node.
4. NMM serializes the backup across all cluster nodes.
5. The primary NMM client creates a final view of the backup based on the partial views received from the secondary clients.
6. The secondary NMM client receives a work order from the primary CSV.
7. The secondary NMM client calls the CSV API to discover the required CSV volume and clear the CSV backup state.
8. The VSS and Hyper-V VSS Writer are called to select the VMs for image backup and to create a shadow copy with the point-in-time copy of the VM files.
9. NMM backs up the VM files from the SCV to the NetWorker server.
10. The secondary NMM client creates a partial view of the backup and forwards it to the primary NMM client.

## Windows Server 2012 and 2012 R2 CSV with no proxy node

During a Windows Server 2012 and 2012 R2 Hyper-V federated backup with no proxy node, the following events occur:

1. The primary NMM client receives the save sets from the NetWorker Server, listing the VMs to backup.
2. The CSV API is called to discover the required CSV volume.
3. The VSS and Hyper-V VSS Writer and CSV VSS Writer are called to select the VMs for the image backup and to create a shadow copy with the point-in-time copy of the VM files.
4. The VM files are backed up from the shadow copy to the NetWorker server.

## Windows Server 2012 and 2012 R2 CSV with proxy node

During a Windows Server 2012 and 2012 R2 Hyper-V federated backup with a proxy node, the following events occur:

1. The primary NMM client receives the save sets from the NetWorker Server, listing the VMs to backup and list of proxy servers.
2. The primary NMM client starts the backup on the proxy sever.
3. The VM files are backed up from the shadow copy to the NetWorker server.

## Recovery overview

The following recovery options are available when you perform regular backups as discussed in [Backup overview on page 27](#):

- Perform a recovery of a VM to its original location on the original Hyper-V Server.
- Perform a redirected recovery of a VM to an alternate Hyper-V server.
- Perform a redirected recovery of a VM on the same Hyper-V server to a different location.
- Perform a GLR of individual files and folders.

## Types of supported recovery

The following table lists the supported types of recoveries.

**Table 7** Types of supported recoveries

Type of recovery	Includes
Disaster recovery for Hyper-V server	Hyper-V Server and role.
Federated Hyper-V image level recovery of Hyper-V clusters	VMs which are stored on clustered shared volumes.
VM recovery	Individual VMs and the Initial Store (or, in Windows Server 2012 and 2012 R2, the host component).
Granular level recovery	Individual files and folders from a VM backup.

## Recovering to the original Hyper-V Server

You might need to recover a VM to its original location on the original Hyper-V Server from which the backup was performed if one certain scenarios occurs.

Recover to the original location when:

- You need to roll back the VM because a patching or virus issue occurred.
- You need to perform disaster recovery of the VM after a disk crash.
- The VM was accidentally deleted.

When you recover a VM to its original location. The recovery process deletes or overwrites all files on the VM, if the host exists.

[Performing Hyper-V recovery to the original machine and location on page 53](#) describes how to recover a VM to the original location.

## Redirected recovery of a VM to an alternate Hyper-V Server

You can recover a VM to an alternate Hyper-V server.

To perform the recovery:

- Select a different Hyper-V server for the VM recovery process.
- Select a different file system location for the files on the original Hyper-V Server.
- In a clustered environment, select the CSV where the files will be placed during a recovery.

[Performing a directed Hyper-V recovery to a different machine or location on page 54](#) describes how to recover a VM to an alternate Hyper-V Server.

## Special character considerations

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

### Supported characters

NMM Hyper-V supports the following characters in VM names and VM configuration paths, including stand-alone, CSV, and SMB 3.0 configurations:

- Alpha numeric (A–Z, a–z, 0–9)
- - . [ ] \_ { } + = ` ~ ! # \$ % ^ & ()
- Space

If a Hyper-V save set contains a VM name or VM configuration path that includes a character not listed above, the backup or recovery might fail. For example, foreign language character sets, such as Japanese or German, are not supported.

### Backups

If a Hyper-V VM name or VM configuration path contains a character not listed above, the CSV backup fails. However, backups of VMs in a stand-alone or SMB configuration succeeds if the VM name or configuration path contains a character not listed above.

### Recovery

Recovery to the original location succeeds if the VM name or configuration path contains a character not listed above. Redirected recoveries to stand-alone or CSV destinations fail when the VM name or VM configuration path contains a character not listed above. NMM does not support redirected recoveries for SMB 3.0.

### Save set names

When specifying a save set name that contains a character not listed above, replace the special character with a URL-encoded value. URL encoding converts non-ASCII and Unicode characters into a format that NMM supports.

The following table lists the most commonly used special characters and their associated URL values.

**Table 8** Common special characters and their URL-encoded values

Special character	URL-encoded value	Special character	URL-encoded value
\	%5C	?	%3F
/	%2F	]	%5D
"	%22	[	%5B
%	%25	}	%7D
#	%23	{	%7B
&	%26	^	%5E
<	%3C	'	%60
>	%3E		%7C

For example, a save set that is named APPLICATIONS:\Microsoft Hyper-V \vmA&B should be specified as APPLICATIONS:\Microsoft Hyper-V\vmA%26B.

# CHAPTER 2

## Backups

This chapter includes the following sections:

- [Planning backups](#).....38
- [Configuring backups](#).....40
- [Configuring a Hyper-V client resource](#).....43

## Planning backups

Before backing up Hyper-V VMs, review the following information:

- [NMM capabilities for Window Server 2012 and 2012 R2 environments on page 38](#)
- [Integration services components version on page 39](#)
- [Viewing valid application data save sets on page 39](#)

## NMM capabilities for Window Server 2012 and 2012 R2 environments

NMM supports Windows Server 2008 R2, 2012, and 2012 R2 Hyper-V stand-alone servers and CSVs. NMM support for Windows Server 2012 and 2012 R2 Hyper-V stand-alone servers is similar to Windows Server 2008 and 2008 R2, and stand-alone server backup and recovery is the same. In Windows Server 2012 and 2012 R2, the Initial Store configuration file is replaced with the Host Component configuration file.

### Backups over SMB 3.0

NMM supports Hyper-V VMs residing on Widows Server 2012 and 2012 R2 SMB 3.0. You back up stand-alone servers and non-CSV failover clusters over SMB the same way you back up local VMs. To back up CSVs over SMB, NMM does not use federated backup architecture. Instead, you configure a backup for each server like a stand-alone Hyper-V server.

### Windows Server 2012 and 2012 R2 Hyper-V CSVs

For Windows Server 2012 and 2012 R2, Microsoft has released several special requirements and special APIs to support 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, and CSVs can be backed up in parallel.

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

In Windows Server 2012 and 2012 R2, the new CSV VSS writer has the capability to report the backup components on the behalf of a remote 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 VM, but also to back up the image located on a remote node. This allows for more configuration options. For example, you can dedicate a single node to back up the cluster.

### Windows Server 2012 and 2012 R2 Hyper-V CSV Continuous Availability

Due to enhancements in Windows Server 2012 and 2012 R2 CSVs, you can back up CSV VMs as part of a highly available (cluster-aware) backup or a physical proxy node backup. The following sections provide an overview of how to configure these backups.

#### Highly Available Backups (cluster-aware backups)

Cluster-aware backups are highly available because you install NMM on each node in the cluster. If one node is not available, NMM initiates 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 a dummy NetWorker client resources 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.

## Windows 2012 and 2012 R2 Scale Out File Server

If a virtual machine is hosted on a Windows 2012 or 2012 R2 Scale Out File Server, then on the standalone machine where the VM runs, install the Failover Clustering feature. The Microsoft documentation provides details.

## Integration services components version

Ensure that the IC version that runs inside the VM is the same as the version of Hyper-V on the host. To determine the version of Hyper-V on the server, start the Hyper-V manager and then select About Hyper-V Manager from the Help menu.

### Procedure

1. In the Device Manager application inside the guest VM, on **System Devices**, select **Device Manager**.
2. Right-click the entry **Hyper-V Volume Shadow Copy**.
3. Select **Properties**.
4. Check the version on the **Driver** tab.
5. If the version does not match the Hyper-V version, insert the integration services disk by choosing that option under the **Action** menu in the VM console.
6. Install the integration components, and then reboot the VM.

## Viewing valid application data save sets

When you configure a client resource, enter the save sets in the Save Set attribute of the client resource.

To display a list of the application data save sets that are available for backup:

### Procedure

1. Open a command prompt on the application server and type the required command.
2. If the Hyper-V Server is a stand-alone host, then type:  
`nsrsnap_vss_save -?`
3. If the Hyper-V Server is configured as a cluster, then type:  
`nsrsnap_vss_save -s networker server -?`

If the application server is on a cluster virtual host, run the command from the physical node that is currently hosting the application server.

Example output:

The following examples show the application data (Hyper-V guest VM) save sets that are available on a Hyper-V system with two VMs, `virtual_machine_name_1` and `virtual_machine_name_2`, on stand-alone and cluster virtual hosts.

Windows Server 2008 and 2008 R2 on a stand-alone virtual host:

`"APPLICATIONS:\Microsoft Hyper-V"`

`"APPLICATIONS:\Microsoft Hyper-V\Initial Store"`

"APPLICATIONS:\Microsoft Hyper-V\virtual\_machine\_name\_1"  
 "APPLICATIONS:\Microsoft Hyper-V\virtual\_machine\_name\_2"  
 Windows Server 2012 and 2012 R2 on a stand-alone virtual host:  
 "APPLICATIONS:\Microsoft Hyper-V"  
 "APPLICATIONS:\Microsoft Hyper-V\Host Component"  
 "APPLICATIONS:\Microsoft Hyper-V\virtual\_machine\_name\_1"  
 "APPLICATIONS:\Microsoft Hyper-V\virtual\_machine\_name\_2"  
 Windows Server 2008 and 2008 R2 on a cluster virtual host:  
 NMM : saveset on node.nmmcsv.com :  
 "APPLICATIONS:\Microsoft Hyper-V"  
 "APPLICATIONS:\Microsoft Hyper-V\Initial Store" nonCSV  
 "APPLICATIONS:\Microsoft Hyper-V\non-csv" nonCSV  
 "APPLICATIONS:\Microsoft Hyper-V\vm27\_rename1" CSV  
 Windows Server 2012 and 2012 R2 on a cluster virtual host:  
 NMM : saveset on node.nmmcsv.com :  
 "APPLICATIONS:\Microsoft Hyper-V"  
 "APPLICATIONS:\Microsoft Hyper-V\Host Component" nonCSV  
 "APPLICATIONS:\Microsoft Hyper-V\non-csv" nonCSV  
 "APPLICATIONS:\Microsoft Hyper-V\vm27\_rename1" CSV  
 Remove the inverted commas when copying the save set name from the output to the save set attribute in the client resource.

4. Press **Enter**.

Each line of output corresponds to a save set entry that you can add to the Save Set attribute of a client resource. Type each entry that you add to the Save Set attribute on a separate line.

## Configuring backups

When configuring backups, the backup tasks differ depending on the items to back up. You must specify the correct save set syntax and Application Information attributes to perform the desired backup type.

The following table describes the backup tasks to perform when you back up Hyper-V parent and VMs.

**Table 9** Backup tasks for Hyper-V

Items to back up	Backup tasks to perform
On the server The Hyper-V role can coexist with other Microsoft applications, such as: <ul style="list-style-type: none"> <li>• SQL Server</li> <li>• SharePoint Server</li> <li>• Exchange Server</li> </ul>	Complete tasks 1 through 7: Task 1: Configure a backup pool Task 2: Configure snapshot policies Task 3: Configure a backup schedule Task 4: Configure a backup group Task 5: Configure a client resource

**Table 9** Backup tasks for Hyper-V (continued)

Items to back up	Backup tasks to perform
<ul style="list-style-type: none"> <li>• Windows Server Cluster</li> </ul>	Task 6: Configure privileges Task 7: Configure a proxy client The <i>NetWorker Module for Microsoft Administration Guide</i> provides details how to perform these tasks.
Hyper-V on the server Hyper-V VMs and Initial Store/ Host Component configuration file	Complete tasks 1 through 5: Task 1: Configure a backup pool Task 2: Configure snapshot policies Task 3: Configure a backup schedule Task 4: Configure a backup group Task 5: <a href="#">Configuring a Hyper-V client resource on page 43</a> The <i>NetWorker Module for Microsoft Administration Guide</i> provides details about tasks 1 through 4.
Hyper-V VM applications Microsoft application data, such as: <ul style="list-style-type: none"> <li>• SQL Server</li> <li>• SharePoint Server</li> <li>• Exchange Server</li> <li>• Windows Server Cluster</li> </ul>	Install NMM on the VM operating system and configure application backups with NMM installed within the VM operating system. Specific instructions for the Microsoft application are provided in the following: <ul style="list-style-type: none"> <li>• Configure Windows application backups.</li> <li>• Configure Windows Server cluster backups.</li> </ul> The <i>NetWorker Module for Microsoft Administration Guide</i> and user guides provide specific instructions about how to back up each application.

Hyper-V VM snapshots are not related to NMM or NetWorker snapshots. Hyper-V VM snapshots are created, viewed, and applied to the VM through Hyper-V Manager. When NMM backs up a Hyper-V VM, the Hyper-V VM snapshots are part of that backup.

The following table lists the Hyper-V save set syntax to specify for supported types of Hyper-V backup.

**Table 10** Hyper-V save set syntax

Type of backup data	Save set syntax
Hyper-V Manager The Hyper-V configuration file and each VM.	APPLICATIONS:\Microsoft Hyper-V The Hyper-V Writer does not support offline backup of the configuration file. You cannot use the APPLICATIONS:\Microsoft Hyper-V save set in a proxy backup group.
Hyper-V configuration file (Initial Store/Host Component) There is one configuration file in the Hyper-V Manager installation. This lists the Hyper-V settings for the host operating system and the guest operating systems.	For Windows Server 2008 and 2008 R2: APPLICATIONS:\Microsoft Hyper-V\Initial Store For Windows Server 2012 and 2012 R2: APPLICATIONS:\Microsoft Hyper-V\Host Component The Hyper-V Writer does not support offline backup of the configuration file. You cannot use the APPLICATIONS:

**Table 10** Hyper-V save set syntax (continued)

Type of backup data	Save set syntax
	\Microsoft Hyper-V\Initial Store and APPLICATIONS:\Microsoft Hyper-V\Host Components save sets in a proxy backup group.
Hyper-V VM There are usually multiple VMs on the host operating system.	APPLICATIONS:\Microsoft Hyper-V\virtual_machine_name Child pertains or VMs can be included in a proxy backup group.

The following table lists the variables that can be specified in the Application Information attribute of the client resource.

**Table 11** Hyper-V application information variable settings

Attribute name	Description	Values
NSR_SNAP_TYPE= <i>value</i>	Specifies the snapshot service provider name.	vss This value is required.
NSR_FEDERATED_BACKUP	Marks the backup for CSV recovery.	Yes
NSR_FEDERATED_PSOL	Optional. Distributes the backup workload across all 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.	Type a comma-separated list of the server names. For example: NSR_FEDERATED_PSOL=server1, server2, server3
NSR_EXCLUDE_SMB	Optional. Excludes VMs that have data stored on SMB file servers. By default, SMB VMs are included in the writer level backup.	Yes
NSR_VSS_FULL_BACKUP	Performs VSS FULL or VSS COPY backups in guest VMs. By default, all Hyper-V VM backups are VSS COPY type.	Yes

**Table 11** Hyper-V application information variable settings (continued)

Attribute name	Description	Values
NSR_VSS_FORCE_SYSTEM_PROVIDER	Required for backups that are hosted over SMB-3.	No
NSR_EXCLUDE_COMPONENTS	Optional. Excludes a VM from the backup. Specify the writer level saveset and the components to exclude from the backup. NMM logs the excluded components in the NMM.raw log file for references.	Type a comma-separated list of the server names. For example: NSR_EXCLUDE_COMPONENTS=VM1, VM2, VM3

## Configuring a Hyper-V client resource

A client resource specifies what to include in a snapshot of an NMM client. Client resources are associated with other backup resources, such as groups and snapshot policies. You can create a client resource by using the Client Configuration Wizard or manually in the NetWorker Management Console.

You can create multiple client resources for the same NMM host. In this way, you can apply different backup attributes to different types of information on the same host.

### Creating the client by using the Client Configuration Wizard

The NMM Configuration Wizard for Hyper-V simplifies configuration of scheduled backups for NMM clients for Hyper-V servers.

The NMM Configuration Wizard for Hyper-V has the following properties:

- The wizard automatically configures Hyper-V save sets, backup commands, Application Information attributes, and backup options.
- The wizard does not require you to have privileges on the local machine root (on the host where the GUI runs) or Administrator privileges.
- You can use the wizard to configure client resources on Windows Server 2008, 2008 R2, 2012, and 2012 R2 platforms, both for stand-alone and federated environments.

#### Procedure

1. In the **Administration** window of the NetWorker Management Console:
2. Click **Configuration > New Client Wizard**.

The Specify the Client Name and Type page appears.

3. In the **Client Name** field, type the name of the client where NMM is installed:
  - For federated backups, type the cluster server name.

- For non-federated backups, type the node name.
4. Select the **Traditional NetWorker Client** option.
  5. Click **Next**.

The Select the Backup Configuration Type page appears. On the Select the Backup Configuration Type page, the wizard automatically detects the applications that are installed on the client specified on the Specify the Client Name and Type page and displays the list from which you can select the backup type. Because Hyper-V is installed, the backup option Microsoft Hyper-V Server appears in Specify the Client Backup Options Type page.

6. Select **Hyper-V Server** and click **Next**.

The Specify the Client Backup Options page appears.

7. To use the Client Direct feature, select the **Client Direct** checkbox.
8. Leave the **Target Pool** field blank.
9. Select one of the following deduplication options:

- **None** — If you have not set up data deduplication.
- **Data Domain backup** — If you are using a Data Domain device for data deduplication.
- **Avamar deduplication backup** — If you are using an Avamar device for data deduplication. After you select this option, select the Avamar node from the available list.

10. Click **Next**.

The Select the Hyper-V Backup Objects page appears. The Select the Hyper-V Backup Objects page automatically discovers the save sets depending on the operating system type and entered client name.

11. Select the save sets:

- **For federated setups** — Select the CSV save sets that are registered as cluster resources. If you select both registered and un-registered CSV VMs, then the backup will fail.  
The wizard creates dummy client resources for all nodes that are not used to perform the backup. The wizard also creates a client resource with the cluster name and specifies the NSR\_FEDERATED\_BACKUP=yes attribute.
- **For stand-alone setups** — Select the non-CSV save sets to back up. The wizard creates a client resource with the physical name using the selected save sets.  
If you do not select any save sets, then the wizard applies the Hyper-V writer level save set.  
If you select both CSV and non-CSV VMs, the backup will fail.

12. Click **Next**.

The Specify Backup Options page appears.

When you use the wizard to configure backups for Windows Server 2012 and 2012 R2 federated setups, the wizard includes preferred server order list (PSOL) options. The PSOL attribute allows you to back up the nodes, specify the attribute NSR\_FEDERATED\_PSOL and type a comma-separated list of user selected node names. For example: NSR\_FEDERATED\_PSOL=node1, node2, node4, node3.

13. In the **Specify Backup Options** page:

- For Windows Server 2008 and 2008 R2 setups, if a VSS hardware provider is installed on the host, then select **Use a data mover for the backup** and from the **Data mover name** dropdown box, select the host name of the data mover.
- For Windows Server 2012 and 2012 R2, under **Select VSS backup type**, if you would like to perform full backups, select **Force VSS FULL backup type** to specify the VSS FULL backup type. By default, this box is unchecked and the backup type is VSS COPY.
- For Windows Server 2012 and 2012 R2 setups, in the **Select proxy servers** section, you can specify the order in which to back up servers. Select a server from in **Available Servers** list and then click the right arrow to move it to the **Proxy Servers** list. To adjust the order of servers, select a server in the **Proxy Servers** list and then click the up and down arrows.

14.Click **Next**.

The Select the Client Properties page appears.

15.Select the browse policy, retention policy, and backup schedule for the backup.

The *NetWorker Administration Guide* describes how to configure these resources.

16.(Optional) Type a comment for the client in the **Client comment** field.

17.In the **Remote access** field, type the user account of the administering machine in the format of `username@hostname`.

18.Click **Next**.

The Specify the NetWorker Backup Group page appears.

19.To choose a group from the existing list, select **Add to existing group**. Only groups without snapshot policies are available for selection.

20.To create a new group, select **Create a new group** and enter the following details:

- Type the group name in the **Group Name** field.
- Select the number of client retries.
- Set the schedule backup time in the **Schedule Backup Start Time** field.
- Select **Automatically start backup at the scheduled time** to start the backup automatically at the designated time.
- Select **Snapshot Pool** and set the **Interval** and **Restart Window**.

21.Click **Next**:

- If you selected **Create a new group**, then the **Specify the Snapshot Policy** page appears.
- If the setup includes a storage node, then the **Specify the Storage Node Options** page appears. Changing the storage node option changes the configuration for this setup.

22.Select the **Add to an existing snapshot policy** option to choose a snapshot policy from the existing list and click **Next**.

23.Select the **Create a new snapshot policy** option to enter the snapshot details and click **Next**:

- **Name**—Type a unique snapshot policy name.
- **Number of Snapshots**—Select the number of snapshots that NetWorker creates and maintains for the Hyper-V backup.
- **Retain Snapshots**—Select the number of snapshots that NetWorker retains.

- **Snapshot Expiration**—Select the length of time NetWorker keeps snapshot entries in the media database and client file index on the NetWorker server. The choices are defined by the set of existing management policies.
- **Backup Snapshots**—Select which snapshots to back up.

The Specify the Storage Node Options page appears.

24. Select either of the following options and click **Next**:

- **Backup to the NetWorker server only**—When the setup does not include a NetWorker storage node.
- **Backup to the following storage nodes**—To select the NetWorker storage node name and other details.

The Backup Configuration Summary page appears.

25. Click **Next**.

The Backup Configuration Summary page appears.

26. Do one of the following:

- Click **Back** to revisit the previous pages.
- Click **Create** to configure the client resources.

The Client Configuration Results page appears with details about the client resources that have been created for an Exchange stand-alone setup, or several client resources that have been created for an Exchange clustered setup.

27. To verify the details for the client, select the client and view the **Client Properties** page in the **NetWorker Management Console**.

28. To make changes to the configuration that you created earlier, click **Modify Client Wizard**.

## Creating the client manually by using the NetWorker Management Console

You can manually create a Hyper-V client resource by using the NetWorker Management Console.

### Procedure

1. On the **Administration** page of the **NetWorker Management Console**, click **Configuration**.
2. In the expanded left pane, select **Clients**.
3. From the **File** menu, select **New**.
4. In the **Name** attribute, type the host name of the NetWorker client computer.
5. In the **Comment** attribute, type a description.

If you are creating multiple client resources for the same NetWorker client host computer, then use this attribute to differentiate the purpose of each resource.

6. From the **Browse Policy** attribute, select a browse policy from the list.

The browse policy determines the time period during which the rolled-over data is available for quick access.

7. From the **Retention Policy** attribute, select a retention policy from the list.

The retention policy determines the time period during which the rolled-over data is available, although not necessarily quickly.

8. Select the **Scheduled Backups** attribute.

9. In the **Save Set** attribute, specify the components to be backed up. Place multiple entries on separate lines. [Table 9 on page 29 on page 41](#) provides the save set syntax.
10. For the **Group** attribute, select the backup group to which this client resource will be added.  
If you add client resources for the same NMM host to different backup groups, then ensure that the Start Time attribute for each backup group is spaced such that the backups for the host's client resources do not overlap.
11. For the **Schedule** attribute, select a backup schedule.
12. Click the **Apps & Modules** tab.
13. In the **Access** area, type the domain administrator username and password.  
For guest VMs hosted over SMB 3.0, the backup fails if you do not provide the domain administrator credentials.
14. In the **Backup command** attribute, type the following backup command:  
`nsrsnap_vss_save.exe`
15. In the **Application Information** attribute, specify the attributes for the backup. [Table 10 on page 30 on page 42](#) lists the application information attributes and settings.
16. Click the **Globals (1 of 2)** tab.
17. In the **Aliases** attribute, type the NETBIOS name for the client.  
NMM client uses the host machine NETBIOS or "short" name when connecting to the NetWorker server to browse backups. If the NETBIOS name is not found, NMM will not be able to display backups.
18. Click the **Globals (2 of 2)** tab.
19. If you are setting up a proxy client for the NMM client, type the host name of the proxy client in the **Remote Access** attribute.  
If the NMM client is part of a cluster, type the names of the physical nodes of the cluster in the Remote Access attribute.
20. Click **OK**.

## Performing cluster-level and CSV VM backups

NetWorker performs CSV VM backups through a client resource created for the cluster virtual server only. You create client resources for all of the nodes in the cluster and for the cluster server. However, the backup will be scheduled against the cluster virtual server client resource only. NetWorker indexes the backup against the cluster server name.

When you perform a backup on a specific cluster node, the backup requestor can only include the clustered VMs that are currently active on that node. This is to ensure that for each cluster node, only VMs active on that node are backed up and that the host is the coordinating node. In addition, the cluster node where the shadow copy is being created must be the coordinating node for the CSV. Therefore, backup of VMs on the same CSV but on different nodes must be serialized.

NMM supports backups for all CSV VMs as well as individual CSV VMs. [Configuring backups on page 40](#) provides details about how to perform Hyper-V backups.

**NOTICE**

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.

Hyper-V CSV supports only a RolloverOnly backup policy or a retention=0 snapshot management policy. The Hyper-V CSV backup is a federated backup, where one Hyper-V server spawns jobs on other Hyper-V servers in cluster environment to perform the backups of all the VMs residing on CSV volumes.

## Performing a CSV-level backup

To perform a CSV-level backup:

- Set the save set attribute “Applications: \Microsoft Hyper-V\” in the client resource for the cluster virtual server to back up all the CSV VMs in the cluster.
- Add NSR\_FEDERATED\_BACKUP=yes to the Application Information attribute in the client resource.

## Performing a CSV VM backup

To perform a CSV VM backup, set the save set attribute **APPLICATIONS:\Microsoft Hyper-V\** in the client resource for the cluster virtual server to back up the CSV VM in the cluster.

NMM supports backup of non-clustered VMs that run on specific cluster nodes. NMM excludes VMs that do not reside on the CSV from the CSV backup. Backup and recovery of non-clustered VMs is managed through the individual physical node name, not the cluster virtual server client resource. The physical node name is the client resource name.

The number of save sessions depends on the number of Hyper-V VMs the volume contains, not on whether the VMs is hosted on single CSV volume or on multiple CSV volumes.

If you configure multiple VMs on a writer level save set, the maximum number of save streams depends on PowerSnap settings.

## Configuring multi-proxy backups

To meet backup windows for 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, you add secondary roles, which must be physical 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 alias. The recovery process for VMs backed up as part of a multi-proxy setup is the same as the recovery process for traditional backups.

In a multi-proxy architecture, 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 and change CSV ownerships to the proxy nodes you select. The backup load is evenly split between multiple nodes. All proxy nodes will perform backups in parallel, increasing backup performance by 60-70% compared with single proxy backups in a normal distributed CSV environment.

CSV ownership can change seamlessly between nodes. The performance of a running VM is not impacted when the underlying CSV ownership changes from one node to another node.

You can add or remove proxy nodes as needed.

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. 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.
- CSV ownership distribution algorithm — A CSV algorithm is used to intelligently reassigned and change CSV ownerships to the proxy nodes you select. NMM evenly distributes the backup load across the proxy nodes, and the proxy nodes perform backups in parallel to maximize the backup performance.
- Application Information attributes — To enable multi proxy backups, 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, server3
```

where *server1*, *server2*, and *server3* will act as proxy servers.

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**—Allows or disallows CSV ownership change during multi-proxy backups for optimal 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. To attain optimal backup performance, set this attribute value to Yes. If you set this attribute value to No, then NMM does not change the CSV ownership.

## 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 optimal backup performance.

## **Load balance VM data on CSVs**

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

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

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

## **Allocate the number of proxy nodes and CSVs**

When determining how many number of 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 optimal performance, the number of CSVs should be multiples of the number of proxy nodes. Each VM 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 will allow the NMM CSV algorithm to intelligently change CSV ownership to the proxy nodes you selected, thereby properly load balancing the backup data. NMM performs data split per CSV, not per VM.

## **Select proxy nodes with good system resources**

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

# CHAPTER 3

## Recoveries

This chapter includes the following sections:

• <a href="#">Overview</a> .....	52
• <a href="#">Turning the VM offline for recovery</a> .....	52
• <a href="#">Hosting the recovered virtual systems</a> .....	52
• <a href="#">Specifying the destinations for the Hyper-V configuration files and virtual system</a> .....	52
• <a href="#">Selecting Hyper-V recovery options</a> .....	53
• <a href="#">Selecting the Hyper-V recovery destination</a> .....	53
• <a href="#">Recovering Hyper-V CSV VMs</a> .....	56

## Overview

Depending on what you specified in the backup save set, you can recover the following from a NMM Hyper-V VM backup:

- All of the Hyper-V components
- The Initial Store/Host Component configuration file
- Individual or multiple VMs
- Individual files and folders (Granular level recoveries provides details)

## Turning the VM offline for recovery

Before you start the recovery operation, take the VM offline. If the VM is online when you start the recovery, the Hyper-V Writer will turn off the VM.

Once the VM is offline, the recovery process destroys the current VM, recovers the recovery version, and registers the VM.

The Hyper-V Writer automatically detects whether the VM is online, and turns it off. No action is required by you.

## Hosting the recovered virtual systems

The Destination Host page allows you to recover the virtual system to the original location.

### Procedure

1. Select **Recover (Overwrite) Virtual System to original location.**

The Finish button is available and Next button is unavailable when you select this option.

2. Click **Finish**.

The Finish button validates the server location and displays the Hyper-V Recovery Options page.

## Specifying the destinations for the Hyper-V configuration files and virtual system

The Destination Host page allows you to specify the destination directory for the Hyper-V configuration files, and the destination host for each virtual system.

### Procedure

1. In the **Destination for Hyper-V configuration files** dialog box, click **Browse** to change the destination path.

The Select Virtual System Destination list displays the destinations for each virtual system VHD. The specified host must have the NMM client installed.

2. To change a destination host:
  - a. Select a virtual system.
  - b. Click **Change Destination**.

The Remote Directory Browser dialog box appears.

3. Click **Finish** to validate the settings.

If the destinations are valid, then the Hyper-V Recovery Options Summary dialog box appears.

## Selecting Hyper-V recovery options

The Recovery Options Summary page lists the Hyper-V Recovery and NetWorker Recovery options. This allows you to review the settings before you start the recovery.

The Hyper-V Recovery Options page appears when you click Recover in the Hyper-V Recover Session view or at the end of the Hyper-V Recovery wizard. This page displays the settings specified in the Destination Host Selection page and Destination Path page.

To change the Hyper-V recovery options:

- If you have reached this page through the Hyper-V Recover Session toolbar:
  1. Click **Cancel > Advanced Options**.  
The Hyper-V Recovery Options wizard starts.
  2. In the **Hyper-V Recovery Options** wizard, click **Back**.
- If you have reached this page through the Hyper-V Recover Session view, click **Advanced Options** on the Hyper-V Recover Session toolbar.

To change the NetWorker recover options, do one of the following:

- Click Recover Options.
- In the Hyper-V Recover Session view, click Recover Options on the Hyper-V Recover Session toolbar.

The NetWorker Recover Options settings are specified on the General, NetWorker, and Security tabs. To validate all pages, click Start Recover. If all pages are valid, the wizard closes and recovery starts.

## Selecting the Hyper-V recovery destination

You can select the destination of the Hyper-V recovery to the original machine or to a different machine or location.

When you perform a Hyper-V recovery on Windows Server Core 2008, 2008 R2, 2012, or 2012 R2 the recovery is a directed recovery. Because the Server Core installation does not provide a GUI, you must use another machine to start the recovery. To recover a Windows Server 2012 and 2012 R2 Server Core VM through directed recovery, you must use the No Proxy option.

[Restrictions and requirements for relocating and recovering data to a different location on page 96](#) provides details about other restrictions.

## Performing Hyper-V recovery to the original machine and location

You might need to recover VMs to their original location if the VMs have been corrupted or deleted.

When you recover VMs to the original Hyper-V server, the original drive letters or mount points for the VMs must exist on the system, and the directory paths are automatically created. If any of the files from the VMs are still on the Hyper-V Server, then the recovery deletes or overwrites the files.

If the recovered VM was a clustered VM, then the recovery creates the VM on one of the cluster nodes. However, you must use Microsoft Cluster Manager to make the VM highly available.

Because Hyper-V recognizes VMs by an internal GUID, you can not move or rename the current VM during the recover if the VM exists on the Hyper-V Server.

The Initial Store file (Windows Serer 2008 and 2008 R2) and the Host Component file (Windows Server 2012 and 2012 R2) contain the authorization configuration for Hyper-V. You might need to recover the Initial Store/Host Component to the original Hyper-V Server if the file has become corrupted or you need to roll back the authorization settings. The NMM system state backups also include the Initial Store/Host Component.

#### **Procedure**

1. Open the NMM client software and select the **NetWorker Server** on which the NMM client software was configured for backup.

If NMM client is part of a cluster, select the physical client to which you are recovering data. The physical client can be selected from the Client list attribute in the application toolbar.

2. From the left pane, select **Recover > Hyper-V Recover Session**.
3. From the navigation tree, select the Hyper-V Writer or individual VMs under the Hyper-V Writer.
4. From the **Hyper-V Recover Session** toolbar, click **Recover**.

The Hyper-V Recover Session Summary dialog box appears.

5. If all the options look correct, then go to [step 11 on page 54](#). If you want to change options, then go to [step 7 on page 54](#).
6. Click **Recover Options**.
7. On the **General** tab, specify the **Diagnostic Output Level**.
8. On the **NetWorker** tab, specify the **Restore Type**, and then select or clear **Terminate recover of item, if errors are encountered**.

To recover Hyper-V components, select the Conventional Restore recovery type.

9. On the **Security tab**, specify **pass phrases** if any are needed. For GLR recovery, mount the VM.

10. Click **OK** to close the **Hyper-V Recover Session Options** dialog box.

11. Click **Start Recover**.

NMM validates all pages. If all pages are valid, then the recovery begins.

## **Performing a directed Hyper-V recovery to a different machine or location**

You can recover a VM to the original Hyper-V Server, but move the VM files to different file system locations. This type of recovery is necessary if the VM files were moved after the selected backup time and you want to preserve the new locations. If the original VM is present, then the VM is overwritten during the recovery.

#### **Procedure**

1. Open the NMM client software and select the **NetWorker Server** on which the NMM client software was configured for backup.

If NMM client is part of a cluster, select the physical client to which you are recovering data. The physical client can be selected from the Client list attribute in the application toolbar.

2. From the left pane, select **Recover > Hyper-V Recover Session**.
3. From the navigation tree, select the Microsoft Hyper-V Writer, or the individual VMs under the Microsoft Hyper-V Writer.
4. From the **Hyper-V Recover Session** toolbar, click **Advanced Recovery**.  
The Hyper-V Recovery wizard starts and the Destination Host page appears.
5. Specify the destination host server for the Virtual System recover:
  - To recover to the same location as the original:
    - a. Select **Recover (Overwrite) Virtual System to original location**.
    - b. Perform the steps in “[To perform validation and start recovery:](#)” on page 43.
  - To recover to a different path on the same Hyper-V Server:
    - a. Select **Recover Virtual System to a different path** and then click **Next**.  
The Destination Path page appears, and you can specify a destination for each Virtual System. The specified host must have the NMM client installed.  
If the VM is online or active, you must recover the VM to the same node on which it is active.
    - b. Click **Browse** to specify the destination location for the configuration files. After you select the destination location for the configuration files, you can change the destination location for the VM’s virtual disks.
    - c. To change the destination location for a virtual disk, select the VM’s virtual disk in the list, and then click **Change Destination**. Repeat as needed for each virtual disk destination that you want to change.
    - d. When you have completed changing destinations, go to “[To perform validation and start recovery:](#)” on page 43.  
The destinations provided on this page are Microsoft’s default configuration file locations and might not match your Hyper-V configuration. Change the destination as needed.  
When you attempt a directed recovery to a different path on the same Hyper-V server, the recovery process takes the recovery VM offline and then recovers the data to the alternate location. NMM registers the VM pointing to the data in the new location.  
You can remove the Virtual Hard Disk (VHD) files of the original VM manually after the recovered Hyper-V child components are up and running.
- To recover to a different Hyper-V server:
  - a. Click **Recover Virtual System to a different Hyper-V Server**.
  - b. From the **Select Remote Host** list, click the server you want to recover to.
  - c. Click **Next**.
  - d. The Destination Path page appears, where you can specify a destination for each Virtual System.
  - e. On the Destination Path page, click **Browse** to specify the destination location for the configuration files. After you select the destination location for the configuration files, you can change the destination location for the VM’s virtual disks.
  - f. To change the destination location for a virtual disk, select the VM’s virtual disk in the list, and then click **Change Destination**. Repeat as needed for each virtual disk destination that you want to change.
  - g. When you have completed changing destinations, go to “[To perform validation and start recovery:](#)” on page 43.

The destinations provided on this page are Microsoft's default configuration file locations and might not match your Hyper-V configuration. Change the destination as needed.

When you attempt a directed recovery to a different Hyper-V Server and the destination Hyper-V Server has an existing VM with the same name as the one being recovered, the recovery process takes the VM offline and recovers the data to the alternate location. NMM registers the VM pointing to the data in the new location.

You can remove the VHD files of the original VM present earlier with the same name manually after the recovered Hyper-V child components are up and running.

6. To perform validation and start recovery, click **Finish**.

NMM performs validation:

- If the validation is not successful, then NMM displays an error message.
- If the validation is successful, then NMM displays a summary page that lists the Hyper-V and NetWorker Recover options that you specified.
- If you need to change any of the options, click the **Recover Options** or **Back** button.

7. Click **Start Recover**.

NMM validates all pages. If all pages are valid, NMM starts the recovery.

## Recovering Hyper-V CSV VMs

The following sections describe the supported types of Hyper-V CSV recoveries.

### SMB 3.0 VM recovery

When you recover Hyper-V VMs over SMB 3.0, use the same steps that you use for other recovery types. Note that NMM does not support redirected recovery of VMs over SMB 3.0.

### CSV VM recovery

NMM supports recoveries for Hyper-V CSV VMs at the cluster level and at the individual CSV VM level. NMM supports recoveries only on supported NMM hosts that run the Hyper-V service. This host might be outside of the cluster.

NMM recovers the CSV VMs that you select on the cluster node where the VM is currently active. If a CSV VM does not exist at the time of recovery, then NMM recovers the VM to the majority node.

If no VMs exist in the cluster, then for the deleted VM:

- If the winclient started within the cluster, then the recovery operation starts on the node that runs the winclient.
- If the winclient is started from outside cluster, then NMM recovers the deleted VM to the cluster owner node.

Relocation recovery of a clustered VM to a node where it is not active is not allowed on the cluster where the VM resides. If you request a relocated recovery of a VM to a node on a cluster, but the VM is already active on another cluster node, then the recovery will fail. EMC recommends that you move the VM to the desired node first and then initiate the VM recovery.

When you select a single VM, NMM supports the following types of recoveries:

- Default recovery (Recover on the cluster node where the VM is active) — This is the default recovery when you click “Recover...” NMM recovers the selected VM to the cluster node where it is active. If the VM does not exist, NMM recovers it to the majority node where most of the VMs are active.
- Advanced recover — When you click the “Advanced Recover...” option, you can select one of the following options:
  - Recover Virtual System to active virtual system cluster node  
If the VM is already active on the destination cluster node, then you can recover the VM to a path that differs from the path at the time of the backup.
  - Recover Virtual System to a different cluster node  
You can recover the selected VM to an alternate node in the cluster if the VM does not exist in the cluster. If the VM exists in the cluster but is not active on the destination node, then the recovery fails. In this case, first migrate the VM to the destination node and then perform either a default recovery or an advanced recovery.
  - Recover Virtual System to a different Hyper-V server  
You can recover the selected VM to Hyper-V server outside of the cluster.

You can choose to recover all CSV VMs or multiple CSV VMs. NMM cannot recover the selected CSV VMs on their respective cluster nodes when the VM is active. If a CSV VM does not exist at the time of recovery, NMM recovers the VM to the majority node.

After you recover a VM, confirm that the recovery process registers the VM. If the VM is not registered as a cluster resource, use Failover Cluster Manager to register the VM as cluster resource.

## Recovering multiple CSV VMs to the original location

You can recover multiple clustered Hyper-V VMs to the same location as the original CSV VM.

### Procedure

1. Open the NMM client software.
2. From the left pane, select **Recover > Hyper-V Recover Session**.
3. From the navigation tree, select **Image Recovery**.
4. Select the cluster name from the **Client** menu.
5. Select the VMs you want to recover.
6. Click **Recover**.
7. Click **Start Recover**.

NMM validates the information on all pages. If all pages are valid, then NMM starts the recovery.

8. Perform the steps in [“To perform validation and start recovery:” on page 43](#).

## Recovering an individual CSV Hyper-V VM to a different location

You can recover an individual CSV VM to a different location.

### Procedure

1. Open the NMM client software.
2. From the left pane, select **Recover > Hyper-V Recover Session**.
3. From the navigation tree, select **Image Recovery**.

4. Select the cluster name from the **Client** menu.

5. Click **Advanced Recover**.

The Hyper-V Advance Recovery wizard starts and the Destination Host page appears. Select the destination to recover the CSV VM.

- To recover the VM to the node on which it is currently active:
  - a. Click **Recover Virtual System to Active Virtual System Cluster Node**.
  - b. Click **Next**. A dialog box displays the following warning: **Proxy based recovery outside Cluster is Not Supported**.
  - c. Click **OK**.
  - d. Click **Recover Options**.
  - e. Select the **CSV Proxy Server** tab.
  - f. Click **No Proxy**.
- To recover the VM to a different node:
  - a. Click **Restore Virtual System to a Different Cluster Node**.
  - b. Select the remote host to which you will recover the CSV VM.
  - c. Click **Next**.
  - d. Specify the destination path for the configuration files.

This option is only available if the selected VM does not exist in the cluster. To recover a VM that exists on the cluster but is not active on the destination node, you must first migrate the VM to the destination node and then perform either a default recovery or an advanced recovery.

- To recover the VM to a different Hyper-V Server:
  - a. Click **Recover Virtual System to a Different Hyper-V Server**.
  - b. Select the remote host to which you will recover the CSV VM.
  - c. Click **Next**.
  - d. Specify the destination path for the configuration files.
  - e. To perform validation, click **Finish**.

NMM performs the validation. If the validation is not successful, then NMM displays an error message. If the validation is successful, then NMM displays a Summary page that lists the specified Hyper-V and NetWorker Recover options.

6. If you need to change any of the options, click the **Recover Options** or **Back** button.
7. Click **Start Recover**.

NMM validates the information on all pages. If all pages are valid, then NMM starts the recovery.

## Recovering with a Windows Server 2012 and 2012 R2 proxy CSV server

You can select a proxy node for the CSV recovery operation.

### Procedure

1. Open the NMM client software.
2. From the left pane, select **Recover > Hyper-V Recover Session**.
3. Click the **CSV Proxy Server** tab and select a proxy server recovery option.

4. For stand-alone setups, select **No Proxy**.
5. For federated setups, select one of the following options:
  - **Local Server or Current Host Server**—Select this option only for nodes configured using the Windows Server 2012 and 2012 R2 GUI.  
When the NMM GUI runs on a server in the cluster, NMM treats the local host as the default server. Otherwise, NMM treats the current host server, the server to which cluster server name resolves, as the default server for recovery operations.
  - **Choose a Server**—Select this option for nodes configured using the Windows Server 2012 and 2012 R2 GUI or nodes configured using Windows Server 2012 and 2012 R2 core.  
Select a server from the drop-down box to choose a proxy for recovery operations. You can select any server in the cluster.
  - **No Proxy**—Select this option for all Windows Server 2012 and 2012 R2 installations.  
No proxy server will be used.
6. Click **OK**.
7. Click **Recover**.
8. Click **Start Recover**.
9. Perform the steps in “[To perform validation and start recovery:](#)” on page 43.

## Troubleshooting RPC service failure messages

The following section describes how to troubleshoot RPC service failure messages.

### Error connecting to RPC server

The following RPC failure message appears: Agent connection failed: Can't connect to RPC server on *NMM\_client\_name* where *NMM\_client\_name* is the NMM client name.

#### Solution

Perform the following steps:

1. In **User Environment Variables**, create two environment variables and type a value to specify the number of seconds for each variable:
  - **NSR\_CSC\_RPC\_TIMEOUT**—The typical value is 6 seconds.
  - **NSR\_CSC\_METHOD\_TIMEOUT**—The typical value is 180 seconds.
2. Restart the host where you launched the NMM recovery GUI.

### Unauthorized client

When you perform a directed recovery, the following error appears: Agent connection failed: Request received from unauthorized client *NMM\_client\_name* where *NMM\_client\_name* is the NMM client name.

#### Solution

Manually stop the nsrscsd.exe process on the destination host and then perform the directed recovery.



# CHAPTER 4

## Granular Level Recoveries

This chapter includes the following sections:

- [Overview](#)..... 62
- [Recovering Hyper-V files and folders](#)..... 62

## Overview

Granular level recovery (GLR) provides the ability to recover specific files from a single backup without recovering the full VM, reducing the recovery time.

NMM can perform a GLR for backups of Hyper-V VMs that you created with NMM 2.4 or above. You can only use GLR for VMs running Windows operating systems.

You perform Hyper-V GLR by using the NMM GLR feature. The NMM GLR feature uses NetWorker Virtual File System (NWFS) to mount the VM that contains the items to recover. When the Hyper-V GLR completes, you can choose to unmount the VM from NWFS or perform another Hyper-V GLR operation. You can only use Hyper-V GLR to recover data that has been rolled over to backup media that supports GLR.

NMM only allows browsing and recovery of one VM at a time, and NMM does not support the recovery of the same backup to multiple clients simultaneously. If you try to mount another VM while you use the Hyper-V GLR, NMM unmounts the first VM and you lose access to the contents of the first VM until you remount it.

EMC recommends that you provision a system to act as the NMM Hyper-V GLR proxy. The Hyper-V GLR proxy has the required setup so that you can granularly recover files from a Hyper-V image backup. NMM supports the use of a 64-bit Windows Server 2008 R2 or 2012 VM or physical machine as the GLR proxy system.

Using this GLR proxy setup, you can recover files from any NMM Hyper-V backup of a Hyper-V VM with a supported operating system installed. After NMM recovers the files or folders that are local to the proxy machine, you must manually move or data mine the files as needed. The NetWorker administrator configures the NetWorker authentication to allow the NMM Hyper-V GLR proxy client the rights to recover any save set that NMM needs to mount by using NWFS.

To enable GLR functionality, install the NMM client software on the system and select the GLR option to install the NWFS features. You can install NMM and NWFS on any systems including any Hyper-V Guest VM, Hyper-V server, or stand-alone server, as long as the operating system is Windows Server 2008 R2 or greater.

The *NetWorker Module for Microsoft Administration Guide* and the *NetWorker Administration Guide* provides more information on NWFS.

## Windows Server 2012 Hyper-V GLR features

NMM supports the following Windows Server 2012 Hyper-V GLR features:

- Recovery of data from a VHDX hard disk
- Recovery of FAT32, NTFS, and ReFS volume data
- GLR of Hyper-V VMs over SMB
- Recovery of deduplicated data, after you enable the Deduplication role

NMM does not support the following Windows Server 2012 Hyper-V GLR features:

- Recovery of data from a Windows Server 2012 Storage Spaces disk on a VM
- Differencing disk with parent and child hard disk on different hard drives

## Recovering Hyper-V files and folders

While you perform the steps in the following procedure, you can switch to the Monitor pane to see the progress of the recovery operation and to check for error messages. NMM

displays any problems attaching hard disks, recognizing VM, or expanding VMs in the Monitor window.

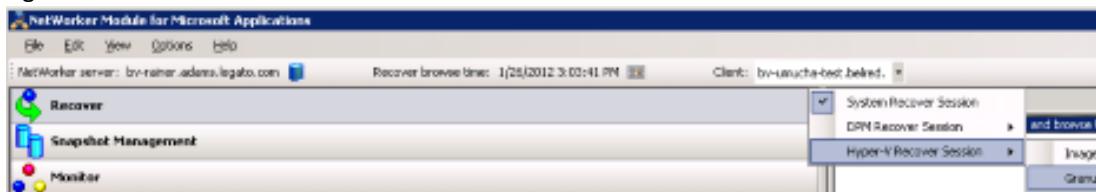
### Procedure

1. Open the NMM client on the GLR proxy server.

If the GLR proxy server is not available, open the NMM client on the physical machine where the backup was taken. This option requires you to install GLR on the physical machine.

2. From the left pane, click **Recover** > **Hyper-V Recover Session** > **Granular Level Recovery**.

NMM displays the GLR-capable VMs for the selected client, as shown in the following figure.



NMM displays the GLR-capable VMs for the selected client.

When you attempt to mount an incompatible backup, such as a backup not created using NMM 2.4 or later, or a backup of a non-Windows VM, an error message appears. To view all compatible backups available and select a compatible backup, launch the version dialog box.

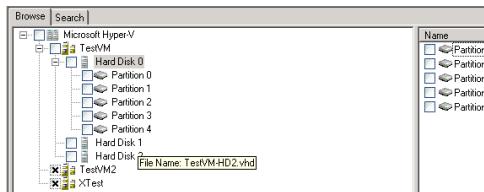
3. In the navigation tree, locate the Hyper-V VM that contains the folders or items you want to recover.

To search for a particular VM or item within the save set, right-click the VM and then click **Search for**.

4. Right-click the target VM and then click **Mount**.

NMM uses NWFS to mount the VM. If another VM is already mounted for GLR recovery, NMM notifies you that it will unmount the first machine. Click **OK** to continue the mounting process with the second VM, or click **Cancel** to leave the first machine mounted.

5. Click a VHD to display a list of the VMs it contains in the right pane, as shown in the following figure. If a VHD does not mount, NMM displays an error message. You can continue to work with the other available VHDs. Hyper-V GLR cannot mount raw virtual hard disks (hard disks that have not been formatted).



6. Locate the folders or items that you want to recover by completing one of the following steps:

- Expand the VMs to view a list of their contents.
- Right-click on a VM and then click **Search for** to search for a specific folder or item.

Depending on VM size, loading a VM can take awhile. You can start other operations in the Recover pane while waiting for a VM to load. After the VMs are loaded, you can perform data mining.

7. Select the check box beside each item that you want to recover. If you select a folder for recovery, then NMM also selects the folder contents for recovery.

You can specify a destination file path to recover the files. NMM retains the files in their original folder hierarchy. NMM overwrites files that already exist on the destination path. If you do not specify a destination path, NMM uses a default recovery path.

8. To specify a destination path for the recovered files, click **Recover Options** in the Hyper-V Recover Session toolbar.
9. In the **Hyper-V Recover Session Options** window, click the **Granular** tab and specify a destination path for the recovered items.
10. Close the **Hyper-V Recover Session Options** window.
11. Click **Recover** in the Hyper-V Recover Session toolbar.
12. In the **Hyper-V Recover Summary** window, click **Start Recover** to start the recovery. NMM recovers the files in the destination path by creating the original folder hierarchy.
13. Close the **Hyper-V Recover Session Summary** window.

To view the recovery progress, click **Monitor** in the left pane. When a recovery is in progress, you cannot perform other tasks in the Recover pane. After a recovery completes you can perform browse and recovery actions on the mounted VM.

## Recovering AES-encrypted data from a Hyper-V VM

You can encrypt backup and archive data on Windows hosts by using an Advanced Encryption Standard (AES) Application Specific Module (ASM) pass phrase. If you did not specify a pass phrase but configured an AES encrypted backup, NetWorker encrypts the data with a default pass phrase.

During data recovery, if you did not use the default or current pass phrase, you must specify the pass phrase used at the time of backup.

### NOTICE

Do not use AES encryption when backing up files that are encrypted by using Windows Encrypting File System (EFS). When NMM applies AES encryption to a file that is also encrypted by using the Microsoft EFS, NetWorker reports the backup as successful. However, recovery of the file will fail.

---

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

### Procedure

1. On the **Options** menu, select **System Recover Session Options**.
2. Click the **Security** tab.
3. Type the pass phrases.
4. Click **OK**.
5. Perform the granular level recovery.

# CHAPTER 5

## EMC Data Protection Add-in for SCVMM

This chapter includes the following sections:

• <a href="#">Overview</a> .....	66
• <a href="#">How the Data Protection Add-in works with SCVMM</a> .....	68
• <a href="#">Installation and uninstallation</a> .....	70
• <a href="#">Configuration</a> .....	75
• <a href="#">Data Protection Add-in overview data</a> .....	77
• <a href="#">Recoveries</a> .....	81
• <a href="#">Monitoring</a> .....	88
• <a href="#">Troubleshooting</a> .....	89

# Overview

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

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

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

## Recoveries

The Data Protection Add-in feature set supports recovery of Hyper-V VMs protected by NetWorker servers. The Data Protection Add-in supports recoveries of conventional backups to the original Hyper-V server on which the VM was backed up or to an alternate Hyper-V server. The Data Protection Add-in does not support recoveries of VMware VMs or recoveries from persistent snapshots.

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

- SCVMM console on the same machine as the SCVMM server
- SCVMM console on a different machine 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 will recover the VM. [Required privileges on page 67](#) provides details about SCVMM privileges. The *NetWorker Administration Guide* provides details about the required privileges.

## Backups

The SCVMM user cannot perform backups from the Data Protection Add-in. A NetWorker administrator must create and configure NetWorker client resources for the Hyper-V servers from which the SCVMM can recover VMs.

After the Hyper-V server has been added to a NetWorker client, the basic workflow for a scheduled backup of a VM managed by SCVMM is the same as that of a standard physical host. If the NetWorker administrator does not create and configure a NetWorker client for the Hyper-V server, then the NetWorker server cannot protect the SCVMM VM and therefore the VM is not available for recovery.

The *EMC NetWorker Module for Microsoft for Hyper-V Guide* Backups chapter and the *EMC NetWorker Administration Guide* provide details on scheduling and managing backups.

## Supported versions

The *NetWorker Online Software Compatibility Guide* on EMC Online Support lists the most up-to-date information about the Windows Server versions that NMM supports. The Data Protection Add-in and the NMM client software must be the same version.

The Data Protection Add-in supports the following component versions of System Center:

- System Center 2012 SP1 Virtual Machine Manager with System Center Update Rollup 4
- System Center 2012 R2 Virtual Machine Manager

The Data Protection Add-in is compatible with the following operating systems when imported into the System Center 2012 SP1 Virtual Machine Manager Console:

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

## Software dependencies

The Data Protection Add-in requires the following software:

- The Data Protection Add-in 8.2 SP1 and the NetWorker 8.2 SP1 client software must be installed on the SCVMM console machine. The Data Protection Add-in and the NMM client software must be the same version.
- The NetWorker and NMM 8.2 SP1 or greater client software must be installed on the Hyper-V server to which the VM is recovered.
- The Data Protection Add-in requires access to a NetWorker 8.2 SP1 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 data files 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 machine 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 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 71](#)
  - b. For SCVMM 2012 SP1,[Installing SCVMM Update Rollups on page 71](#)
  - c. [Installing the Data Protection Add-in on page 71](#)
2. Each user must perform the following steps:
  - a. [Importing the Data Protection Add-in on page 72](#)
  - b. [Activating the Data Protection Add-in on page 72](#)

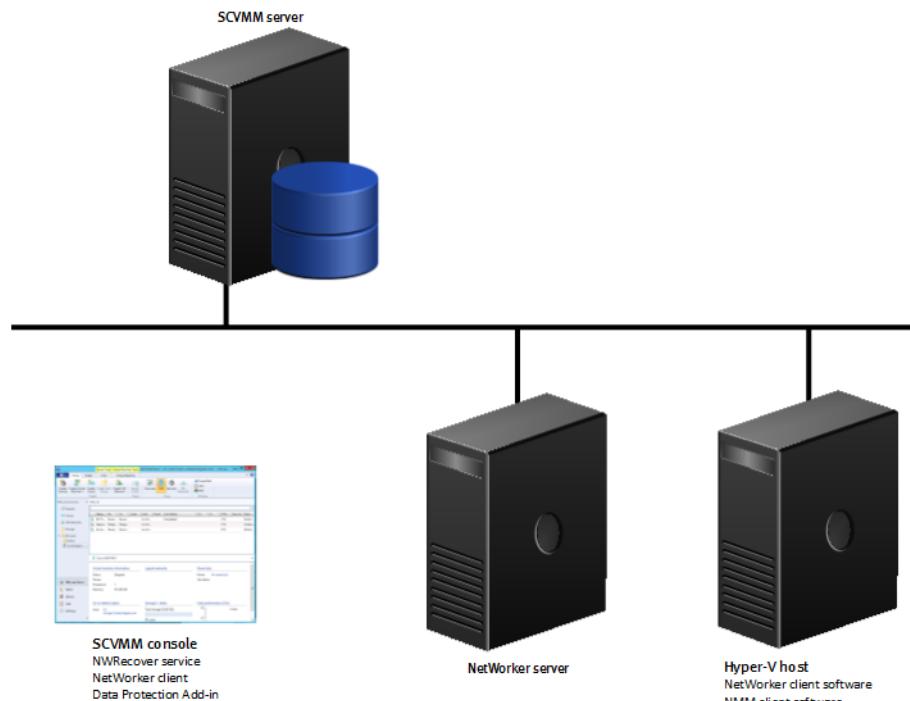
### After you finish

To uninstall the Add-in, an SCVMM administrator and each user must perform the steps described in [Uninstalling the Data Protection Add-in on page 73](#).

## How the Data Protection Add-in works with SCVMM

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

**Figure 11** Data Protection Add-in architecture



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 machine or on the

SCVMM server. However, the NetWorker client must be installed on the SCVMM console machine.

## 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 VMs for the context you selected.
2. Displays VMs 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 VMs in the current context.
4. Displays the updated protection information on the Overview page and VMs 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 VM, backup time, and destination options you selected to the NWRecover service. The NWRecover service initiates the recovery process.
2. The NWRecover service invokes 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.

4. 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 following pages:

- Overview - Displays the protection status for all VMs 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 VMs 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 VMs 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 12** SCVMM user roles and actions allowed by the Data Protection Add-in

Role	Actions allowed
Fabric Administrator (Delegated Administrator)	Can see all VMs, hosts, and clouds. Can recover all VMs managed by SCVMM to original and alternate locations.
Tenant Administrator	Can see and recover VMs 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 VMs and hosts within the private cloud they manage. No recovery operations are allowed.
Application Administrator (Self-Service Administrator)	Can see and recover VMs 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 & standalone)
- HostGroup
- VM

The Microsoft website provides more information about SCVMM scopes.

## Installation and uninstallation

To install the Data Protection Add-in, an SCVMM administrator installs the NetWorker client, SCVMM and the SCVMM console, SCVMM update rollups, and the Data Protection Add-in. To begin using the Data Protection Add-in, 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 SCVMM Update Rollups

Before using the SCVMM 2012 SP1 console, you must install [Update Rollup \(UR\) 4](#) or later. Specifically the Data Protection Add-in requires the Virtual Machine Manager Administrator Console (KB2888946) update. The Microsoft website provides the update rollups and instructions on how to install them.

## 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.

### 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 Data Protection Add-in software .zip file, either the 32-bit or 64-bit version as appropriate for the system, 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 is set to auto and will automatically start 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 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 will be used in [Importing the Data Protection Add-in on page 72](#). 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 data files 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.ScvmmAddIn.zip** and click **Open**. For example: C:\\Users\\Public\\Documents\\EMC NetWorker\\nsr\\addins\\VMM\_DataProtection.
8. To continue installing, select the check box 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 2012\\Virtual Machine Manager\\bin\\AddInPipeline\\AddIns\\<domain\_username>.

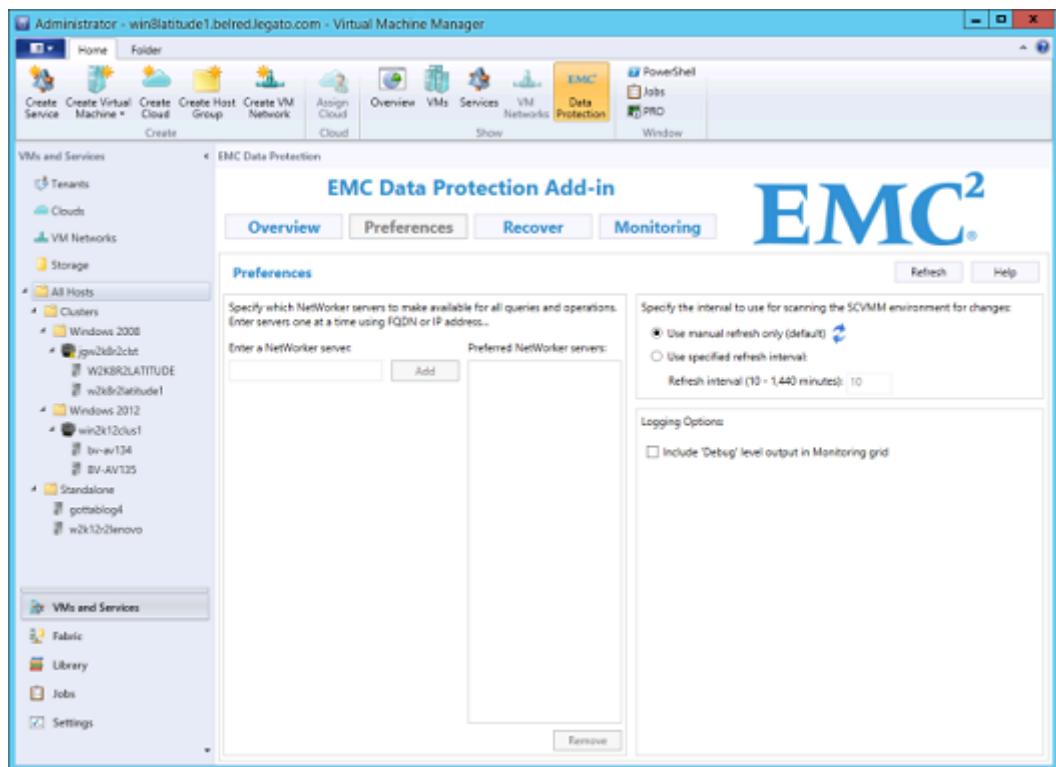
## Activating the Data Protection Add-in

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, as shown in the following figure:



The first time a user launches the Data Protection Add-in, 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 will perform other NetWorker operations on this machine, 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 will remove all components 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 will not be 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 machine 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 will be able to 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 and the NMM client software must be the same version. The *NetWorker Online Software Compatibility Guide* on EMC Online Support lists the most up-to-date information about supported versions.

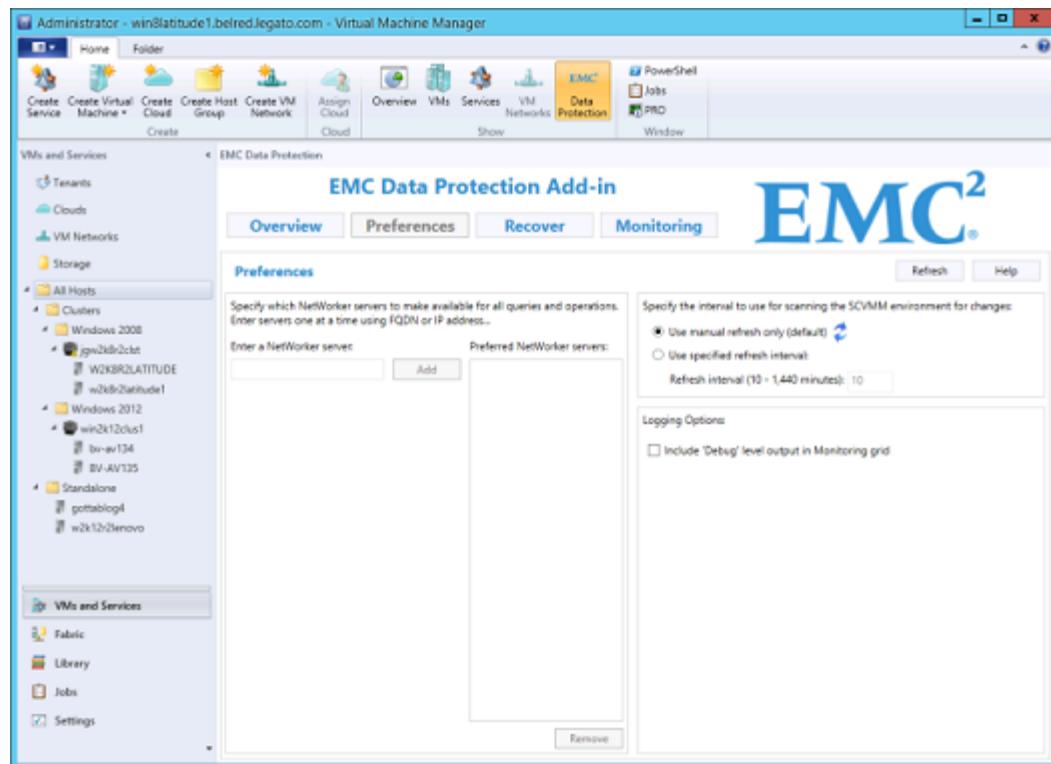
#### Procedure

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 described in [Uninstalling the Data Protection Add-in on page 73](#).
4. Follow the steps described in [Installing the Data Protection Add-in on page 71](#).
5. Follow the steps described in [Importing the Data Protection Add-in on page 72](#).
6. Follow the steps described in [Activating the Data Protection Add-in on page 72](#).
7. In the **Upgrade Successful** window, click **OK**.
8. Click **Refresh** to repopulate the Data Protection Add-in.

# Configuration

After completing the installation process, you must configure the Data Protection Add-in to access the NetWorker servers that contain VM 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 from within the Data Protection Add-in to ensure the Add-in is displaying current information.

**Figure 12** Data Protection Add-in for SCVMM Preferences page



## 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 navigation pane, select the host or cloud you wish 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**.

---

### Note

The Data Protection Add-in does not support IPv6 addresses.

The NetWorker FQDN or IP address displays in the Preferred NetWorker servers list.

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 VM protection status on the Overview page and the available VM 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 wish 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 VM 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 your 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 VM savesets

The Data Protection Add-in learns about protected virtual machines by querying the NetWorker servers that are specified on the Preferences page. If there is conflicting data regarding a Hyper-V server and its virtual machine protection because the Hyper-V server is a client of multiple NetWorker servers, the Data Protection Add-in might display inconsistent data.

Therefore, if you are using multiple NetWorker servers that define the same Hyper-V clients and VM save sets, it is best to change the Preference page's to one NetWorker

server at a time to reduce 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 VM save sets, then the Data Protection Add-in arbitrarily chooses information from one NetWorker server if conflicts exist. This 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 VMs 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 VMs that are protected on the NetWorker servers listed on the Preferences page.

The pie chart provides the following data about virtual machines:

#### **VMs excluded from protection**

These VMs 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 VMs are not configured for a scheduled backup as part of a NetWorker client resource and not explicitly excluded for backup.

#### **VMs protected**

These VMs are configured for scheduled backup as part of a NetWorker client resource.

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

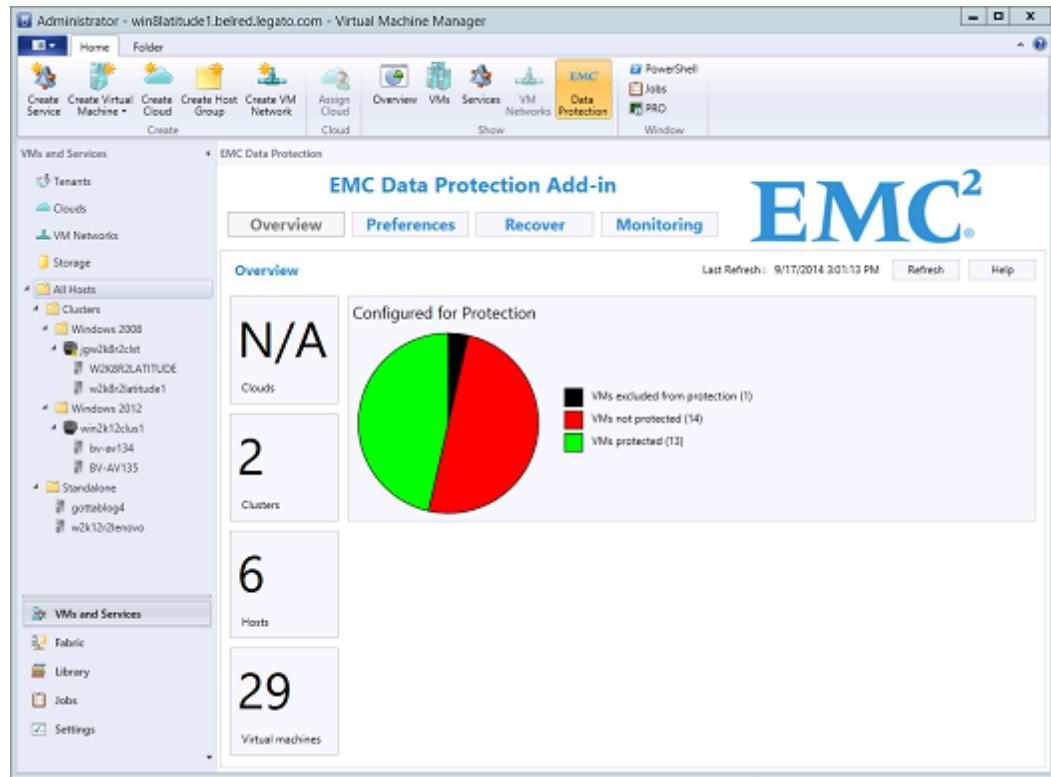
---

#### **Note**

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

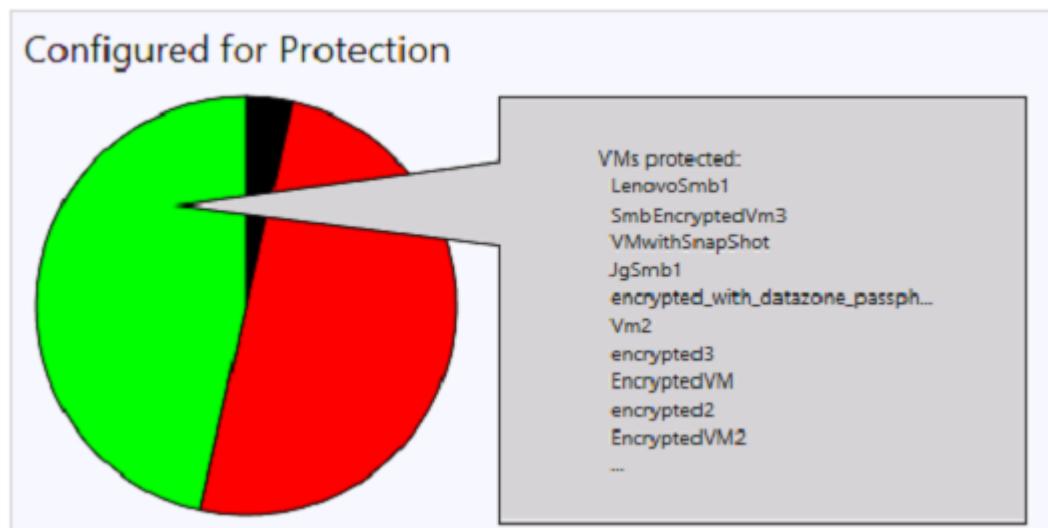
The following figure shows the Data Protection Add-in Overview page for Administrator, Fabric Administrator, and Read-Only Administrator user roles.

**Figure 13** Data Protection Add-in Overview page for Administrator, Fabric Administrator, and Read-Only Administrator user roles



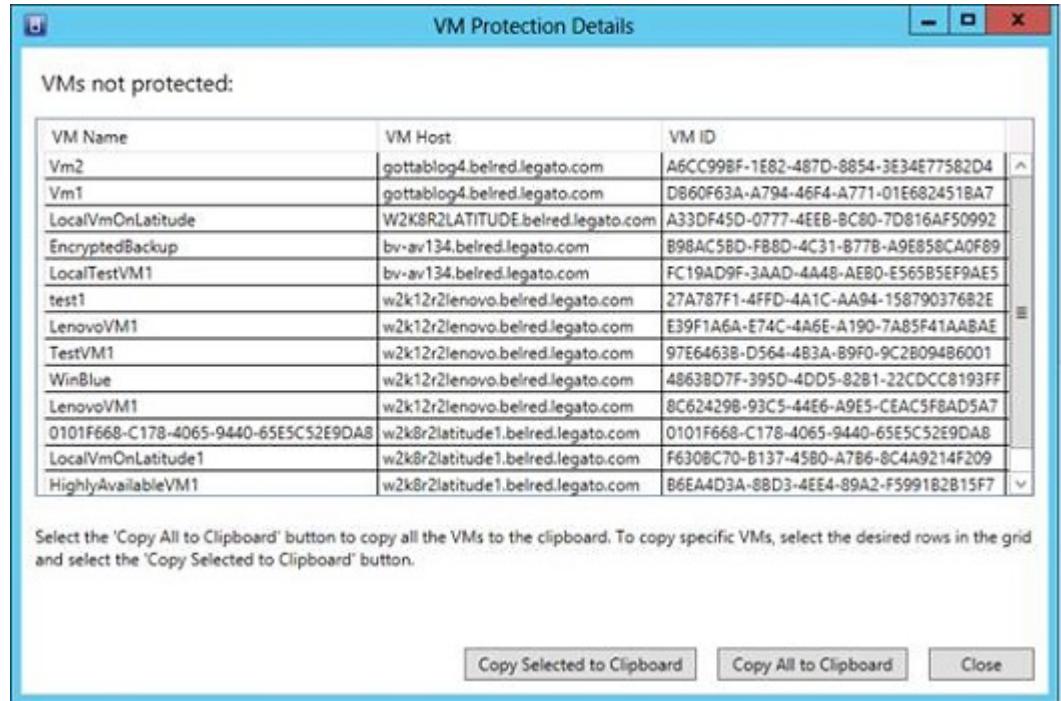
When you position the mouse over a protection category in the pie chart, a tooltip lists the first 10 VMs for that protection category. If there are more than 10 VMs in that category, the list is truncated with an ellipsis. To view the full list, click the desired section of the pie chart. If the VM name is more than 30 characters in length, the tooltip truncates the VM name with an ellipsis. The following figure shows the pie chart and tooltip for Administrator, Fabric Administrator, and Read-Only Administrator user roles.

**Figure 14** VM Protection Details tooltip for Administrator , Fabric Administrator, and Read-Only Administrator user roles



When you click a protection category in the pie chart, the VM Protection Details window displays. This window contains a table that lists the name, host, and ID for each VM in the selected protection category. To copy data for all of the VMs to the clipboard, click the Copy All to Clipboard button. To copy data for specific VMs, select the desired 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 15** VM Protection Details window for Administrator, Fabric Administrator, and Read-Only Administrator user roles



The screenshot shows a Windows application window titled "VM Protection Details". Inside, there's a section labeled "VMs not protected:" containing a table with three columns: "VM Name", "VM Host", and "VM ID". Below the table is a note about copying data to the clipboard. At the bottom are three buttons: "Copy Selected to Clipboard", "Copy All to Clipboard", and "Close".

VM Name	VM Host	VM ID
Vm2	gottablog4.belred.legato.com	A6CC99BF-1E82-487D-8854-3E34E77582D4
Vm1	gottablog4.belred.legato.com	D860F63A-A794-46F4-A771-01E682451BA7
LocalVmOnLatitude	W2K8R2LATITUDE.belred.legato.com	A33DF45D-0777-4EEB-BC80-7D816AF50992
EncryptedBackup	bv-av134.belred.legato.com	B98AC5BD-FB8D-4C31-B77B-A9E858CA0F89
LocalTestVM1	bv-av134.belred.legato.com	FC19AD9F-3AAD-4A48-AEBO-E565B5EF9AE5
test1	w2k12r2lenovo.belred.legato.com	27A787F1-4FFD-4A1C-AA94-15879037682E
LenovoVM1	w2k12r2lenovo.belred.legato.com	E39F1A6A-E74C-4A6E-A190-7AB5F41AABAE
TestVM1	w2k12r2lenovo.belred.legato.com	97E6463B-D564-4B3A-B9F0-9C2B094B6001
WinBlue	w2k12r2lenovo.belred.legato.com	4863BD7F-395D-4DD5-8281-22CDCC8193FF
LenovoVM1	w2k12r2lenovo.belred.legato.com	8C624298-93C5-44E6-A9E5-CEAC5F8AD5A7
0101F668-C178-4065-9440-65E5C52E9DA8	w2k8r2latitude1.belred.legato.com	0101F668-C178-4065-9440-65E5C52E9DA8
LocalVmOnLatitude1	w2k8r2latitude1.belred.legato.com	F630BC70-B137-45B0-A7B6-8C4A9214F209
HighlyAvailableVM1	w2k8r2latitude1.belred.legato.com	B6EA4D3A-8BD3-4EE4-89A2-F599182B15F7

Select the 'Copy All to Clipboard' button to copy all the VMs to the clipboard. To copy specific VMs, select the desired rows in the grid and select the 'Copy Selected to Clipboard' button.

### 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 VMs are not currently backed up as part of a NetWorker client resource.

#### VMs backed up

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

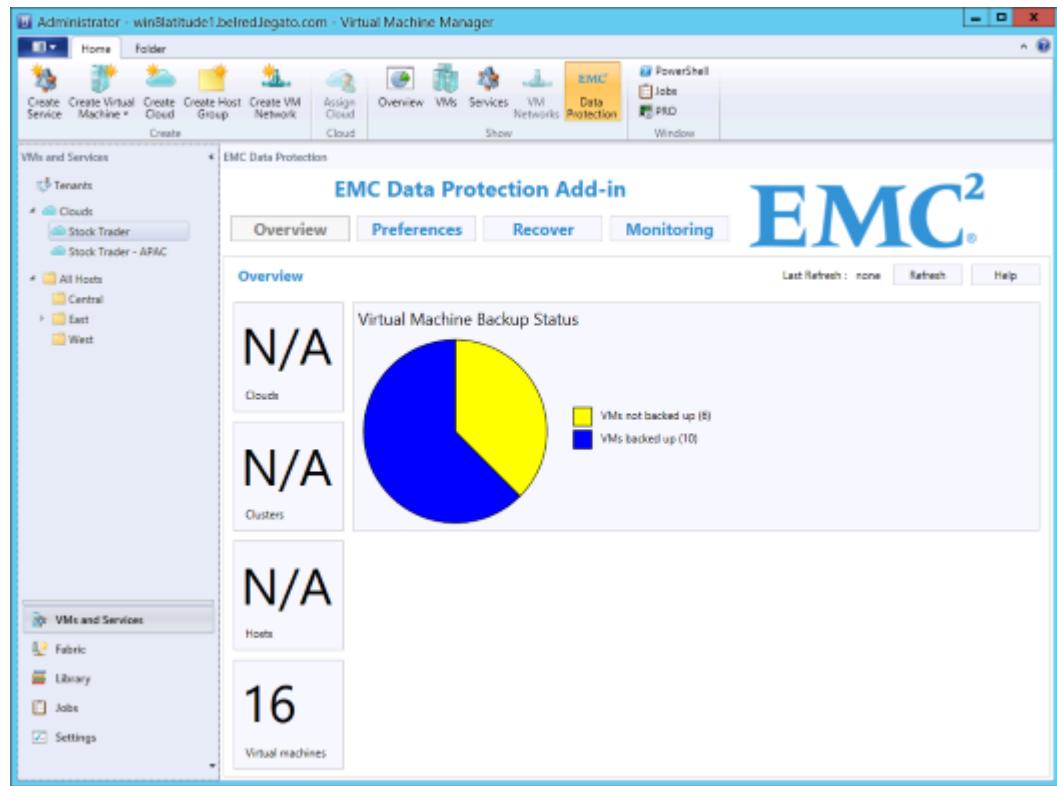
---

#### Note

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, and any of these VMs are backed up, the Data Protection Add-in shows all of the VMs as backed up.

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

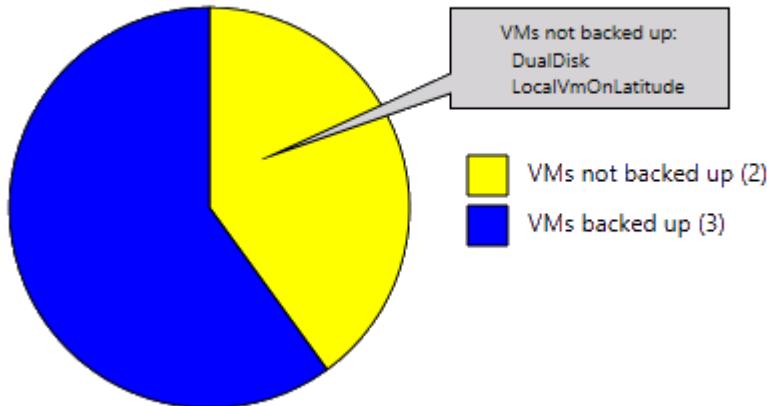
**Figure 16** Data Protection Add-in Overview page for Tenant Administrator and Application Administrator user roles



When you position the mouse over a backup status category in the pie chart, a tooltip lists the first 10 VMs for that backup status category. If there are more than 10 VMs in that category, the list is truncated with an ellipsis. To view the full list, click the desired section of the pie chart. If the VM name is more than 30 characters in length, the tooltip truncates the VM name with an ellipsis. The following figure shows the pie chart and tooltip for Tenant Administrator and Application Administrator user roles.

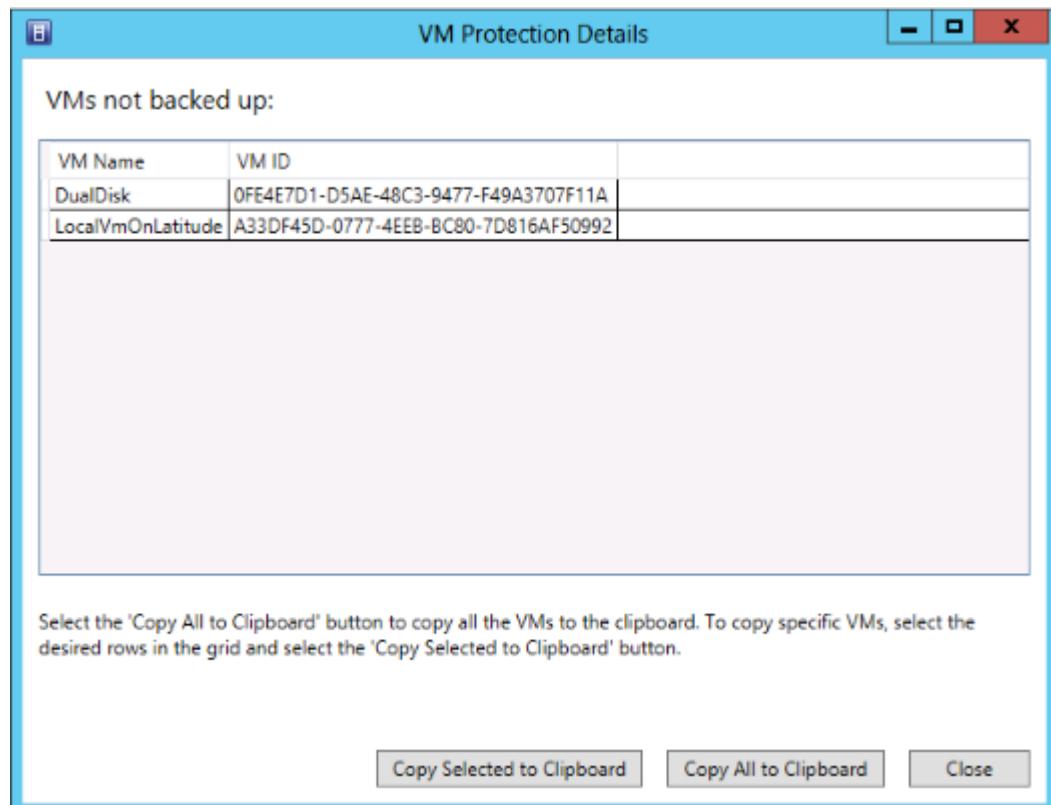
**Figure 17** Virtual Machine Backup Status tooltip for Tenant Administrator and Application Administrator user roles

### Virtual Machine Backup Status



When you click a backup status category in the pie chart, the VM Protection Details window displays. This window contains a table that lists the VM name and VM ID for each VM in the selected backup status category. To copy data for all of the VMs to the clipboard, click the Copy All to Clipboard button. To copy data for specific VMs, select the desired 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 18** VM Protection Details window for Tenant Administrator and Application Administrator user roles

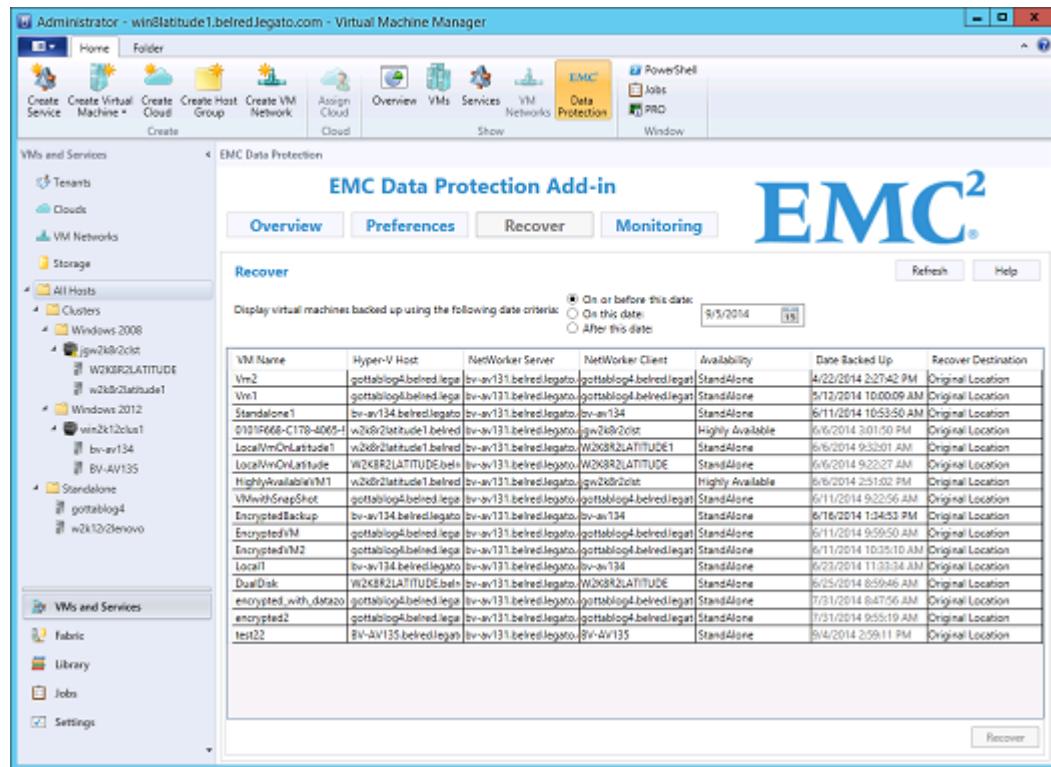


## Recoveries

The Recover page displays a list of all VMs managed by SCVMM that are backed up by a NetWorker server in the Preferred Servers list and in the current selected context and

match the date filtering criteria. When you access the Recover page for the first time, click Refresh to populate the grid with backups performed on the VMs in the current context.

**Figure 19** Data Protection Add-in for SCVMM Recover page



If you make a change in the SCVMM environment, such as adding a VM, adding a NetWorker server on the Preferences page, or performing a redirected recovery, click Refresh to update the list of VMs on the Recover page.

When performing VM recoveries by using the Data Protection Add-in, 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 at the same time, 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. For example: C:\Program Files\Emc Networker\nsr\applogs\nmm.raw.

**Figure 19** Data Protection Add-in for SCVMM Recover page (continued)

- Open the Windows Event Viewer on the machine that is hosting the SCVMM console. To access the event logs, navigate to Application and Services Logs > Networker Recovery Service.
- For cluster configurations, recovery to the original location is always to the active node of the cluster, regardless of the existing VM physical host location. Before starting the recovery, confirm that the cluster active node is the same as the VM physical host. After the recovery is complete, you might need to use Microsoft Cluster Manager to make the VM highly available again.

**Note**

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

- For highly available VM 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 VM 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.

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

## Viewing available VMs

The Recover page displays a list of all VMs 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.

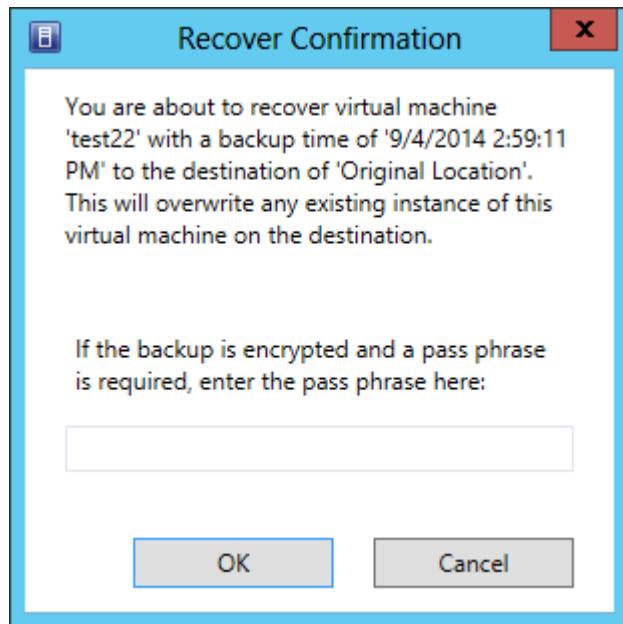
By default, the Recover page shows all VMs backed up on or before the current date. You can filter the VMs by selecting one of the date criteria options and choosing a date on the calendar. Only VMs with backup times that match the specified date filtering criteria are displayed. If you select criteria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

## VM Encrypted Recoveries

The Data Protection Add-in allows you to recover VM data that was backed up using AES encryption. When a NetWorker administrator configures a backup with AES encryption, NetWorker encrypts the backup data with a pass phrase stored on the NetWorker server. Before performing a recovery, contact your NetWorker administrator to determine whether a pass phrase is required for recovery.

The *NMM Administrator Guide*, *NMM Hyper-V VSS Guide*, and the *NetWorker Administration Guide* provide more details about AES and pass phrase usage.

Entering a pass phrase is not required if the backup is not encrypted. The Data Protection Add-in does allow you to specify an AES pass phrase and recover AES-encrypted VM data to the original location or to an alternate location. On the Recover page, after you select a VM and click the Recover button, a confirmation message dialog box allows you to enter a pass phrase, as shown in the following figure.



In the Recover Confirmation window:

- If the backup is encrypted and you need to specify a pass phrase, type the pass phrase and click **OK**.
- If the backup is encrypted and the current pass phrase on the NetWorker server is sufficient, you do not need to enter a pass phrase. Leave the pass phrase blank and click **OK**.
- If the backup is not encrypted, leave the pass phrase blank and click **OK**.

The Data Protection Add-in provides this pass phrase to NMM, and NMM performs the VM recovery. The Data Protection Add-in does not detect whether an AES password is required for a recovery. You will need to know whether a backup used AES encryption and which pass phrase was used for that backup.

When a pass phrase is required and a correct pass phrase is not provided, encrypted data is not recovered. Instead, file names are created without data; for example, the VHD file of a VM is created but there is no data associated with it. The recovery is failed and an error message with "*Invalid decryption key specified*" is displayed in the Monitor log.

The Data Protection accepts the same characters that the NMM UI accepts.

## Recovering a VM to the original location

The recovery operation runs on the Hyper-V server that is hosting the VM or, if the VM 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 VM you wish 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 the desired VM in the table.

7. Select the **Date Backed Up** cell, click again to activate the drop-down list, and select the desired backup date and time.
8. Click the **Recover** button.
9. In the **Recover Confirmation** dialog box that displays, do one of the following:
  - If the backup is encrypted and you need to specify a pass phrase, type the pass phrase and click **OK**.
  - If you do not need to specify a pass phrase, click **OK**.

[VM Encrypted Recoveries on page 83](#) provides details.

The recovery starts on the Hyper-V host. The Monitoring page displays the status of the recovery.

## Redirected recoveries

The Data Protection Add-in supports redirected recovery of VMs 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 VM placement path properties contain one or more paths. The redirected recovery location will be 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 taken before an NMM 8.2 upgrade.

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

## VM IDs after redirected recovery

NMM assigns a new VM ID in certain redirected recovery scenarios. The redirected recovery proceeds 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 VM will appear in both the source and destination hosts.

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

**Table 13** VM IDs after redirected recovery

Source operating system	Destination host		Destination VM ID assigned
	Operating system	Configuration type	
Windows Server 2008 R2	Windows Server 2008 R2	N/A	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

**Table 13** VM IDs after redirected recovery (continued)

Source operating system	Destination host		Destination VM ID assigned
	Operating system	Configuration type	
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 VMs 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 VMs with the same name on different source hosts to the same destination host, the Data Protection Add-in recovers these VMs to two different folders with unique subfolders by appending `vmname_timestamp` to the folder names. For example, if two VMs that are both named `Virtual_Machine` are recovered to the default SCVMM placement path property "`C:\ProgramData\Microsoft\Windows\Hyper-V`", the VMs 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 VM has multiple disks with the same name, the Data Protection Add-in recovers these disks to separate folders. For example, if a VM 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 VMs 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 VM file paths to 260 characters. If the appended file path exceeds 260 characters, the recovery fails.

---

## Performing a redirected recovery

### Procedure

1. (Recommended) Take the original VM 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 navigation pane, select the host or cloud that contains the VM you wish 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 the desired VM in the table.
8. Select the **Date Backed Up** cell, click again to activate the drop-down list, and select the desired backup date and time.
9. Select the **Recover Destination** cell, click again to activate the drop-down list, and select the desired 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. In the Recover Confirmation dialog box that displays, do one of the following:
  - If the backup is encrypted and you need to specify a pass phrase, type the pass phrase and click **OK**.
  - If you do not need to specify a pass phrase, click **OK**.
- [VM Encrypted Recoveries on page 83](#) provides details.
- The recovery starts on the destination Hyper-V host. The Monitoring page displays the status of the recovery.
12. If an Action Needed message displays, click **OK** to clear the message.
13. Confirm the VM is successfully recovered by verifying that the VM appears in the hypervisor on the Hyper-V server where you recovered the VM.
14. If the Action Needed message displayed, delete the original VM from its original host by using the SCVMM console. Alternatively, delete the VM by using Hyper-V Manager or PowerShell, and then refresh the SCVMM console.
15. In the navigation pane of the SCVMM console, right-click the destination host and click **Refresh Virtual Machines**.
16. If a new ID was assigned to the VM as described in [VM IDs after redirected recovery on page 85](#), in the navigation pane of the SCVMM console, right-click the source host and click **Refresh Virtual Machines**.
17. Ask the NetWorker administrator to perform a backup of the VM from its new Hyper-V host.

## Viewing VMs after a redirected recovery

If you perform a redirected recovery of a VM to a Hyper-V host, then the VM will not meet the criteria listed in [Viewing available VMs on page 83](#) until after a new backup of the VM is completed. Therefore, the Recover page does not immediately display the redirected

VMs. After you perform a redirected recovery for a VM, ask the NetWorker administrator to perform a backup of the Hyper-V host where the VM currently resides.

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

## Recovering a deleted VM

The Data Protection Add-in does not support recovering VMs 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 20** Data Protection Add-in for SCVMM Monitoring page

Date/Time	Event Type	Message
9/5/2014 12:51:07 PM	Info	===== Loading the 'EMC Data Protection Add-in for SCVMM' =====
9/5/2014 12:51:08 PM	Info	Initial SCVMM parameters: ScopeType='None', ContextObject='<none>', PowerShellContext is 'valid', UserDetails='BELRED\gr...
9/5/2014 12:51:08 PM	Info	The required service 'EMCNWRecoverSvc' is currently 'Running'.
9/5/2014 12:51:09 PM	Info	The SCVMM scope has changed to 'Host' and the context is 'w2k12r2lenovo.belred.legato.com'
9/5/2014 11:27:25 AM	Info	---- Initiating manual refresh ----
9/5/2014 11:27:25 AM	Debug	Transitioning from 'Idle' to 'Refreshing'.
9/5/2014 11:27:25 AM	Info	---- Refresh start time is 9/5/2014 11:27:25 AM. ----
9/5/2014 11:27:25 AM	Info	Beginning refresh of 1 EMC server(s) data.
9/5/2014 11:27:25 AM	Debug	Initiating retrieval of all VM hosts..
9/5/2014 11:27:25 AM	Debug	Start getting all the VM hosts from the PowerShell context.
9/5/2014 11:27:25 AM	Debug	Start executing the PowerShell callback to get all the VM hosts.
9/5/2014 11:27:25 AM	Debug	Start populating VM host table.
9/5/2014 11:27:25 AM	Info	Finished retrieving 6 VM hosts.
9/5/2014 11:27:25 AM	Debug	Finished executing the PowerShell callback to get all VM hosts.
9/5/2014 11:27:25 AM	Info	Callback complete for all VM hosts
9/5/2014 11:27:25 AM	Info	Refreshing EMC server bv-av131.belred.legato.com.
9/5/2014 11:27:25 AM	Info	Beginning refresh of local cache data of 'bv-av131.belred.legato.com'.
9/5/2014 11:27:25 AM	Info	Importer executing command C:\Program Files\EMC NetWorker\bin\nsadmin.exe -s bv-av131.belred.legato.com.
9/5/2014 11:27:25 AM	Info	Found 17 clients on server bv-av131.belred.legato.com.
9/5/2014 11:27:25 AM	Info	Begin refresh of client jgw2k8r2clst.
9/5/2014 11:27:25 AM	Info	Importer executing command C:\Program Files\EMC NetWorker\bin\nsinfo.exe -v -X all -x m -s bv-av131.belred.legato.co...
9/5/2014 11:27:36 AM	Info	Addition successful. ADDITION INFORMATION: Microsoft Windows\Windows\MS taken at 3/2/2014 8:22:27 AM on client jgw2k8r2clst

The Monitoring page shows 3 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 VM is restored. The exported log file

name is `MonitorExportFile` and is located at `C:\Users\<current user>\AppData\Local\EMC\NetWorker\SCVMM`.

## Troubleshooting

The following 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 VM doesn't start

If a recovered VM doesn't start, perform the following steps:

#### Procedure

1. Select the recovered VM, then right-click the VM and select **Discard Saved State**.
2. Right-click the recovered VM and then select **Properties**.
3. In the Properties dialog box, click **Hardware Configuration** and verify the Network Adapter settings of the VM.

### 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 noted above.

### Importing fails due to access issue

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

- `C:\Program Files\Microsoft System Center 2012 R2\Virtual Machine Manager\bin\AddInPipeline\AddInViews`
- `C:\Program Files\Microsoft System Center 2012 R2\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.ScvmmAddIn.AddInView` could not be found in the add-in package. Make sure 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 paths noted above.

## VM 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 VM.
2. In the Data Protection Add-in, click **Refresh**.

## Redirected recovery appears to succeed but no VM appears in Hyper-V Manager

If a redirected recovery appears to succeed but no VM 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 VM is disconnected from the network. The recovery succeeds according to the Data Protection Add-in monitor log and the NMM.RAW log, and the VM files are stored on the target host and volume, but Hyper-V Manager does not display or recognize the VM.

### Solution

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

```
$sw=Get-VMSwitch;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 VM might fail due to VM network or saved state incompatibility between the original Hyper-V host and the target Hyper-V host. The Hyper-V writer cannot register the VM because of errors in VM 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 VM files: look in the Monitor logs for the "redirected restore cmd line options:" output. If the VM files are there, then try to register the VM manually by using the SCVMM UI.

## Avoid VM 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 VMs. Although not required, it is considered best practice for virtual machine names to be unique.

## Cluster VM backups do not display on the Recover page

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

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

## Unable to recover if 'localhost' used as NetWorker server name

The use of 'localhost' as a NetWorker server name in the Data Protection Add-in Preferences page might cause recoveries to fail. Replace the 'localhost' entry with any of the following: 127.0.0.1, the machine's actual IP address (IPv4), or the machine's fully qualified domain name.

Sample error message for failed recovery:

```
12/9/2014 1:55:35 PM      Info      44498:nsrsnap_vss_recover:common\apputil(288): Cannot
connect to resdb database on localhost 102357:nsrsnap_vss_recover:Unknown Application
Information parameter: PROXY, may not be supported 38006:nsrsnap_vss_recover:Program not
registered. 102803:nsrsnap_vss_recover:Terminating restore due to error.
50338:nsrsnap_vss_recover: usage: nsrsnap_vss_recover [<options>]
80291:nsrsnap_vss_recover: options: [-v] [-c client-name] 43597:nsrsnap_vss_recover:
[-s server] 50422:nsrsnap_vss_recover:           [-A key=value] 80292:nsrsnap_vss_recover:
[-d destination] 80293:nsrsnap_vss_recover:           [-S SSID] or [-I -]
```

## Redirected recovery fails when the VM name or VM configuration path contains special characters

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

NMM Hyper-V supports the following characters in VM names and VM configuration paths, including stand-alone, CSV, and SMB 3.0 configurations:

- Alpha numeric (A–Z, a–z, 0–9)
- - . [ ] \_ { } + = ` ~ ! # \$ % ^ & ()
- Space

When you attempt to recover a Hyper-V save set, VM name, or file path that contains a character not listed above, the Data Protection Add-in checks the name and path of the VM objects and displays an error.

If a redirected recovery fails, you can use SCVMM to perform the following workaround:

### Procedure

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



# CHAPTER 6

## Best Practices and Recommendations

This chapter includes the following sections:

• <a href="#">Overview</a> .....	94
• <a href="#">Hyper-V Server backup and recovery best practices</a> .....	94
• <a href="#">Hyper-V VM applications backup and recovery best practices</a> .....	95
• <a href="#">Improving backup performance in Windows Server 2012 and 2012 R2 clusters with CSV</a> .....	95
• <a href="#">Data mining using Hyper-V granular level recovery</a> .....	96
• <a href="#">Restrictions and requirements for relocating and recovering data to a different location</a> .....	96
• <a href="#">Restrictions for backup and recovery of Hyper-V VMs in a Windows Server Failover Cluster</a> .....	97
• <a href="#">Restrictions for Hyper-V VM granular level recovery</a> .....	97

## Overview

This chapter includes best practices for backing up and recovering Hyper-V VMs.

## Hyper-V Server backup and recovery best practices

This type of backup uses the Hyper-V Writer on the Hyper-V server:

- To get the most benefit from the Hyper-V role, create separate VMs for each application, so that the application-type backup and recovery performed at the host level is only for Hyper-V.
- After disaster recovery of the Hyper-V server, you might need to recover applications within each VM:
  - If you are performing separate VM backups.
  - These backups are more recent than the complete Hyper-V server backups.
- Best use for this type of backup is Bare Metal Recover of a guest and for recovery of operating system roles.
- Best practice for Initial Store backup is to back up when Hyper-V configuration changes are made. Initial Store does not need to be backed up each time a VM guest is backed up.
- In NMM, the Hyper-V Writer does not support backup of the configuration file Initial Store to a proxy client.
- The primary purpose for recovering Initial Store in NMM is for disaster recovery of the Hyper-V Server.
- Back up DPM Server, Exchange Server, SQL Server, or SharePoint Server applications from within the VM.

EMC does not recommend that you use the Hyper-V backup to back up applications on a VM for the following reasons:

- The backup process uses the VM copy backup method.
- The backup process does not manage VM logs. For example, the backup process does not truncate Exchange logs.
- Roll-forward recovery is not available for VM level disaster scenarios. From a Hyper-V server, a roll-forward recovery of a VM is not possible. Recoveries from a Hyper-V server are point-in-time (disaster recovery).
- VM pass-through disks are skipped during Hyper-V server backup.
- Basic disks are supported only within VMs.  
NMM does not support backups of dynamic disks within VMs. NMM mounts the guest snapshot during the Hyper-V backup process, and this changes the disk signature on dynamic disks in a guest.
- NMM supports Windows Server 2008, 2008 R2, 2012, and 2012 R2 Failover Clustering, which allows you to configure a failover of VM.
- EMC storage connected to Fibre Channel or iSCSI storage can be used in the parent to host VMs.
- Do not take a Hyper-V VSS server snapshot of Hyper-V VMs that are part of a SharePoint farm.

To back up SharePoint on the Hyper-V VM:

1. Install the NMM client on the VM.
2. Perform the Share Point backup locally from within the VM.

The Microsoft website provides recommendations and requirements about using SharePoint and Hyper-V together.

## Hyper-V VM applications backup and recovery best practices

This type of backup and recovery is performed within the VM, and uses application and system components Writers available on that VM:

- Microsoft recommends using backups within the VM as the preferred method for Exchange backup and recovery.
- NMM supports roll-forward recovery for Exchange, when Exchange is backed up within the VM.
- Within VMs, standard application backup and recovery rules and capabilities apply, including roll-forward recoveries.
- NMM skips VM pass-through disks in Hyper-V backups. NMM supports pass-through disks backups within the VM.
- Windows Server Failover Clustering with iSCSI storage is supported.
- VM Windows Server failover clustering with fibre channel storage is not supported because SCSI-3 is not supported in Hyper-V VMs.

## Improving backup performance in Windows Server 2012 and 2012 R2 clusters with CSV

When you use the NMM cluster proxy client to perform image-level backups of VMs in a Windows Server 2012 and 2012 R2 cluster with the NMM cluster proxy client, NMM performs the backup by using the data is serviced from the cluster node that owns the CSV where the VM files reside. For example, Cluster Node 1 owns CSV 1 where the VM 1 files reside, and the cluster proxy client is currently executing on Cluster Node 2. When the cluster proxy node backs up VM 1, the backup process:

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

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

If you use a cluster proxy client for VM backups, then consider the following recommendations to improve performance for image-level backups:

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

To maximize 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:

- 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, this might

impact the performance of all the associated VMs on the CSV. The cluster proxy client node should have the capacity to be the owner of all 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 ownership of a CSV 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 is capable of granular level recovery for backups of Hyper-V VMs created with NMM 2.4 or above. NMM with Hyper-V also supports data mining the information from the VM image drives by using a third party tool such as Kroll OnTrack PowerControls.

To prepare to mine the data, use the NMM GUI to mount the VM, attach the VHDs, and load the VMs. Then you can access the data on the mounted VHDs using the third party tool, outside of the NMM GUI.

For example, if the VM guest is running SharePoint, first use the NMM GUI to mount the Hyper-V VM image, attach the VHSs, and load the VMs. Then use Kroll OnTrack PowerControls to recover SharePoint sites, lists, libraries, and items.

You must keep the NMM GUI open while you explore and recover files on the mounted VM VHDs. If you close or change the focus of the NMM GUI, you lose access to the mounted VHDs. A warning displays anytime a closure or focus change will cause loss of access to a mounted VM image.

## 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 redirect recovery of virtual machines to Hyper-V Servers of the same or higher release as the Hyper-V Server where the virtual machine was created. Although a virtual machine can be redirect recovered to an older release of a Hyper-V Server, it may not fully function on that server. For mixed environments, you might not be able 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 VMs to non-Hyper-V Servers.
- Recovering the Initial Store to a different location.
- Relocating or redirecting Hyper-V backups taken before an NMM upgrade.

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

- The parent partition must run Windows Server 2008 SP2 or higher in order to recover, with relocation of files, a VM that has Hyper-V snapshots.

- The destination host must have the NMM client installed.
- When you perform a directed recovery of a VM to a second Hyper-V Server, you must update the Network Adapter settings of the VM with the Hyper-V Manager before you start the VM.

## Restrictions for backup and recovery of Hyper-V VMs in a Windows Server Failover Cluster

When a Hyper-V VM resides on a physical host, which is part of a Windows 2008, 2008 R2, 2012, and 2012 R2 Failover Cluster, you cannot back up or recover the VM as part of the cluster virtual server.

For example, consider the following Failover Cluster setup:

- A cluster, Cluster\_Virtual\_Name, contains two physical machines, Physical\_Machine\_1 and Physical\_Machine\_2.
- Physical\_Machine\_1 contains two VMs, VM1 and VM2.

You want to back up and recover VM1.

If you create a NetWorker client resource for:

- Cluster\_Virtual\_Name, NMM does not support backup and recovery of VM1 through that client resource.
- Physical\_Machine\_1, you can specify the following values in the save set attribute:
  - APPLICATIONS:\Microsoft Hyper-V to back up the Hyper-V application. This includes all VMs on the physical machine.
  - APPLICATIONS:\Microsoft Hyper-V\VM1 to back up an individual VM, such as VM1.

You can recover VM1 from a backup of the NetWorker client resource of the physical machine, Physical\_Machine\_1. [Performing Hyper-V recovery to the original machine and location on page 53](#) describes how to perform this type of recovery.

You can perform a directed recover to recover VM1 from the NetWorker client resource of the physical machine, Physical\_Machine\_1 to the Physical\_Machine\_2. [Performing a directed Hyper-V recovery to a different machine or location on page 54](#) describes how to perform this type of recovery.

## Restrictions for Hyper-V VM granular level recovery

The following restrictions apply when you perform a granular level recovery of a Hyper-V VM.

Windows 2008 R2 and earlier does not support:

- Recovery of data from a VHDX hard disk.
- Recovery of deduplicated data. To recover deduplicated volume data, you must enable the Deduplication role.
- Recovery of ReFs volume data.

NMM Hyper-V GLR does not support differencing disk with parent and child hard disk on different hard drives.



# CHAPTER 7

## Troubleshooting

This chapter includes the following sections:

- [Troubleshooting backups](#).....100
- [Troubleshooting recovery](#).....100

## Troubleshooting backups

The following topics explain issues that might occur during the backup process for a Hyper-V environment, as well as steps to resolve or work around the issues.

### Redirected I/O status does not update after CSV backup

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 performing the backup.

If the redirected I/O status does not update properly after the NMM CSV backup is complete, you must clear the status by performing one of the following steps:

- Type the following commands at the command prompt to delete the stale shadows:  
**diskshadow**

```
DISKSHADOW> list shadows all
```

```
DISKSHADOW> delete shadows all
```

- Type the following command at the Windows PowerShell command prompt:  
**Test-ClusterResourceFailure "volume name"**

---

#### Note

This command might clear the "backup in progress" status only.

- If the "redirected access" status is not cleared after performing steps 1 and 2, change the coordinator node by moving the volume to another node in the cluster and verifying that the volume is online.
- Use nsrsvutil.exe to clear the backup state for the affected volume by typing the following:

```
nsrsvutil -c <csv_volume_path>
```

For example: **nsrsvutil -c "c:\ClusterStorage\Volume1"**

### Hyper-V pass-through disks may not be backed up in a child partition backup

For Hyper-V backups, the child partition pass-through disks are skipped in Hyper-V parent partition backup, and child partition pass-through disks are supported by backups within the child partition.

However, in some cases, Hyper-V parent partition backup of a child partition with a pass-through disk might fail completely. If this occurs, contact Microsoft support for assistance because the problem might be with the hardware configuration or the Microsoft Hyper-V writer.

## Troubleshooting recovery

The following topics explain issues that might occur while performing a Hyper-V recovery, as well as steps to resolve or work around the issues.

### Where the NMM GUI is opened for Hyper-V Server Core recovery or other CSV setup, recovery needs two environmental variables

In a Hyper-V Server core or other CSV setup, if you encounter an RPC service failure message, perform the following steps:

1. Create the following variables under **User environment variables**:

- **NSR\_CSC\_RPC\_TIMEOUT**

The typical value is 6. This value will be treated in seconds.

- NSR\_CSC\_METHOD\_TIMEOUT  
The typical value is 180. This value will be treated in seconds.
2. To apply the changes, restart the host where you launch the NMM recovery GUI.

**When recovering multiple Hyper-V CSV VM through proxy, all the VMs are recovered but all the VMs are not getting registered**

#### Problem

In a Hyper-V CSV setup, when recovery of multiple Hyper-V CSV VMs through proxy is performed, all the VMs are recovered although only one VM is registered.

#### Solution

After recovery of multiple Hyper-V CSV VMs through proxy is complete, NMM recovers .VHD and .XML files. Manually run the following Powershell command to register the VMs 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 unable to browse CSV mount point**

#### Problem

After disaster recovery, if NMM is used to recover Hyper-V data, the following issues are observed:

- Hyper-V recovery of VMs located in a shared disk (but non-CSV volume) fails.
- The CSV volumes are not browsable, and recovery of VMs located in 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 cross-check that the failover cluster virtual network name account of Hyper-V Virtual Server is enabled.

**Through Advanced Recovery option, recovery of online VM to other node in same cluster setup completes**

#### Problem

In a Hyper-V CSV setup, when a child partition is up and running, the same child partition can be recovered to another node by using the Advanced Recovery option. This creates multiple VMs in different CSV nodes.

#### Solution

If the VM is online or active, recover the VM to the same node.

**For Hyper-V image recovery, in the differencing disk type configuration, recovery to the alternate location fails with error**

#### Problem

In the differencing disk type configuration, image recovery of child partition to the alternate location fails with error.

### Solution

Follow the example solution provided:

If the source of the backup is:

- D:\VMDiff\Diff\VMdiff01\diff.vhd
- D:\VMDiff\Base\SA\_Gold\SA\_Gold.vhd

During the recovery, if the alternate location is P:\, then manually create the following path in P:\ and then perform the recover:

- P:\VMDiff\Diff\VMdiff01
- P:\VMDiff\Base\SA\_Gold

The recovery completes without any error.

### Recovery by using the GUI fails for a child partition in Hyper-V Server 2008 cluster

#### Problem

Recovery by using the GUI for a child partition that is available in Hyper-V Server 2008 cluster fails.

## Solution

Use the following procedures to recover a child partition in a Hyper-V Server 2008 cluster:

- If the child partition .vhdx file is corrupt and the child partition configuration file is intact:
  1. Locate the child partition (\*.vhdx) file path.
  2. Move the corrupted .vhdx file to a safe location (network share drive) for future reference, if needed.
  3. Create a dummy .vhdx file that is identical to the original child partition that was created before the corruption occurred.
  4. Use the NMM GUI to select the child partition and perform a recovery to the original location.
  5. Start the child partition after recovery is complete.
  6. After the child partition is functional, discard the .vhdx file that was placed in a safe location on the network share drive.
- Both the child partition .vhdx file and the child partition configuration file are corrupt:
  1. Locate the child partition .vhdx file and configuration file .xml path.
  2. Copy the child partition .vhdx file path and the configuration file to a safe location on a network share drive.
  3. Remove the child partition from the cluster service.
  4. From the Hyper-V Management Console, delete the child partition.
  5. Delete the corrupted .vhdx files from the local hard drive.
  6. Create a dummy child partition that has the identical details to the original child partition that was created before the corruption occurred. This includes the following:
    - Observe that two child partitions with the same name are available.
    - .vhdx file location

---

### Note

Do not install the operating system on the child partition.

7. From the Failover Cluster Manager Console, add this dummy child partition to the cluster service. This makes the child highly available.
8. Use the NMM GUI to select the child partition and perform a recovery to the original location.
9. From the Hyper-V Management Console, after the recovery operation:
  - a. Observe that two child partitions with the same name are available.
  - b. Note that one of the child partitions is the duplicate or corrupted instance that needs to be deleted.
10. From the Failover Cluster Manager Console, remove the child partition from the cluster service.
11. From the Hyper-V Management Console:
  - a. Start both instances of the child partitions. One of the instances will start successfully and the other will fail.
  - b. Identify the corrupted instance, the one that fails to start.
  - c. Delete the corrupted instance.
  - d. Shut down the working instance of the child partition.
12. From the Failover Cluster Manager Console, make the child partition highly available.
13. Start the child partition after the recovery operation is complete.
14. After the child partition is functional, discard the .vhdx files and the configuration file that were placed in a safe location on the network share drive.

### NMM registers corrupted Hyper-V child partition to Hyper-V Server

#### Problem

Even if a recovery operation for a Hyper-V child partition fails, NMM still 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.

# APPENDIX A

## Recovering SQL Server, Exchange Server, and SharePoint Server Items from a Hyper-V VM

This appendix includes the following sections:

- [Overview](#).....106
- [Recovering items](#).....106

## Overview

This appendix describes how to recover Microsoft SQL Server, Exchange Server, and SharePoint Server items stored in Hyper-V VMs by using GLR.

The Hyper-V writer of Microsoft Hyper-V Server (Windows Server 2008, 2008 R2, 2012, or 2012 R2) supports only full backups (VSS\_BT\_FULL). The Hyper-V requestor performs a full backup of VMs that run a Microsoft application (SQL, Exchange, or SharePoint). If a requestor specifies VSS\_BT\_COPY, then the Hyper-V writer still performs a full backup, as per the VSS MSDN documentation.

The following table shows the backup types set by the requestor by using the SetBackupState on the host and the backup type set by the Hyper-V requestor inside the guest.

**Table 14** Backup types

Backup type set by requestor via 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

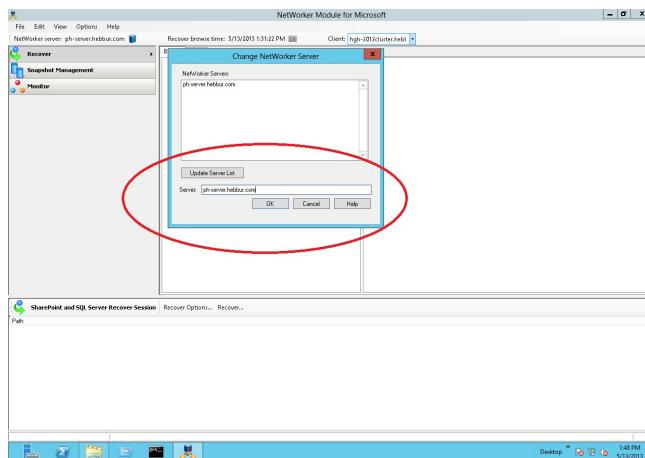
VM image backups are copy-type backups in-guest for applications. Log grooming requires a separate in-guest application backups. The Microsoft documentation provides information about the VSS\_BT\_FULL backup type.

## Recovering items

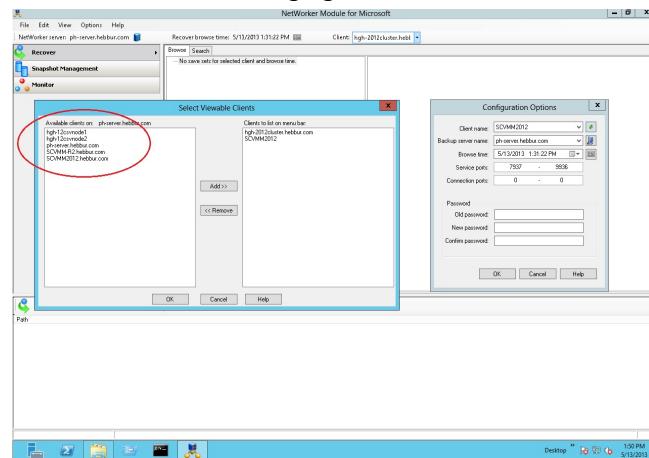
To recover items stored on a Hyper-V VM, you must first perform the following steps:

### Procedure

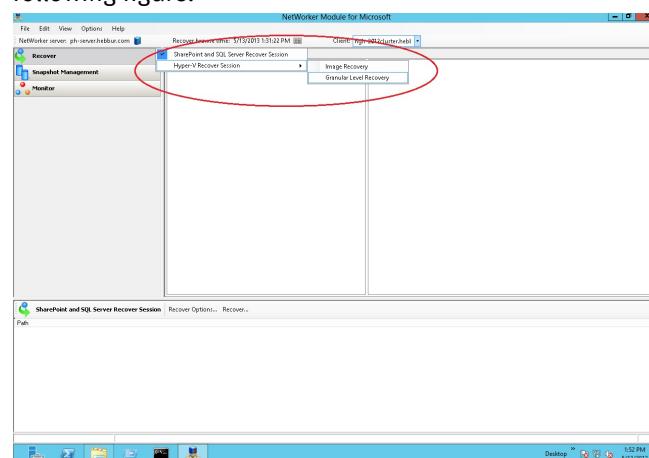
1. Configure the Hyper-V Client resources on NetWorker server and choose the Hyper-V Writer save set for backup.
2. Perform a full backup.
3. Open the NMM GUI on the FLR proxy server that you configured for GLR.
4. Select the NetWorker server where you performed the Hyper-V Server backup as shown in the following figure.



5. Use the Configure Option in the NMM GUI to select the Hyper-V Server Client resources as shown in the following figure.



6. Click Recover > Hyper-V Recover Session > Granular Level Recovery as shown in the following figure.



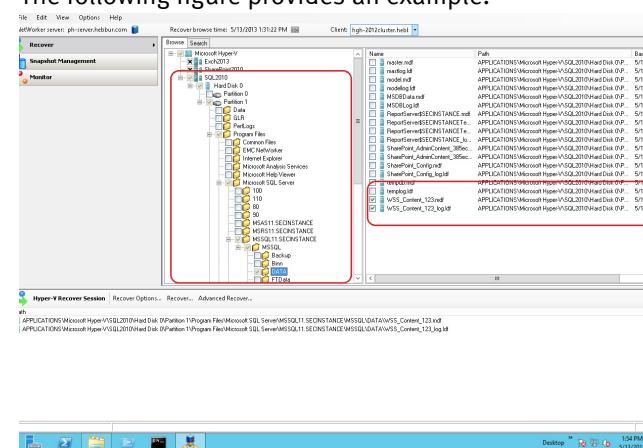
## Recovering SQL Server items

You can recover SQL Server items from a Hyper-V VM.

### Procedure

1. Mount the VM 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.



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 Management Studio, then 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 Management Studio with some data corruption or loss, then perform the following steps:

- d. Bring the database offline.
- e. Replace the existing database and logs with the recovered database and logs files.
- f. Bring the database online.
- g. Check that the recovered data is intact.

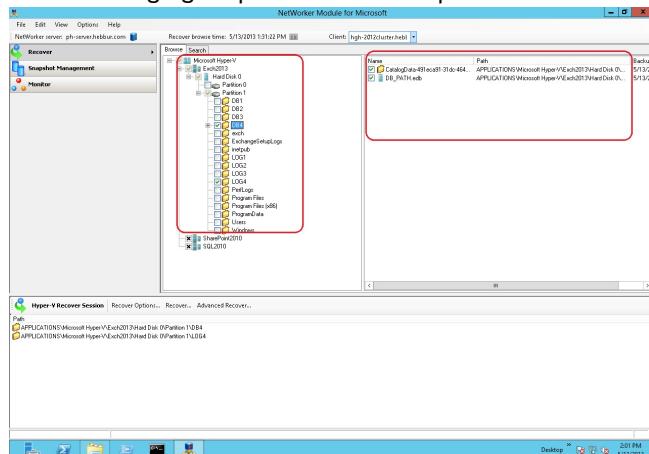
## Recovering Exchange Server items

You can recover Exchange Server items from a Hyper-V VM.

### Procedure

1. Mount the VM 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.



2. Select the database and logs 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:

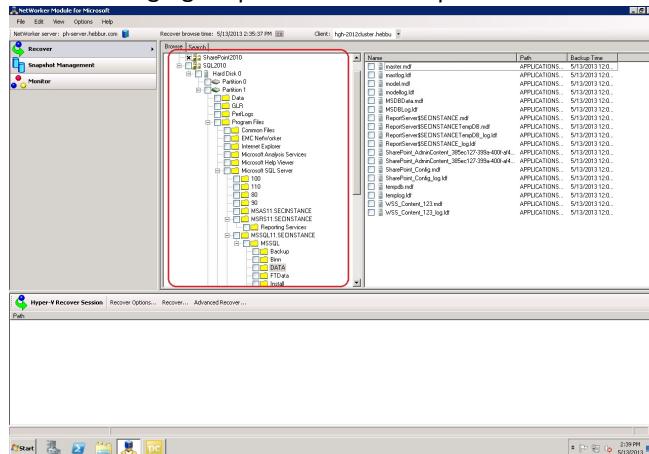
- e. Replace the existing database and logs folder with recovered database and logs folder in the actual path.
- f. Bring the database online.
- g. Check that the recovered data is intact.

## Recovering SharePoint Server items

### Procedure

1. Mount the VM 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.

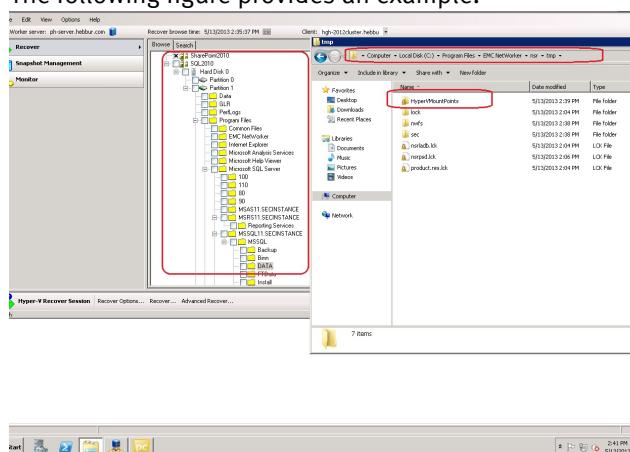
The following figure provides an example.



2. Select the database and logs 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\.

The following figure provides an example.



3. Use Kroll Ontrack Power Control Software to perform the SharePoint GLR.

You must install Kroll on the SharePoint Server and on the FLR proxy server where you mount the Hyper-V VM. 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.

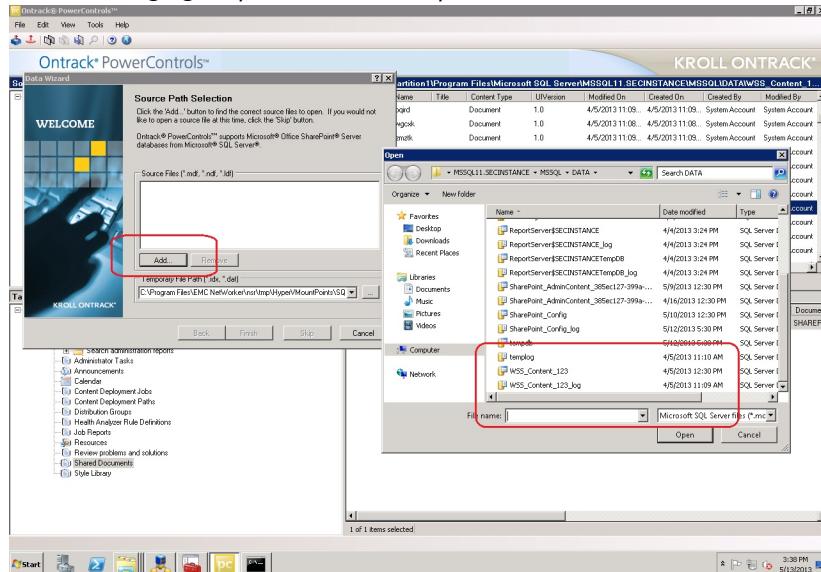
However, to recover items from a VM that hosts the SharePoint Server, you must configure Kroll differently. In Add the Source Path for the database, select the path

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\sqlfirstinst
\MSSQL11.FIRSTINSTANCE\MSSQL\DATA
```

The following figure provides an example.



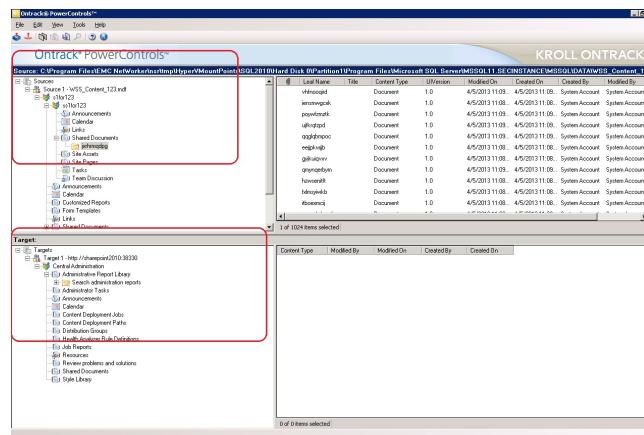
4. Provide the target SharePoint Server with credentials as shown in the following figure.

The following figure provides an example.



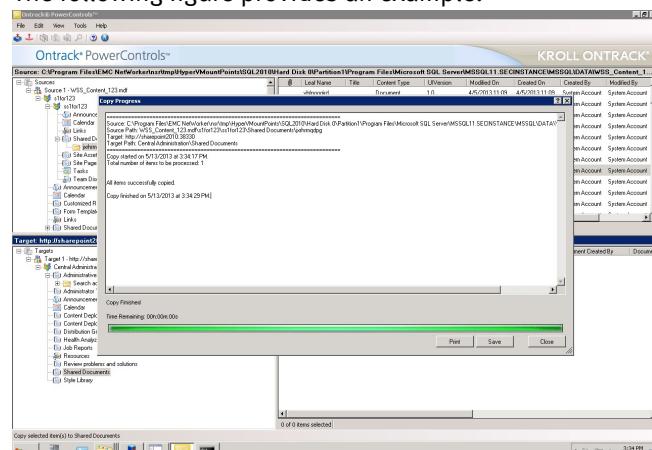
Kroll Ontrack Software configures itself with the SharePoint Server and FLR proxy server by scanning the logs, prescanning the logs, hashing the logs, and retrieving the content database.

The following figure provides an example.



- After the Kroll OnTrack configuration completes, copy the content to be recovered from the source to the target location.

The following figure provides an example.





# GLOSSARY

This glossary contains terms related to disk storage subsystems. Many of these terms are used in this manual.

## A

<b>ad hoc backup</b>	See <a href="#">manual backup</a> .
<b>administrator</b>	The person normally responsible for installing, configuring, and maintaining NetWorker software.
<b>administrators group</b>	Microsoft Windows user group whose members have the rights and privileges of users in other groups, plus the ability to create and manage the users and groups in the domain.
<b>Application Specific Module (ASM)</b>	Program that is used in a directive to specify how a set of files or directories is to be backed up or recovered. For example, compressasm is a NetWorker directive used to compress files.
<b>archive</b>	Backing up directories or files to an archive volume to free disk space. Archived data is not recyclable.
<b>archive volume</b>	Volume used to store archive data. Archived data cannot be stored on a backup volume or a clone volume.
<b>ASR writer</b>	The VSS Writer, which is responsible for identifying critical data that is needed to perform an offline restores.
<b>autochanger</b>	See <a href="#">library</a> .
<b>autochanger sharing</b>	See <a href="#">library sharing</a> .
<b>auto media management</b>	Feature that enables the storage device to automatically label, mount, and overwrite an unlabeled or recyclable volume.

## B

<b>backup</b>	Operation that saves data to a volume. See <a href="#">conventional backup</a> . See <a href="#">snapshot</a> .
<b>backup components</b>	See <a href="#">metadata document</a> .
<b>backup group</b>	See <a href="#">group</a> .
<b>backup level</b>	See <a href="#">level</a> .
<b>backup volume</b>	Volume used to store backup data. Backup data cannot be stored on an archive volume or a clone volume. See <a href="#">volume</a> .

<b>Boot Configuration Data (BCD)</b>	The ASR Writer component that identifies the location of the boot configuration database. This is required to perform an offline restore.
<b>bootstrap</b>	Save set that is essential for NetWorker disaster recovery procedures. The bootstrap consists of three components that reside on the NetWorker server. The media database, the resource database, and the server index.
<b>browse policy</b>	NetWorker policy that specifies how long backed-up data will be readily available for recovery. Backed-up data that has not exceeded its browse policy time can be recovered more quickly than data that has exceeded its browse policy time but not its retention policy time.

**See also** retention policy

## C

<b>carousel</b>	<b>See</b> <a href="#">library</a> .
<b>client</b>	Computer, workstation, or fileserver whose data can be backed up and recovered.
<b>client file index</b>	Database that tracks every database object, file, or file system that is backed up. The NetWorker server maintains a single client index file for each client.
<b>client-initiated backup</b>	<b>See</b> <a href="#">manual backup</a> .
<b>client resource</b>	NetWorker server resource that identifies the save sets to be backed up on a client. The client resource also specifies information about the backup, such as the schedule, browse policy, and retention policy for the save sets. <b>See</b> <a href="#">client</a> . <b>See</b> <a href="#">resource</a> .
<b>clone</b>	Reliable copy of backed up data. Unlike volumes created with a simple copy command, clone volumes can be used in exactly the same way as the original backup volume. Single save sets or entire volumes can be cloned.
<b>clone volume</b>	Exact duplicate of a backup volume. One of four types of volumes that NetWorker software can track (backup, archive, backup clone, and archive clone). Save sets of these different types may not be intermixed on one volume.
<b>cluster</b>	<ol style="list-style-type: none"> <li>Two or more independent network servers that operate and appear to clients as if they are a single unit. The cluster configuration enables work to be shifted from one server to another, providing "high availability" that allows application services to continue despite most hardware or software failures. Also known as an agent (Sun), logical server (HP TruCluster), package (HP-UX), and virtual server (Microsoft).</li> <li>Group of disk sectors. The operating system assigns a unique number to each cluster and keeps track of files according to which clusters they use.</li> </ol>
<b>cluster shared volume (CSV)</b>	A shared disk that contains an NTFS volume that is accessible for read and write operations by all nodes within the cluster. A virtual machine stored on CSV can change ownership from one node to another.
<b>Cluster VSS Writer</b>	In a Windows Server 2012 or 2012 R2 cluster with virtual machine storage on CSV, the Cluster VSS Writer reports components for backup for virtual machines that are owned by nodes other than the proxy or local node.

<b>command line</b>	Line on a display screen, also known as a command prompt or shell prompt, where you type software commands.
<b>component</b>	<ol style="list-style-type: none"> <li>1. Group of related data that must be treated as a single unit for backup and recovery.</li> <li>2. In Microsoft VSS terminology, a component is a subordinate unit of a writer.</li> </ol>
<b>components metadata document</b>	<a href="#">See metadata document</a> .
<b>consistent</b>	State of a dataset that is fully and immediately available to an application view.
<b>console server</b>	Software program that is used to manage NetWorker servers and clients. The Console server also provides reporting and monitoring capabilities for all NetWorker processes.
<b>conventional backup</b>	<a href="#">See nonpersistent snapshot</a> .
<b>critical volume</b>	Any volume containing system state files or files for an installed service, including volumes mounted as NTFS directories which contain such files. The volume where a critical volume is mounted is also considered to be critical. This is required to perform an offline restore, however maybe optional for this release depending upon the difficulties of implementing this feature.
<b>CSV Shadow Copy Provider</b>	The VSS provider that performs the snapshot for virtual machines that are owned by nodes other than the proxy node in a Windows Server 2012 or 2012 R2 cluster with virtual machine storage on CSV.

## D

<b>Data Mover (DM)</b>	Client system or application, such as NetWorker, that moves the data during a backup, recovery, or snapshot operation. <a href="#">See proxy client</a> .
<b>data retention policy</b>	<a href="#">See retention policy</a> .
<b>datawheel</b>	<a href="#">See library</a> .
<b>datazone</b>	Group of hosts administered by a NetWorker server.
<b>device</b>	<ol style="list-style-type: none"> <li>1. Storage unit that reads from and writes to backup volumes. A storage unit can be a tape device, optical drive, autochanger, or file connected to the server or storage node.</li> <li>2. When dynamic drive sharing (DDS) is enabled, refers to the access path to the physical drive.</li> </ol>
<b>directed recovery</b>	Method of recovery that recovers data that originated on one client computer and re-creates it on another client computer.
<b>directive</b>	Instruction that directs NetWorker software to take special actions on a given set of files for a specified client during a backup or recovery operation. Directives are ignored in manual (unscheduled) backups.
<b>disk subsystem</b>	Integrated collection of storage controllers or HBAs, disks, and any required control software that provides storage services to one or more hosts, such as CLARiiON arrays.

<b>Distributed File System (DFS)</b>	Microsoft Windows add-on that allows you to create a logical directory of shared directories that span multiple machines across a network.
<b>domain controller</b>	Computer that stores directory data and manages user interactions within a domain, including logon, authentication, directory searches, and access to shared resources.
<b>Dynamic Drive Sharing (DDS)</b>	Feature that allows NetWorker software to recognize shared drives.

## F

<b>file index</b>	See <a href="#">client file index</a> .
<b>File server VSS agent</b>	Responds to commands from the file server VSS provider to back up server applications that store data on a file server.
<b>File Server VSS Agent Service</b>	A service that lets you create volume shadow copies of applications that store data on a file server.
<b>File server VSS provider</b>	Sends commands through the VSS infrastructure to enable VSS operations on remote file servers.
<b>file system</b>	<ol style="list-style-type: none"> <li>1. The software interface used to save, retrieve, and manage files on storage media by providing directory structures, data transfer methods, and file association.</li> <li>2. The entire set of all files.</li> </ol>
<b>full backup</b>	See <a href="#">level</a> .

## G

<b>granular recovery</b>	Granular recovery provides the ability to recover specific files in seconds from a single backup. This dramatically reduces the recovery time and the footprint of the backup on storage resources.
<b>group</b>	Client or group of client computers that are configured to back up files at a designated time of day.
<b>guest operating system</b>	The operating system on a virtual machine.
<b>GUID</b>	The globally unique identifier of a virtual machine.

## H

<b>high-available system</b>	System of multiple computers configured as cluster nodes on a network that ensures that the application services continue despite a hardware or software failure. Each cluster node has its own IP address with private resources or disks that are available only to that computer.
<b>host ID</b>	Serial number that uniquely identifies a host computer.

## I

<b>inactivity timeout</b>	Number of minutes to wait before a client is considered to be unavailable for backup.
<b>incremental backup</b>	Backup level in which only files that have changed since the last backup are backed up. See <a href="#">level</a> .
<b>Initial Store</b>	An XML file on the management operating system that contains role-based security configuration details for Hyper-V.
<b>instant backup</b>	Process of creating a point-in-time copy (snapshot) of data from a single client and saving it on a primary storage volume, which can be immediately recovered as a backup copy.
<b>instant restore</b>	Process of copying data created during an instant backup to its original location, or to an alternate location, during a recover operation.
<b>Integration Components</b>	A collection of services and software drivers that maximize performance and provide a better user experience within a virtual machine. Integration services are only available through Integration Components for supported guest operating systems.

## J

<b>jukebox</b>	See <a href="#">library</a> .
----------------	-------------------------------

## L

<b>label</b>	Electronic header on a volume used for identification by NetWorker or other Data Mover application.
<b>legacy method</b>	Use of special-case Microsoft APIs to back up and recover operating system components, services, and applications.
<b>level</b>	Backup configuration option that specifies how much data is saved during a scheduled or manual backup. A full (f) backup backs up all files, regardless of whether they have changed. Levels one through nine [1-9] backup files that have changed since the last lower numbered backup level. An incremental (incr) backup backs up only files that have changed since the last backup.
<b>library</b>	Hardware device that contains one or more removable media drives, as well as slots for pieces of media, media access ports, and a robotic mechanism for moving pieces of media between these components. Libraries automate media loading and mounting functions during backup and recovery. The term library is synonymous with autochanger, autoloader, carousel, datawheel, jukebox, and near-line storage.
<b>library sharing</b>	Shared access of servers and storage nodes to the individual tape drives within a library.
<b>local cluster client</b>	NetWorker client that is not bound to a physical machine, but is instead managed by a cluster manager. It is also referred to as a logical or virtual client.
<b>locale settings</b>	Settings that specify the input and output formats for date and time, based on local language conventions.

**LUN (logical unit)** Logical unit of storage on a CLARiiON system. This refers to a device or set of devices, usually in a CLARiiON storage array.

**LUN address** SCSI identifier of a logical unit number (LUN) within a device target. Each LUN address identifies a device on a SCSI bus that can perform input/output (I/O) operations.

## M

**manual backup** Backup that a user performs from the client, also known as an unscheduled backup or an ad hoc backup. The user specifies the files, file systems, and directories to back up.

**media** Physical storage medium, such as magnetic tape, optical disk, or file system to which backup data is written.

**media database** Database that contains indexed entries of storage volume location and the life cycle status of all data and volumes managed by the NetWorker server. See [volume](#).

**media index** See [media database](#).

**metadata document** VSS Information stored in an XML document that is passed from the writer to the requestor. Metadata includes the Writer name, files, and components to back up, a list of components to exclude from the backup, and the methods to use for recovery. See [shadow copy set](#).

**mount** To make a database available for use or to place a removable tape or disk volume into a drive for reading or writing.

**mount point** See [volume mount point](#).

## N

**Network Data Management Protocol (NDMP)** TCP/IP-based protocol that specifies how heterogeneous network components communicate for the purposes of backup and recovery.

**NetWorker administrator** User who can add to or change the configuration of the NetWorker server, media devices, and libraries. NetWorker administrators must have their usernames included in the NetWorker server Administrator list.

**NetWorker client** See [client](#).

**NetWorker Console server** See [console server](#).

**NetWorker Management Console** See [console server](#).

**NetWorker server** Computer on a network running the NetWorker software, containing the online indexes, and providing backup and recover services to the clients on the same network.

**NetWorker storage node** See [storage node](#).

**nonclone pool** Pools that contain data that has not been cloned.

<b>noncritical volume</b>	A volume containing files that are not part of the system state or an installed service. The backup of non-critical volumes is not supported by either product for their initial releases.
<b>nonpersistent snapshot</b>	Snapshot backup that is moved to secondary storage on the NetWorker server or storage node and is no longer available for instant restore from a supported type of primary storage.

## O

<b>offline restore</b>	A restore operation performed from the Windows PE environment.
<b>online indexes</b>	Databases located on the NetWorker server that contain all the information pertaining to the client backups ( <a href="#">See client file index</a> ) and backup volumes ( <a href="#">See media database</a> ).
<b>online restore</b>	A restore operation performed using the normal recover UI, and the computer has been booted from an installed operating system.
<b>operator</b>	Person who monitors the server status, loads backup volumes into storage devices, and executes day-to-day NetWorker tasks.

## P

<b>pathname</b>	Set of instructions to the operating system for accessing a file. An <i>absolute pathname</i> indicates how to find a file starting from the root directory. A <i>relative pathname</i> indicates how to find the file starting from the current directory.
<b>persistent snapshot</b>	Snapshot that is retained on disk. A persistent snapshot may or may not be rolled over to tape.
<b>point-in-time copy (PiT)</b>	Fully usable copy of a defined collection of data, such as a consistent file system, database, or volume, which contains an image of the data as it appeared at a single point in time. A PiT copy is also called a shadow copy or a snapshot.
<b>policy</b>	Set of constraints that specify how long the save sets for a client are available for recovery. Each client has a browse policy and a retention policy. When the retention policy expires, the save sets associated with that policy are marked recyclable.
<b>pool</b>	Feature to sort backup data to selected volumes.
<b>PowerSnap</b>	EMC technology that provides point-in-time snapshots of data to be backed up. Applications that are running on the host system continue to write data during the snapshot operation, and data from open files is included in the snapshots.
<b>provider</b>	Software component defined by Microsoft VSS, that plugs in to the VSS environment. A provider, usually produced by a hardware vendor, enables a storage device to create and manage snapshots.
<b>proxy client</b>	Surrogate client that performs the NetWorker save operation for the client that requests the backup. A proxy client is required to perform a serverless backup.
<b>proxy node</b>	The node with the proxy cluster client in a Windows Server 2012 or 2012 R2 cluster with virtual machine storage on CSV.

## Q

**quiescing** Process in which all writes to disk are stopped and the file system cache is flushed. Quiescing the database prior to creating the snapshot provides a transactionally consistent image that can be remounted without file system checks or database consistency checks. Quiescing a database is the most common way of creating a database snapshot.

## R

**recover** To recover files from a backup volume to a client disk.

**Registry** Microsoft Windows database that centralizes all Windows settings and provides security and control over system, security, and user account settings.

**replica** See [shadow copy](#).

**requestor** Interface with the Microsoft VSS infrastructure to initiate the creation and destruction of [See shadow copy](#). NetWorker software is a requestor.

**resource** Component that describes the NetWorker server or its clients. Clients, devices, schedules, groups, and policies are all NetWorker resources. Each resource has attributes that define its properties.

**restore** Process of retrieving individual datafiles from backup storage and copying the files to disk.

**retention policy** NetWorker policy that specifies the minimum period of time that must elapse before backed-up data is eligible to be overwritten on the backup media. Backed-up data that has not exceeded its browse policy time can be recovered more quickly than data that has exceeded its browse policy time but not its retention policy time. [See browse policy](#).

**retrieve** To locate and recover archived files and directories.

**rollover** Process of backing up a snapshot to a conventional backup medium such as tape. Whether or not the snapshot is retained on disk depends on the snapshot policy.

**root** Highest level of the system directory structure.

## S

**save set** Group of files or a file system from a single client computer, which is backed up on storage media.

**save set ID (SSID)** Internal identification number assigned to a save set.

**save set recover** To recover data by specifying save sets rather than by browsing and selecting files or directories.

**save set status** NetWorker attribute that indicates whether a save set is browsable, recoverable, or recyclable. The save set status also indicates whether the save set was successfully backed up.

<b>save stream</b>	The data and save set information being written to a storage volume during a backup.
<b>server index</b>	See <a href="#">client file index</a> .
<b>serverless backup</b>	Backup method that uses a proxy client to move the data from primary storage on the application server host to secondary storage on another host. Serverless backups free up resources on the application server by offloading the work of processing snapshots to a secondary host.
<b>Server Message Block (SMB) 3.0 file share</b>	File share that uses the SMB 3.0 protocol. You can store virtual machines on a SMB 3.0 file share for a stand-alone Windows Server 2012 or 2012 R2 computer with Hyper-V or for a Windows Server 2012 or 2012 cluster with Hyper-V.
<b>service port</b>	Port used to listen for backup and recover requests from clients through a firewall.
<b>shadow copy</b>	Temporary, point-in-time copy of a volume created using VSS technology. See <a href="#">Volume Shadow Copy Service (VSS)</a> .
<b>shadow copy set</b>	Complete roadmap of what was backed up at a single instant in time. The shadow copy set contains information about the Writers, their components, metadata, and the volumes. A backup components metadata document containing that information is created and returned to the requestor after the snapshot is complete. NetWorker uses this document with the corresponding save set at recover time.
<b>shadow copy technology</b>	Defined and standard coordination between business application, file system, and backup application that allows a consistent copy of application and volume data to exist for replication purposes.
<b>skip</b>	Backup level in which designated files are not backed up.
<b>snap clone</b>	Exact copy of a snap set data backup. The clone operation is an archive operation without the deletion of the source data. A new snap ID is assigned to the cloned copy.
<b>snap ID</b>	Also known as a snapid, a unique 64-bit internal identification number for a snap set.
<b>snap set</b>	Group of files, volumes, or file systems from a single client, describing the collection of data for which a point-in-time copy is created on an external disk subsystem, such as a storage array.
<b>snapshot</b>	Point in time, read-only copy of data created during an instant backup. In Microsoft applications, this is known as a shadow copy or replica.
<b>snapshot expiration policy</b>	Policy that determines how long snapshots are retained before their storage space is made available for the creation of a new snapshot.
<b>snapshot policy</b>	Set of rules that control the lifecycle of a snap set. The snapshot policy specifies the frequency of snapshots, and how long snapshots are retained before recycling.
<b>snapshot retention policy</b>	Policy that determines how many PIT copies are retained in the media database and thus are recoverable.
<b>staging</b>	Moving data from one storage medium to a less-costly medium, and later removing the data from its original location.

<b>stand-alone device</b>	Storage device that contains a single drive for backing up data. Stand-alone devices cannot store or automatically load backup volumes.
<b>storage device</b>	<a href="#">See device</a> .
<b>storage node</b>	Storage device physically attached to a computer other than the NetWorker server, whose backup operations are administered from the controlling NetWorker server.
<b>system state</b>	All files that belong to VSS Writers with a usage type of BootableSystemState or SystemService. This is required to perform an offline restore.

## V

<b>virtual hard disk (VHD) files</b>	The file format for a virtual hard disk, which is the storage medium for a virtual machine. A VHD file can reside on any storage topology that the management operating system can access, including external devices, storage area networks, and network-attached storage. For Windows Server 2008 R2 and earlier, the file extension is .vhd. For Windows Server 2012 and later, the file extension is .vhdx.
<b>volume</b>	<ol style="list-style-type: none"> <li>1. A unit of physical storage medium, such as a magnetic tape, optical disk, or file system to which backup data is written.</li> <li>2. An identifiable unit of data storage that may reside on one or more host disks.</li> </ol>
<b>volume ID</b>	Internal identification that NetWorker software assigns to a backup volume.
<b>volume mount point</b>	Disk volume that is grafted into the namespace of a host disk volume. This allows multiple disk volumes to be linked into a single directory tree, and a single disk or partition to be linked to more than one directory tree.
<b>volume name</b>	Name assigned to a backup volume when it is labeled. <a href="#">See label</a> .
<b>volume pool</b>	<a href="#">See pool</a> .
<b>Volume Shadow Copy Service (VSS)</b>	Microsoft technology that creates a point-in-time shadow copy of a disk volume. NetWorker software backs up data from the shadow copy. This allows applications to continue to write data during the backup operation, and ensures that open files are not omitted.
<b>VSS</b>	<a href="#">See Volume Shadow Copy Service (VSS)</a> .
<b>VSS component</b>	Subordinate unit of a writer.

## W

<b>writer</b>	Database, system service, or application code that provides metadata document information about what to back up and how to handle VSS component and applications during backup and recovery operations. A Writer provides information to requesters to ensure that application data is consistent, application files are closed and ready for a slight pause to make a Shadow Copy.
---------------	---