

Veritas Storage Foundation™ and Disaster Recovery Solutions Guide for Microsoft Hyper-V™

Windows Server 2012 (x64)

6.0.2

Veritas Storage Foundation™ and Disaster Recovery Solutions Guide for Microsoft Hyper-V™

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product version: 6.0.2

Document version: 6.0.2 Rev 0

Legal Notice

Copyright © 2013 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, the Checkmark Logo, Veritas, Veritas Storage Foundation, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. See the Third-party Legal Notices document for this product, which is available online or included in the base release media.

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043
<http://www.symantec.com>

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our website at the following URL:

www.symantec.com/business/support/index.jsp

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/contact_techsupp_static.jsp

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level

- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan	customercare_apac@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportolutions@symantec.com

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

<http://www.symantec.com/connect/storage-management>

Documentation

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

doc_feedback@symantec.com

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

<https://www-secure.symantec.com/connect/storage-management/forums/storage-and-clustering-documentation>

Contents

Technical Support	4
Section 1 Storage Foundation for Windows (SFW)	11
Chapter 1 Introduction to Storage Foundation for Windows solutions for Hyper-V environments	13
Advantages of running Storage Foundation for Windows in the Hyper-V parent	13
How Storage Foundation for Windows manages storage for virtual machines	14
Related documentation on Storage Foundation for Windows	15
Chapter 2 Live migration support for SFW dynamic disk group	17
Tasks for deploying live migration support for Hyper-V virtual machines	17
About implementing Hyper-V virtual machine live migration on SFW storage	19
Requirements for Storage Foundation for Windows (SFW)	20
Installing Windows Server 2012	21
Preparing the host machines	21
Adding the Hyper-V role	21
Adding the Failover Cluster feature	21
Configuring failover cluster nodes	22
Installing Storage Foundation Cluster Option for Microsoft Failover Cluster	22
About using the SFW Fast Failover Configuration utility for enabling Hyper-V live migration support	27
Configuring Fast Failover using the SFW Fast Failover Configuration Utility	28
Unconfiguring Fast Failover using the SFW Fast Failover Configuration Utility	31

	Reconfiguring Fast Failover using the SFW Fast Failover Configuration Utility	32
	Configuring the SFW storage	33
	Creating dynamic cluster disk groups	33
	Creating dynamic volumes	35
	Managing disk groups and volumes	37
	Adding the Volume Manager Disk Group (VMDg) resource	38
	Creating a virtual machine service group	40
	Setting the dependency of the virtual machine on the VMDg resource	41
Chapter 3	Administering storage migration for SFW and Hyper-V virtual machine volumes	43
	About storage migration	44
	About performance tunables for storage migration	44
	Setting performance tunables for storage migration	44
	About performing online storage migration	46
	Storage migration limitations	47
	About changing the layout while performing volume migration	47
	Migrating volumes belonging to SFW dynamic disk groups	47
	Migrating volumes belonging to Hyper-V virtual machines	49
	Migrating data from SFW dynamic disks of one enclosure to another	52
	Converting your existing Hyper-V configuration to live migration supported configuration	54
Chapter 4	Optional Storage Foundation for Windows features for Hyper-V environments	57
	About using optional Storage Foundation for Windows features in the Hyper-V parent	57
	Dynamic Multi-Pathing for the virtual environment	58
	Replicating virtual machines	59
	Virtual machine volume snapshots	60
	Campus clusters	61

Section 2	Veritas Cluster Server for Windows (VCS)	63
Chapter 5	Overview of the Disaster Recovery for Hyper-V solution	65
	About wide-area disaster recovery for Microsoft Hyper-V	65
	Advantages of Disaster Recovery Manager for Microsoft Hyper-V	66
	About the Disaster Recovery Manager for Microsoft Hyper-V configuration	66
	How disaster recovery with Disaster Recovery Manager works	68
Chapter 6	Deploying Hyper-V disaster recovery	71
	Requirements for Disaster Recovery Manager	71
	Ports used by Disaster Recovery Manager	75
	Workflow for deploying Hyper-V disaster recovery	76
	Configuration requirements for the DR site disaster recovery cluster	78
	Setting up the hardware replication for disaster recovery	78
	Prerequisites for EMC SRDF replication with Disaster Recovery Manager	79
	Prerequisites for Hitachi TrueCopy replication with Disaster Recovery Manager	80
	Preparing a VM for Hyper-V DR installation	81
	Installing the Disaster Recovery Manager for Hyper-V	82
	Configuring disaster recovery using the Disaster Recovery Configuration Wizard for Microsoft Hyper-V	83
	Modifying the application VM configuration after configuring disaster recovery	86
	Connecting to the Disaster Recovery Manager with the Java Console	86
	Manually failing over the VMs between sites	89
	Bringing the DR site up if the primary site fails	92
Chapter 7	Hyper-V DR agent	97
	About the Hyper-V DR agents	97
	MonitorVMs agent functions	98
	MonitorVMs agent state definitions	98
	MonitorVMs agent attribute definitions	99
	MonitorVMs agent resource type definition	99

Storage Foundation for Windows (SFW)

- Chapter 1. Introduction to Storage Foundation for Windows solutions for Hyper-V environments
- Chapter 2. Live migration support for SFW dynamic disk group
- Chapter 3. Administering storage migration for SFW and Hyper-V virtual machine volumes
- Chapter 4. Optional Storage Foundation for Windows features for Hyper-V environments

Introduction to Storage Foundation for Windows solutions for Hyper-V environments

This chapter includes the following topics:

- [Advantages of running Storage Foundation for Windows in the Hyper-V parent](#)
- [How Storage Foundation for Windows manages storage for virtual machines](#)
- [Related documentation on Storage Foundation for Windows](#)

Advantages of running Storage Foundation for Windows in the Hyper-V parent

Veritas Storage Foundation for Windows (SFW) is a host-level volume manager that provides a means to virtualize storage seen by the host it runs on. SFW provides central-point control of that storage space.

By running SFW in the Hyper-V parent partition, SFW features and functionality extend to virtual machines (VMs), offering benefits that would otherwise be unavailable at the guest level.

See [“How Storage Foundation for Windows manages storage for virtual machines”](#) on page 14.

SFW has added the following features and functionality specifically to support Hyper-V VMs when running SFW in the Hyper-V parent:

- SFW live migration support

You can configure the SFW storage on which the VMs reside to support VM live migration between nodes of a Microsoft failover cluster.

Detailed instructions are available on how to implement live migration for VMs on SFW storage.

- SFW storage migration for VMs

The SFW storage migration feature enables you to view and select VMs to migrate to different storage.

Detailed instructions are available on how to implement the storage migration solution for VMs.

SFW also offers advanced features and functionality, such as multi-pathing, replication, and snapshots, which further extend the capabilities of Windows in the datacenter. More information is available on how to use the following features and the benefits they provide in a Hyper-V environment:

- Using Dynamic Multi-pathing (DMP) to provide failover and load-balancing to the LUNs that host the VMs in the child partition (DMP Device Specific Modules option)
- Replicating VMs between sites (Veritas Volume Replicator option)
- Maintaining Quick Recovery snapshots of the VMs (FlashSnap option)

See [“About using optional Storage Foundation for Windows features in the Hyper-V parent”](#) on page 57.

How Storage Foundation for Windows manages storage for virtual machines

In virtual environments, managing the storage that is used by guests is not an easy task. Typically, the guest is separated from the physical storage. Veritas Storage Foundation for Windows (SFW) provides several solutions to make it easier to manage storage requirements for virtual machines.

With Hyper-V, guests reside on virtual hard disk (VHD) files, which in turn are located on volumes that reside on physical storage. Direct access to those volumes or the LUNs they reside on is not available from the guest. The VHD files are provisioned by the parent on storage accessed by the parent partition. As storage needs change in the guest VHDs, they may require additional space. It can be difficult to effectively manage space requirements or to relocate a guest from one storage location to another.

Running Veritas Storage Foundation for Windows (SFW) in the parent provides the following storage management solutions for VHDs:

- The SFW storage migration feature enables you to view and select VMs to migrate to different storage.
 For details on using SFW for migrating VM to new storage, see the following: See “[About storage migration](#)” on page 44.
- SFW allows for dynamically growing the volumes that host the guest VHDs. As SFW allows for growth of all volume types, the volumes that host the VHD files can be configured for performance via RAID-5, striping or mirrored-stripes.
- In environments using thin provisioned storage, SFW can be configured to automatically grow volumes based on user- defined space thresholds and policies that set the amount to grow the volumes by and whether that growth should be restricted or unrestricted. This counters the effects of NTFS uncontrolled growth tendencies in a thin environment, by allowing the creation of small volumes on the thin storage, which will grow automatically as needed, triggering corresponding growth in the hardware.
- As a host-level volume manager, SFW also allows for mirroring volumes across arrays and, with its support for dynamic disk operations in a cluster, the creation of stretch or campus clusters.

Related documentation on Storage Foundation for Windows

This guide covers information specific to deploying Microsoft Hyper-V virtual machines with SFW.

The following table describes related documentation on SFW.

Table 1-1 Related documentation on SFW solutions

For information about	Refer to
Installation information	<i>Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide</i>
Information on all SFW features	<i>Veritas Storage Foundation Administrator's Guide</i>
Information on Veritas Volume Replicator (VVR)	<i>Veritas Volume Replicator Administrator's Guide</i>

Table 1-1 Related documentation on SFW solutions (*continued*)

For information about	Refer to
Implementing SFW snapshot solutions for applications	<p>SFW Quick Recovery solutions guides. Guides are available for:</p> <ul style="list-style-type: none"> ■ Microsoft Exchange 2007 and 2010 ■ Microsoft SQL 2005, 2008, and 2008 R2 ■ Microsoft SharePoint 2007 ■ Enterprise Vault
Implementing Microsoft cluster solutions with SFW for applications on physical machines	<p>SFW Microsoft Clustering solutions guides. Guides are available for:</p> <ul style="list-style-type: none"> ■ Microsoft Exchange 2007 ■ Microsoft SQL 2005, 2008, and 2008 R2 <p>In addition, the following guide provides general guidelines for using Microsoft clustering with SFW storage for other applications or server roles:</p> <p><i>Veritas Storage Foundation and High Availability Solutions, Solutions Guide</i></p>

Live migration support for SFW dynamic disk group

This chapter includes the following topics:

- [Tasks for deploying live migration support for Hyper-V virtual machines](#)
- [About implementing Hyper-V virtual machine live migration on SFW storage](#)
- [Requirements for Storage Foundation for Windows \(SFW\)](#)
- [Installing Windows Server 2012](#)
- [Preparing the host machines](#)
- [Installing Storage Foundation Cluster Option for Microsoft Failover Cluster](#)
- [About using the SFW Fast Failover Configuration utility for enabling Hyper-V live migration support](#)
- [Configuring the SFW storage](#)
- [Creating a virtual machine service group](#)
- [Setting the dependency of the virtual machine on the VMDg resource](#)

Tasks for deploying live migration support for Hyper-V virtual machines

To deploy SFW Hyper-V virtual machine live migration on the host side, perform the following tasks in the sequence shown.

Table 2-1 Process for configuring SFW Hyper-V virtual machine live migration

Action	Description
Review the requirements	See “Requirements for Storage Foundation for Windows (SFW)” on page 20.
Install Windows Server 2008 R2	
Prepare the host machines	<ul style="list-style-type: none"> ■ Adding the Hyper-V role ■ Adding the Failover Cluster feature on the host side See “Adding the Failover Cluster feature” on page 21. ■ Configuring failover cluster nodes on the host side See “Configuring failover cluster nodes” on page 22.
Install the SFW Microsoft Failover Cluster option	See “Installing Storage Foundation Cluster Option for Microsoft Failover Cluster” on page 22.
Configure SFW support for Hyper-V live migration using the SFW Fast Failover Configuration Utility from the Solutions Configurations Center (SCC)	<p>See “About using the SFW Fast Failover Configuration utility for enabling Hyper-V live migration support” on page 27.</p> <p>See “Configuring Fast Failover using the SFW Fast Failover Configuration Utility” on page 28.</p> <p>See “Unconfiguring Fast Failover using the SFW Fast Failover Configuration Utility” on page 31.</p> <p>See “Reconfiguring Fast Failover using the SFW Fast Failover Configuration Utility” on page 32.</p>

Table 2-1 Process for configuring SFW Hyper-V virtual machine live migration (continued)

Action	Description
Configure the storage	<p>Use the Veritas Enterprise Administrator (VEA) console to create disk groups and volumes.</p> <p>Note: Setting up a Microsoft failover cluster creates physical disk resources for all the basic disks on the shared bus. To use these disks when you create your SFW cluster disk groups, you must first remove the physical disk resources from the cluster. Otherwise, a reservation conflict occurs.</p> <ul style="list-style-type: none"> ■ Creating dynamic cluster disk groups See “Creating dynamic cluster disk groups” on page 33. ■ Creating dynamic volumes See “Creating dynamic volumes” on page 35. ■ See “Managing disk groups and volumes” on page 37. ■ Adding a Volume Manager Disk Group (VMDg) resource See “Adding the Volume Manager Disk Group (VMDg) resource” on page 38.
Create a virtual machine service group	See “Creating a virtual machine service group” on page 40.
Set the dependency of the virtual machine on the VMDg resource	See “Setting the dependency of the virtual machine on the VMDg resource” on page 41.

About implementing Hyper-V virtual machine live migration on SFW storage

This chapter provides steps for configuring SFW Microsoft Hyper-V virtual machine live migration on a highly available failover cluster disk group resource.

Live Migration of a Hyper-V virtual machine (VM) is achieved through the use of Windows Server 2012 Failover Cluster feature. Live Migration significantly increases availability of the virtual machines during planned and unplanned downtime.

Live migration produces significantly less downtime for the virtual machines that are being migrated. Users can have uninterrupted access to the migrating virtual machine. The guest operating system in the migrating virtual machine is unaware that the migration is taking place. In addition, physical host maintenance can be carried out with no effect on virtual machine availability; this maintenance can occur during normal business hours.

You can configure the SFW storage on which the virtual machine (VM) resides to support VM live migration between nodes of a Microsoft failover cluster.

Note: If you are upgrading SFW, and there is a disk group resource already configured in the cluster, then after the upgrade you must take the disk group resource offline and then restart the cluster service. **FastFailOver** attribute for the disk group resource is displayed on the Failover Cluster Manager console. To successfully set it to **True**, configure SFW Fast Failover Configurations for Live Migration support which is between the cluster nodes, using the SFW Fast Failover Configuration Utility through the Solution Configuration Center (SCC).

See [“About using the SFW Fast Failover Configuration utility for enabling Hyper-V live migration support”](#) on page 27.

See [“Adding the Volume Manager Disk Group \(VMDg\) resource”](#) on page 38.

Requirements for Storage Foundation for Windows (SFW)

Before installing Veritas Storage Foundation for Windows (SFW), review the following requirements:

See [“Converting your existing Hyper-V configuration to live migration supported configuration”](#) on page 54.

- Review the general installation requirements for installing SFW in the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Memory must be a minimum 4 GB of RAM per server for SFW.
- Processor can be either a x64 architecture-based computer with Intel processor that supports Intel Extended Memory 64 Technology (Intel EM64T) or an AMD processor that supports the AMD64 platform; Intel Itanium family IA64 processors are not supported.
- Disk partitions must be formatted for the NTFS file system.
- Memory must be a minimum 4 GB of RAM per server for SFW HA.

- Do not install SFW on servers that are assigned the role of a domain controller. Configuring a cluster on a domain controller is not supported.
- You must be a domain user.
You must be a member of the local Administrators group on all nodes where you are installing.
- A minimum of two NICs are required. One NIC will be used exclusively for private network communication between the nodes of the cluster. The second NIC will be used for both private cluster communications and for public access to the cluster. Symantec recommends three NICs.
- Ensure that all systems used for a highly available solution have a shared storage.
- At least two systems are required for setting up Microsoft failover cluster nodes between host machines.
- Ensure that your firewall settings allow access to ports used by SFW wizards and services. For a detailed list of ports and services used by SFW, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.

Installing Windows Server 2012

Install Windows Server 2012 on the desired host systems. Refer to Microsoft documentation for details.

Preparing the host machines

For virtual machine live migration support, you need to add the Hyper-V role and configure a failover cluster on your host machines. Perform the following tasks in the order shown.

Adding the Hyper-V role

After installing Windows Server 2012, the next step is to add the Hyper-V role to enable the live migration feature. To enable the Hyper-V role, refer to Microsoft Hyper-V documentation for details.

Adding the Failover Cluster feature

Install the Microsoft Failover Cluster feature on all host systems by using the Add Features option from the Server Manager. Refer to Microsoft documentation for details on installing the Failover Cluster feature.

If you are installing the Microsoft Failover Clustering feature on a server on which Veritas Storage Foundation for Windows is already installed, then restart Veritas Enterprise Administrator Service (vxsvc) manually. Issue the following CLI commands to restart the vxsvc service:

- `net stop vxsvc`
- `net start vxsvc`

Configuring failover cluster nodes

Configure a failover cluster on all of your host nodes. Refer to Microsoft documentation for details on how to add a failover cluster node. Ensure that you have fulfilled the clustering network requirements before you start creating the failover cluster nodes.

Verify that you have at least three network adapters (two NICs exclusively for the private network and one for the public network). However, when using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.

Installing Storage Foundation Cluster Option for Microsoft Failover Cluster

Administrative privileges are required on both host and guest operating systems for installing the Veritas Storage Foundation for Windows Cluster Option for Microsoft Failover Cluster.

Before you install SFW, you must install the Microsoft Failover Cluster feature and configure failover cluster nodes on all the systems that are part of the live migration configuration.

Installing SFW requires a restart, but a restart on the active cluster node causes it to fail over. Hence, it is advisable to use a "rolling install" procedure to install SFW first on the inactive cluster node. Then move the cluster resources to the other node and install on the inactive node.

During SFW installation using the product installer, make the following selections:

- Select Storage Foundation for Windows as the product to install.
- When selecting the available options from the server components, ensure that you select the **Cluster Option for Microsoft Failover Cluster** option.
 - Leave the client components selected (the default).

During installation, the installer will display a message box about Quorum Arbitration. The Quorum Arbitration time settings are adjusted to ensure optimal functionality of a dynamic quorum resource on a mirrored dynamic volume.

The quorum arbitration minimum and maximum time settings are used to set the limits of the time period that Microsoft clustering allows for quorum arbitration. Quorum arbitration is the process that occurs when the controlling node of the cluster is no longer active and other nodes of the cluster attempt to gain control of the quorum resource and thus control of the cluster. Refer to the *Veritas Storage Foundation for Windows Administrator's Guide* for information on the settings.

For additional details on using the product installer or command line installation, see the *SFW HA Installation and Upgrade Guide*.

To install SFW/Cluster Option for Microsoft Failover Cluster

- 1 Insert the software disc containing the installation package into your system's disc drive or download the installation package from the following location:
<https://fileconnect.symantec.com>
- 2 Allow the autorun feature to start the installation or double-click **Setup.exe**.
The CD browser appears.

Note: If you are installing the software using the product software disc, the CD browser displays the installation options for all the products specified earlier. However, if you are downloading the installation package from the Symantec Web site, the CD browser displays the installation options only for the product to be installed.

3 Click to download the required contents.

Note: The client components are installed by default along with the server components. However, on a server core computer, the client components will not be installed.

Veritas Storage Foundation	Click to install the server components for Storage Foundation for Windows.
Late Breaking News	Click to access the latest information about updates, patches, and software issues regarding this release.
Windows Data Collector	Click to verify that your configuration meets all the software and hardware requirements.
SORT	Click to access the Symantec Operations Readiness Tools (SORT) site. In addition to the product download you can also download the custom reports about your computer and Symantec enterprise products, a checklist providing configuration recommendations, and system and patch requirements to install or upgrade your software.
Browse Contents	Click to view the software disc contents.
Technical Support	Click to contact Symantec Technical Support.

4 On the Welcome panel, review the list of prerequisites and click **Next**.

5 On the License panel, read the license terms, select **I accept the terms of License Agreement**, and then click **Next**.

6 On the System Selection panel, select the systems and the desired Installation and Product options:

You can select the systems in one of the following ways:

- In the System Name or IP text box, manually type the system name or its IP address and click Add.

Note: The wizard does not support the Internet Protocol version 6. To add the systems having Internet Protocol version 6, you must type the system name.

The local host is populated by default.

- Alternatively, browse to select the systems.
The systems in the domain to which you are logged on are listed in the **Available Systems** list. Select one or more systems and click the right arrow to move them to the **Selected Systems** list.
Click **OK**.
The selected systems are validated and details are displayed in **Verification Details**. Select a system to review the details.
To select the installation and product options, perform the following tasks on each of the selected system.

Note: To apply the selection to multiple systems, select the system for which you have selected the installation and product options and then click **Apply to multiple systems**.

- The product is installed by default at the %ProgramFiles%\Veritas location. To customize the installation directory, click **Browse** and select a location of your choice.
Click **OK**.

Note: The installation directory is selected by default on systems where the product is being upgraded.

- Select the required license type from the **License key** drop-down list.

Note: The default license type is **Keyless**.

If you select the **Keyless** license type, all the available product options are displayed and are selected by default.

If you select **User entered license key** as your license type, the License Details panel appears by default. On the **License Details** panel, enter the license key and then click **Add**. You can add multiple licenses for the various product options you want to use.

Validation check is done for the entered license keys. If validation fails, an error message is displayed.

After successful validation, click **OK**.

- From the list of product options, select the **Cluster Option for Microsoft Failover Cluster** which provides support for Microsoft Failover Cluster.

7 On the System Selection panel, click **Next**.

All the selected systems must pass the validation check. In case the validation checks have failed on any of the system, review the details and resolve the issue.

Click **Re-verify** to run the validation check on the system again.

8 On the Pre-install Summary panel, review the summary and click **Next**.

9 The **Automatically reboot systems after installer completes operation** check box is selected by default. This option restarts all the selected remote systems immediately after installation completes on selected systems.

However, if you do want to initiate the auto restart option at this stage, uncheck the checkbox **Automatically reboot systems after installer completes operation**.

10 On the Installation panel, review the progress of installation.

Click **Next** after the installation completes.

If installation is not successful on any of the selected systems, then a failed installation status is shown.

11 On the Post-install Summary panel, review the installation result and click **Next**.

Refer to the log file for details, if installation has failed on any of the selected system. You may need to reinstall the software.

12 On the Summary page, click **Finish**.

If you selected the auto restart as shown in step **9**, a confirmation message to restart the local system appears.

Click **Yes** to restart immediately.

Click **No** to restart later.

However, if the auto restart option was not selected in step **9**, then you must ensure to manually restart the selected systems.

If you are installing the Microsoft Failover Clustering feature on a server on which Veritas Storage Foundation for Windows is already installed, then restart Veritas Enterprise Administrator Service (vxsvc) manually. Issue the following CLI commands to restart the vxsvc service:

- `net stop vxsvc`

- `net start vxsvc`

About using the SFW Fast Failover Configuration utility for enabling Hyper-V live migration support

After configuring the cluster, use the SFW Fast Failover Configuration Utility to set up the SFW components for enabling support for live migration. SFW support for Hyper-V live migration is supported on a network within a Microsoft failover cluster. This utility also enables fast failover support for SFW disk group resources in the cluster. After you run this utility you can set the FastFailover attribute of the SFW VMDg resources.

Launch this utility from the Solutions tab in the Solutions Configuration Center (SCC). The SFW Fast Failover Configuration Utility can be run from any node of the Microsoft failover cluster (FOC) which has Hyper-V role added to it.

Use the SFW Fast Failover Utility for the following scenarios:

- After configuring the Microsoft failover cluster for the first time, or after adding a new node to the FOC cluster.
- For enabling Fast Failover for disk group resources.
- After removing a node from the FOC cluster.
- Reconfiguring live migration support in case of any change in the network properties.

Refer to the following topics:

See [“Configuring Fast Failover using the SFW Fast Failover Configuration Utility”](#) on page 28.

See [“Unconfiguring Fast Failover using the SFW Fast Failover Configuration Utility”](#) on page 31.

See [“Reconfiguring Fast Failover using the SFW Fast Failover Configuration Utility”](#) on page 32.

Note: In synchronous mode of replication, if fast failover is set, then the RVG cannot be stopped and started when a disk group fails over to another node. If the RLINK is in hard synchronous mode, it may not be connected when the volume arrives, and the I/Os may fail. In such case, the Event Viewer displays NTFS errors and file system reports the volume as RAW. Therefore, fast failover is not supported if the RLINK is in hard synchronous mode.


Configuring Fast Failover using the SFW Fast Failover Configuration Utility

Configure fast failover using the SFW Fast Failover Configuration Utility.

Note: You must perform this step before adding Volume Manager Disk Group (VMDg) resources to the cluster.

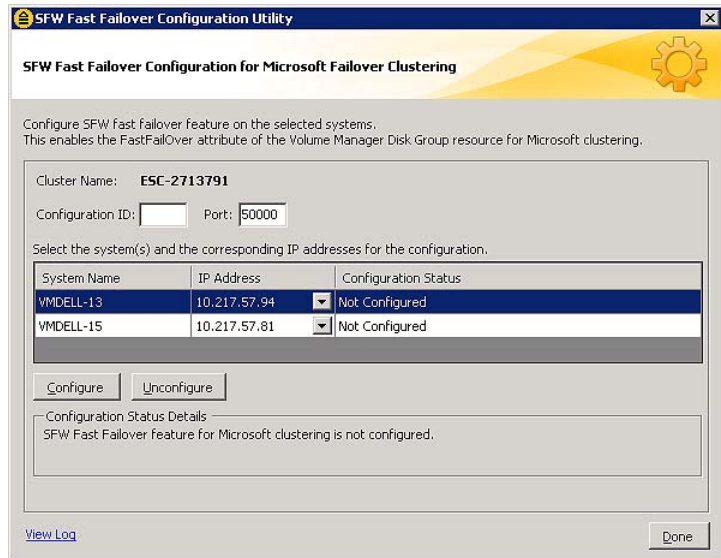
To configure live migration support using the SFW Fast Failover Configuration Utility

- 1 Click **Start > Run**, type **scc**, and then click **OK** to launch the Solutions Configuration Center (SCC).
- 2 From the Solutions view in the SCC, launch the **SFW Fast Failover Configuration Utility**.



- 3 On the SFW Fast Failover Configuration Utility page, complete the following:

About using the SFW Fast Failover Configuration utility for enabling Hyper-V live migration support

**Field****Description**

Configuration ID

Enter a unique cluster configuration ID of your choice. The value can range from 0 to 65535.

Port

Specify a valid port number.

The port values can range from 49152 to 65535.

Ensure that the UDP ports used by you are unblocked if using a firewall.

System Name

This table lists all the hosts in a cluster configuration as well as hosts which are already part of live migration configuration.

Select the required cluster nodes from this column.

Note: Please select at least two systems from the System list to configure live migration support.

The SFW Fast Failover Configuration Utility will configure and unconfigure live migration support for selected systems.

IP Address	<p>Corresponding IP address for the selected host or hosts is displayed.</p> <p>If multiple addresses for a selected host are configured, then a drop-down list is displayed. Select an IP address within the same cluster network.</p>
Configuration Status	<p>The following configuration status is displayed:</p> <ul style="list-style-type: none">■ All the nodes which are already part of Fast Failover configuration are shown as Configured.■ For newly added node in a cluster, the status is displayed as Not Configured.■ For invalid configuration, status is displayed as Invalid Configuration.■ Not Supported status is displayed when SFW Cluster option for Microsoft Failover Cluster is not installed.
Configure	<p>Click on this button to configure SFW Fast Failover Configuration Utility for a selected system that has a Not Configured status.</p> <p>Select a system and click on Configure to enable Live Migration support on it.</p>
Configuration status details	<p>Displays fast failover configuration status for selected systems.</p>

Click **Done** to configure live migration support for the selected systems and then exit the SCC console.

Note: If a cluster node crashes or shuts down abruptly, then it is noticed that on subsequent restart of the other remaining cluster nodes, the SFW Fast Failover Configuration Utility shows the crashed node as **Invalid Configuration**.

In such cases, the SFW messaging for Live Migration support will not work between the remaining nodes and the VMDg **FastFailOver** attribute cannot be set to **True** for any new VMDg resource. To resolve this issue, it is recommended to first **Unconfigure** and then **Configure** the remaining cluster nodes using the SFW Fast Failover Configuration Utility through the Solutions Configuration Center (SCC).

Unconfiguring Fast Failover using the SFW Fast Failover Configuration Utility

To remove a node from a cluster when the status is shown as **Configured** or **Invalid Configuration**, use the SFW Fast Failover Configuration Utility to unconfigure live migration support.

Perform the following steps to unconfigure SFW Hyper-V live migration support for selected cluster nodes.

To unconfigure fast failover using the SFW Fast Failover Configuration Utility

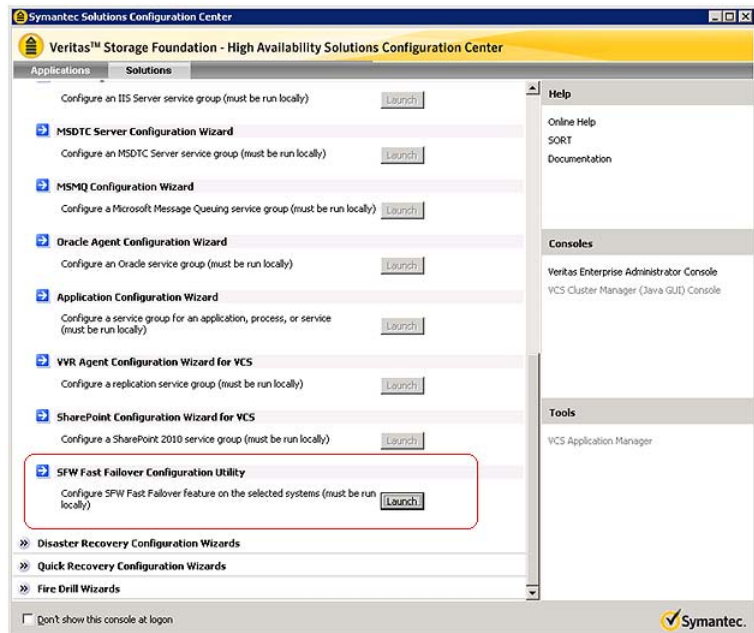
- 1 Right-click a VMDg resource on which the virtual machine is dependent from the Failover Cluster Manager console. Select **Properties** tab from the context menu.

From the Properties window, select the **Properties** tab and click to edit the VMDg **FastFailOver** attribute value to **FALSE**.

See “[Adding the Volume Manager Disk Group \(VMDg\) resource](#)” on page 38.

- 2 Now, using the Solutions Configuration Center (SCC) unconfigure live migration support for the cluster nodes that show a status as **Configured** or **Invalid Configuration**.

Click **Start > Run**, type **scc**, and click **OK** to launch the Solutions Configuration Center (SCC).



- 3 From the **Solutions** tab in the SCC, launch the **SFW Fast Failover Configuration Utility**.
- 4 Select **Configured** or **Invalid Configuration** cluster nodes that you need to remove from fast failover configuration.

If total number of cluster nodes is two, you can unconfigure or configure live migration support for these two cluster nodes. However, if total number of cluster nodes displayed is three, then select at least two systems to unconfigure fast failover support.
- 5 Now click the **Unconfigure** button to unconfigure live migration support for the selected cluster nodes.

Additionally, you can also use **Alt+U** shortcut keys to unconfigure live migration support.
- 6 Validate that unconfiguring live migration support on selected cluster nodes is successful in the **Configuration Status Details** panel.
- 7 Click **Done** to unconfigure live migration support, and then exit the SCC console.

Reconfiguring Fast Failover using the SFW Fast Failover Configuration Utility

In case of a change in network configuration (IP or port changes), use the SFW Fast Failover Configuration Utility to unconfigure and reconfigure fast failover support.

Note: If SFW Hyper-V live migration configuration is not reconfigured, then cluster nodes that display **Invalid Configuration** status fail to receive Read-Only import messages and hence, will not have an up-to-date configuration. In such cases, the eventual SFW-Hyper-V live migration will not be instantaneous and will be slow.

To unconfigure and configure SFW Fast Failover support again on selected cluster nodes, refer to the following sections:

See [“Unconfiguring Fast Failover using the SFW Fast Failover Configuration Utility”](#) on page 31.

See [“Configuring Fast Failover using the SFW Fast Failover Configuration Utility”](#) on page 28.

Configuring the SFW storage

You use Veritas Storage Foundation for Windows to create dynamic cluster disk groups and volumes for a cluster environment. You then add Volume Manager Disk Group resources to the failover cluster.

See [“Creating dynamic cluster disk groups”](#) on page 33.

See [“Creating dynamic volumes”](#) on page 35.

See [“Managing disk groups and volumes”](#) on page 37.

See [“Adding the Volume Manager Disk Group \(VMDg\) resource”](#) on page 38.

Creating dynamic cluster disk groups

You create a dynamic cluster disk group with volumes on shared storage so that they can be shared between nodes in the cluster. Part of the process of creating a dynamic disk group is assigning it a name. You must choose a name that is unique to your environment. Make note of this name, as it will be required later.

To create dynamic cluster disk groups, use the Veritas Enterprise Administrator (VEA). The VEA can be invoked on one of the servers and can be used to connect to all the other servers. However, VEA can also be launched on client system and can be used to manage all the servers remotely.

Note: Setting up a Microsoft failover cluster creates physical disk resources for all the basic disks on the shared bus. To use these disks when you create your SFW cluster disk groups, you must first remove the physical disk resources from the cluster. Otherwise, a reservation conflict occurs. After creating the SFW cluster disk groups, you will add Volume Manager Disk Group resources to the cluster, instead of physical disk resources.

Note that dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

Note: For SFW Hyper-V live migration support, only one virtual machine (VM) per disk group is a mandatory requirement.

Note: If User Access Control (UAC) is enabled on Windows Server 2008 systems, then you cannot log on to VEA GUI with an account that is not a member of the Administrators group, such as a guest user. This happens because such user does not have the "Write" permission for the "Veritas" folder in the installation directory (typically, C:\Program Files\Veritas). As a workaround, an OS administrator user can set "Write" permission for the guest user using the Security tab of the "Veritas" folder's properties.

To create a dynamic (cluster) disk group

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.

To connect to the local system, select **localhost**.

Provide the user name, password, and domain if prompted.

- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right-click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.
- 6 Provide information about the cluster disk group:
 - Enter the name of the disk group (for example, DG1).
 - Check the **Create cluster group** check box.
 - Select the appropriate disks in the **Available disks list**, and use the **Add** button to move them to the **Selected disks list**.
 - Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier.
For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

Click **Next** to accept the confirmation screen with the selected disks.

- 7 Click **Finish** to create the new disk group.

Creating dynamic volumes

This section will guide you through the process of creating a volume on a dynamic disk group.

Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

To create dynamic volumes

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.

To connect to the local system, select **localhost**.

Provide the user name, password, and domain if prompted.

- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.

You can right-click the disk group you have just created, for example **DG1**.

- 5 At the New Volume wizard opening screen, click **Next**.
- 6 Select the disks for the volume.

Make sure the appropriate disk group name appears in the **Group name** drop-down list.

- SFW automatically selects the disks based on the following criteria:
 - Their port assignment (disks with two different ports are selected). Note that in the list of available disks, the entry after each disk name starts with the port number. For example, the “P3” in the entry P3C0T2L1 refers to port 3.
 - Amount of available space on the disks. SFW will pick two disks (one from each array) with the most space.

- To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list.
- You may also check **Disable Track Alignment** to disable track alignment for the volume.
Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.

Click **Next**.

7 Specify the volume attributes and complete the following.

- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
- Provide a size for the volume. If you click the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
- Select a layout type.
- If you are creating a striped volume, the Columns and Stripe unit size boxes need to have entries. Defaults are provided.
- To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
- In the Mirror Info area, select the appropriate mirroring options.
- Verify that **Enable logging** is not selected.

Click **Next**.

8 Assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.

- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.
- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.

Click **Next**.

9 Create an NTFS file system.

- Make sure the **Format this volume** checkbox is checked and click **NTFS**.
- Select an allocation size or accept the default.
- The file system label is optional. SFW makes the volume name the file system label.

- Select **Perform a quick format** if you want to save time.
- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.

Click **Next**.

10 Click **Finish** to create the new volume.

11 Repeat these steps to create additional volumes.

Create the cluster disk group and volumes on the first node of the cluster only.

Managing disk groups and volumes

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Importing a disk group and mounting a volume

Use the Veritas Enterprise Administrator (VEA) console to import a disk group and mount a volume.

To import a disk group

- 1** From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2** From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1** If the disk group is not imported, import it.
- 2** To verify if a disk group is imported, from the VEA Console, click the **Diskstab** on the right pane and check if the status is imported.

- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - To assign a drive letter
Select **Assign a Drive Letter**, and select a drive letter.
 - To mount the volume as a folder
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

- 1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2 In the Drive Letter and Paths dialog box, click **Remove**.
Click **OK** to continue.
- 3 Click **Yes** to confirm.
- 4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5 Click **Yes**.

Adding the Volume Manager Disk Group (VMDg) resource

Perform the following steps to add a Volume Manager Disk Group (VMDg) resource.

Note: For SFW Hyper-V live migration support, only one virtual machine (VM) per disk group is a mandatory requirement. If multiple virtual machines reside on the same disk group, then before configuring live migration, use the Storage Migration wizard to migrate virtual hard disks and split the disk group using SFW to create separate disks groups.

See “[Converting your existing Hyper-V configuration to live migration supported configuration](#)” on page 54.

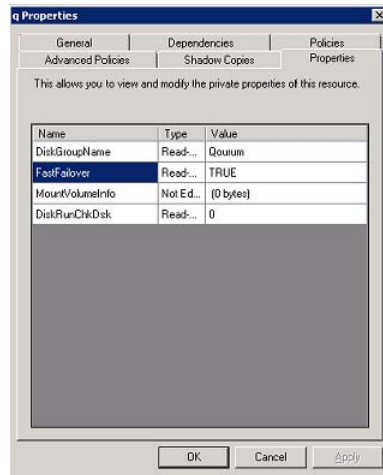
To add a Volume Manager Disk Group (VMDg) resource

- 1 Click **Start > Administrative Tools > Failover Cluster Manager** to open the failover cluster snap-in.
- 2 Right-click **Services and Applications**. Select **More Actions > Create Empty Service or Application**. This creates a service group, for example, **SG1**.
- 3 Right-click the new service group and select **Add a Resource > More Resources > Add a Volume Manager Disk Group** from the context menu.

A new Volume Manager Disk Group (VMDg) resource is created with a default name

- 4 Right-click the VMDg resource and select **Properties**.

Complete the following on the **Properties** window:



- Select the **General** tab to change the default name of the **New Volume Manager Disk Group** to a name of your choice. Say, for example: **VMDg1**
- Now select the **Properties** tab and perform the following steps:
 - In the **DiskGroupName** box enter the dynamic cluster disk group name created earlier in this document. Say, for example **GuestDG1**.
 - Edit the **FastFailOver** attribute value to **TRUE**. Displayed default value is **FALSE**.

Note: To enable live migration, you must set the **FastFailOver** attribute to **TRUE** for all VMDg resources in a cluster.

- 5 Right-click the VMDg resource and select **Bring this resource online** option from the center pane of the failover cluster snap-in.

Creating a virtual machine service group

After adding a Volume Manager Disk Group (VMDg) resource, proceed with adding a virtual machine on the active failover cluster node.

Note: Virtual machine and virtual hard disk (.vhd) must be stored in the VMDg resource. This is required to make the virtual machine highly available.

To make the shared storage, i.e, the VMDg resource, available to the virtual machine, you must create the virtual machine on a cluster node that owns the storage. You can either create a new virtual hard disk (.vhd), use an existing .vhd as shown in the procedure below, or you can simply create it later.

To create a virtual machine

- 1 Click **Start > Administrative Tools > Failover Cluster Manager**.

If you are not connected to the cluster node that owns the shared storage connect to it.

- 2 Click on **Service and Applications > Virtual Machine > New > Virtual Machine** from the left pane of the Failover Cluster Manager console.

- 3 The New Virtual Machine Wizard is launched. Review the information on the welcome page.

Click **Next**.

- 4 On the Specify Name and Location page, specify a name for the virtual machine, for example, **VM1**.

- 5 Enable the checkbox **Store the virtual machine in a different location**, and then type the full path or click **Browse** and copy the virtual hard disk (VHD) file to the Volume Manager Disk Group (VMDg1) resource created in earlier section of this document for storing the virtual machine.

See “[Adding the Volume Manager Disk Group \(VMDg\) resource](#)” on page 38.

- 6 On the Memory page, specify the amount of memory required for the operating system that will run on this virtual machine.

- 7 On configure Networking page, select the network adapter configured to use for Hyper-V.
- 8 On the Connect Virtual Hard Disk page, three options are shown:
 - **Create a new virtual hard disk**
 - **Use existing virtual hard drive**
 - **Create virtual hard disk later**Select the required option and give the path to the VHD (.vhd) file on the Volume Manager Disk Group (VMDg) resource.
The VHD (.vhd) file must be stored in a location where the VMDg resource is located.
- 9 Click **Finish** to close the wizard.

Setting the dependency of the virtual machine on the VMDg resource

As the virtual machine configuration and Virtual Hard Disk (VHD) file lie on the VMDg resource, you need to set a dependency of the virtual machine resource on the VMDg resource to make it highly available and live migrate it.

See [“Adding the Volume Manager Disk Group \(VMDg\) resource”](#) on page 38.

See [“Creating a virtual machine service group”](#) on page 40.

Perform the following steps to set a dependency of the VMDg resource on the virtual machine:


To set the dependency of the VMDg resource on the virtual machine

- 1 Right click the VMDg resource from the center pane of the Failover Cluster Manager console. Select **Actions > More Actions > Move this resource to another service group or application** from the menu.

VMDg resource is now moved to the selected virtual machine.

- 2 From the center pane of console, select and expand the virtual machine resource to display the **Virtual Machine Configuration Resource**.

Right-click and select **Properties > Dependency** tab to make this Virtual Machine Configuration Resource dependent on the Volume Manager Disk Group resource (VMDg).



Select the VMDg resource that is displayed and click **Ok** to set the dependency.

Now the virtual machine is ready to be live migrated to other cluster node.

Refer to Microsoft documentation for details regarding live migrating a virtual machine.

Administering storage migration for SFW and Hyper-V virtual machine volumes

This chapter includes the following topics:

- [About storage migration](#)
- [About performance tunables for storage migration](#)
- [Setting performance tunables for storage migration](#)
- [About performing online storage migration](#)
- [Storage migration limitations](#)
- [About changing the layout while performing volume migration](#)
- [Migrating volumes belonging to SFW dynamic disk groups](#)
- [Migrating volumes belonging to Hyper-V virtual machines](#)
- [Migrating data from SFW dynamic disks of one enclosure to another](#)
- [Converting your existing Hyper-V configuration to live migration supported configuration](#)

About storage migration

SFW provides the ability to move volumes to new storage locations via the Storage Migration Wizard. The storage migration feature facilitates moving multiple volumes to different set of disks while the volumes are still online. Volumes associated with a Hyper-V Virtual Machine (VM) or an SFW disk group can be moved in a single administrative operation while the volumes are online without stopping the applications or Hyper-V VMs.

Storage migration provides administrators great flexibility when deploying new arrays or moving to LUNs that are configured for better performance.

Note: Volume layout and site boundary constraints are preserved during storage migration.

About performance tunables for storage migration

Performance tunables are provided to optimize the time taken for the storage migration operation. Using them, you can enhance the storage migration performance, in migrating data from source disks to target disks. These new performance tunables gives you the flexibility to trade between storage migration performance and application performance.

Usage of performance tunables is limited to storage migration operation only.

The following new performance tunables are introduced:

- I/O size
- Task Threads

You can access these tunables from **VEA > Control Panel > Task Settings > Configuration Task Performance Tuning** tab.

For more details about these tunables,

See [“Setting performance tunables for storage migration”](#) on page 44.

Setting performance tunables for storage migration

The performance tunables are set at default values. However you can define these values, to optimize the storage migration performance:

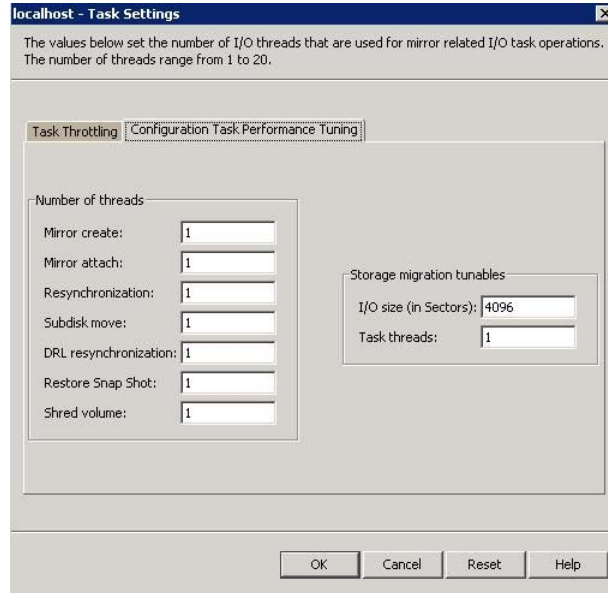
I/O size	<p>To read or write data to the disks during migration.</p> <p>To trade off between application performance and storage migration performance.</p> <p>The I/O size tuning is not applicable for migration of striped volumes, if you do not change the layout.</p> <p>IO size, default value (in sectors) = 4096, Lower Bound = 128, Upper Bound = 16384</p>
Task threads	<p>To perform the subdisk move</p> <p>Can be specified separately for storage migration tasks</p> <p>Operations initiated from array migration and storage migration wizards, use task threads value, specified under storage migration tunables.</p> <p>Task threads, default value = 1, Lower Bound = 1, Upper Bound = 20</p>
Task Throttling	<p>No new parameters are added to the task throttling tab. However, to further optimize the application performance during storage migration, you can modify the parameters for MirrorCreate (in case of volume relay layout) and SDMove (without volume relay layout and array migration). These parameters are available on the Task Throttling tab.</p>

Perform the following steps from the system where the storage (disk group volumes or array enclosure disks) is mounted and online.

To set the storage migration performance tunables

- 1 Launch the VEA console and connect to the local system.
Click **Start > Run**, type **vea-host localhost**, and then click **OK**.
- 2 Click **Control Panel** in the left panel and then double-click **Task Settings** in the right pane.

- 3 On the localhost - Task Settings window, click the **Configuration task performance tuning** tab to view the new performance tunables.



- 4 In the Storage migration tunable box, define appropriate values for the new tunables, I/O size and Task threads, and then click **OK**.

The attribute values take effect immediately and apply to all the storage migration operations run subsequently.

These settings do not apply to storage migration operations that are already in progress.

About performing online storage migration

The storage migration wizard migrates data on the disks from one enclosure to another, by moving volumes and optionally changing their layout. This is achieved in a single administrative operation, while the volumes are online, and without stopping the applications.

The following operations are available to perform online storage migration:

- Migrate volumes belonging to SFW dynamic disk groups.
See [“Migrating volumes belonging to SFW dynamic disk groups”](#) on page 47.
- Migrate volumes belonging to Hyper-V virtual machines.

See “[Migrating volumes belonging to Hyper-V virtual machines](#)” on page 49.

- Migrate data from SFW dynamic disks of one enclosure to another. See “[Migrating data from SFW dynamic disks of one enclosure to another](#)” on page 52.

Storage migration limitations

The limitations while performing storage migration are as follows:

- For migrating volumes, the target disks selected for migration must belong to the same disk group.
- If the Hyper-V virtual machine configuration spans across multiple disk groups, then storage migration should be performed on a per disk group basis.
- RAID-5 volumes are not supported for migration.

About changing the layout while performing volume migration

You can change the volume layout while migrating it, from concatenated to striped or vice versa. In addition, you can also change the number of columns or stripe unit size of a stripe volume.

For a mirrored volume, the specified target layout applies to all the plexes of that particular volume, except the snap ready plexes.

Migrating volumes belonging to SFW dynamic disk groups

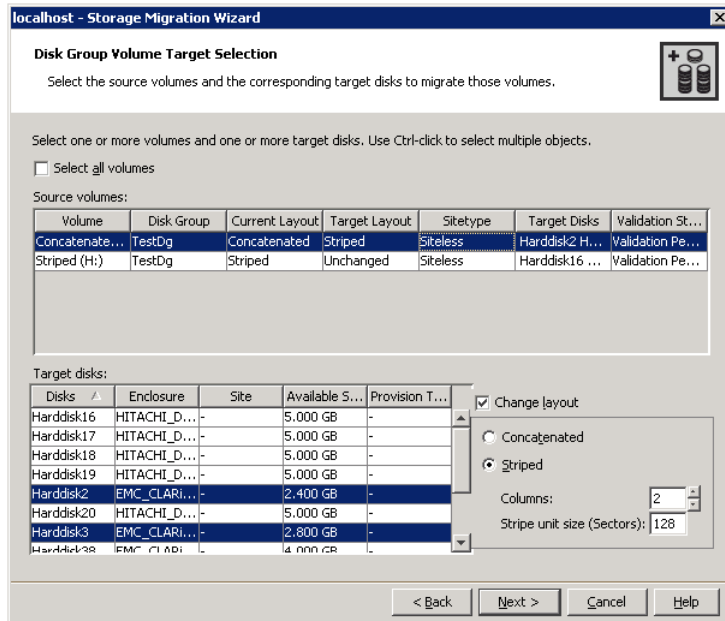
Use this procedure to perform storage migration for SFW dynamic disk groups.

Ensure that the disk groups that you want to migrate are imported on the system where you run this procedure.

To migrate volumes belonging to SFW dynamic disk groups

- 1 Launch the VEA console and connect to the local system.
Click **Start > Run**, type **vea -host localhost**, and then click **OK**.
- 2 In the left panel expand **Disk Groups** and then either right-click and select **Migrate Storage** or from the menu bar click **Actions > Migrate Storage** to start the Storage Migration Wizard.

- 3 Review the prerequisites on the Welcome panel and then click **Next**.
- 4 On the Storage Component Selection panel, Select **Migrate disk group volumes**, and then select the desired disk group from the drop-down list and click **Next**.
- 5 On the Disk Group Volume Selection panel, you can select one or more volumes for online migration. The **Select all volumes** option is selected by default. Clear the **Select all volumes** option and press Ctrl+click if you want to select one or more volumes for online migration. Click **Next**.



- 6 On the Disk Group Volume Target Selection panel, you can select the source volumes and the corresponding target disks to migrate volumes. The **Select all volumes** checkbox is selected by default. Clear the **Select all volumes** checkbox to map target disks to individual volumes and **Change layout** if desired. Press Ctrl+click to select multiple target disks.
 - All selected source volumes can migrate to any target disks if the **Select all volumes** option is selected.
 - If you select the **Select all volumes** option, you will not have the option to change the layout. The layout can be changed only on selecting individual volumes.
 - If you select the **Change layout** checkbox and opt to change the layout of the volume, the target layout column gets populated with the selected layout, otherwise **Unchanged** will be displayed.

Click **Next** to go to the Summary panel.

Unless the validation status of all the volumes shows **Success**, the GUI will not navigate to the next panel. The validation status column will be populated when you click **Next** on the Disk Group Volume Target Selection panel.

- 7 On the Summary panel, the mapping of the source volumes and target disks is shown based on the selection done on the previous panel. The Storage migration summary panel displays the name of the disk group, the source volumes, target layout and the target disks.

If the select all option is selected then the summary is shown in the form of a grouping of source volumes and target disks. If the select all option is unchecked, the individual selected volumes and assigned target disks will be displayed on the summary panel, according to the selection made by you.

The **Save Summary** button on the Summary panel lets you to save user selection summary in an HTML file.

The default summary file name is

```
StorageMigration_Summary_yyyy_MM_dd_HH_mm_ss.html.
```

The default location of the file is %allusersprofile%\veritas.

- 8 Click **Finish** to begin the volume migration.
- 9 Now check the storage migration status (whether successful or failed) by completing the following on the VEA:
 - Click the **Console** tab on the bottom panel of the VEA.
 - Verify that separate Subdisk move tasks are created per subdisk.
 - Click the **Tasks** tab on the VEA to check the tasks progress in the bottom panel of the console.
 - Wait for all Subdisk move tasks to complete.
 - From the **Disk View** tab verify that all selected volumes are now migrated to the selected destination disks.

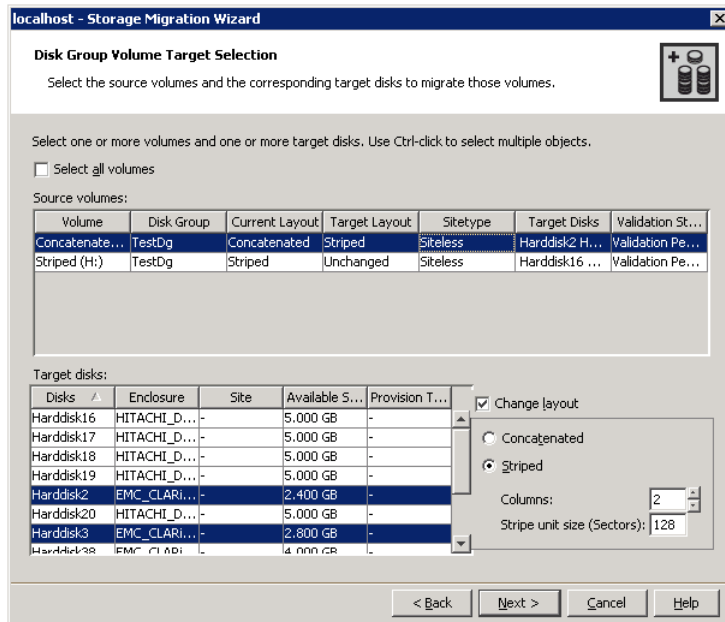
Migrating volumes belonging to Hyper-V virtual machines

Use this procedure to migrate volumes belonging to Hyper-V virtual machines.

Ensure that the disk groups that you want to migrate are imported on the system where you run this procedure.

To migrate volumes belonging to Hyper-V virtual machines

- 1 Launch the VEA console and connect to the local system.
 Click **Start > Run**, type **vea -host localhost**, and then click **OK**.
- 2 Launch the Storage Migration Wizard.
 Right-click **localhost** and then click **Migrate Storage**.
 Or, select **localhost** and then on the menu bar, click **Actions > Migrate Storage**.
- 3 Review the prerequisites on the welcome panel and then click **Next**.
- 4 On the Storage Component Selection panel, select **Migrate volumes used by Microsoft Hyper-V virtual machines**, and select the desired virtual machine from the drop-down list.
- 5 Click **Next** to navigate to the Disk Group Volume Selection panel. You can select one or more volumes for online migration. The **Select all volumes** option is selected by default. Clear the **Select all volumes** option and press Ctrl+click if you want to select one or more volumes for online migration. Click **Next**.



- 6 On the Disk Group Volume Target Selection panel, you can select the source volumes and the corresponding target disks to migrate volumes. The **Select all volumes** checkbox is selected by default. Clear the **Select all volumes**

checkbox to map target disks to individual volumes and **Change layout** if desired. Press Ctrl+click to select multiple target disks.

- All selected source volumes can migrate to any target disks if the **Select all volumes** option is selected.
- If you select the **Select all volumes** option, you will not have the option to change the layout. The layout can be changed only on selecting individual volumes.
- If you select the **Change layout** checkbox and opt to change the layout of the volume, the target layout column gets populated with the selected layout, otherwise **Unchanged** will be displayed.

Click **Next** to go to the Summary panel.

Unless the validation status of all the volumes shows **Success**, the GUI will not navigate to the next panel. The validation status column will be populated when you click **Next** on the Disk Group Volume Target Selection panel.

- 7 On the Summary panel, the mapping of the source volumes and target disks is shown based on the selection done on the previous panel. The Storage migration summary panel displays the name of the disk group, the source volumes, target layout and the target disks.

If the select all option is selected then the summary is shown in the form of a grouping of source volumes and target disks. If the select all option is unchecked, the individual selected volumes and assigned target disks will be displayed on the summary panel, according to the selection made by you.

The **Save Summary** button on the Summary panel lets you save the selection summary in an HTML file.

The default summary file name is

`StorageMigration_Summary_yyyy_MM_dd_HH_mm_ss.html`.

The default location of the file is `%allusersprofile%\veritas`.

- 8 Click **Finish** to begin the volume migration.
- 9 Now check the storage migration status (whether successful or failed) by completing the following on the VEA:
 - Click the **Console** tab on the bottom panel of the VEA.
 - Verify that separate Subdisk move tasks are created per subdisk.
 - Click the **Tasks** tab on the VEA to check the tasks progress in the bottom panel of the console.
 - Wait for all Subdisk move tasks to complete.

- From the **Disk View** tab verify that all selected volumes are now migrated to the selected destination disks.

Migrating data from SFW dynamic disks of one enclosure to another

Use this procedure to migrate data from SFW dynamic disks of one enclosure to another.

To migrate data from SFW dynamic disks of one enclosure to another

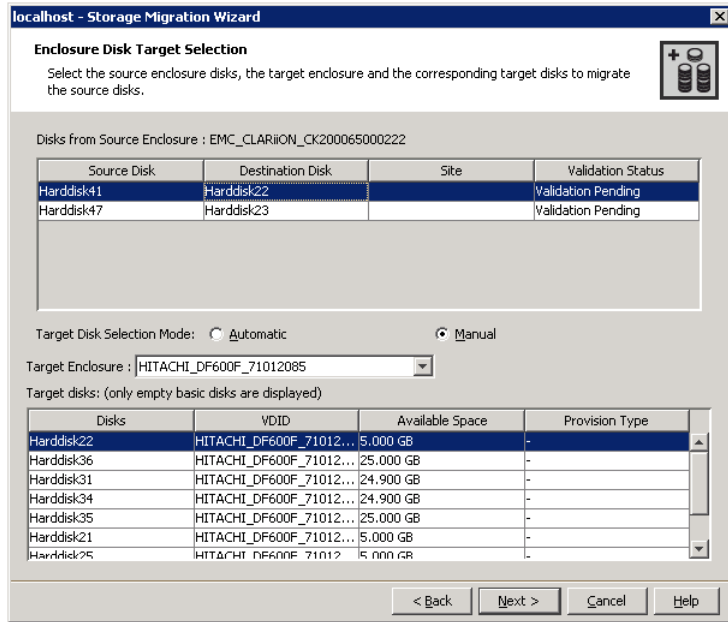
- 1 Launch the VEA console and connect to the local system.
Click **Start > Run**, type **vea -host localhost**, and then click **OK**.
- 2 Launch the Storage Migration Wizard.
Right-click the **Enclosures node**, and then on the menu bar click **Actions > Migrate Storage**.
Or, right-click **Enclosures Node** or an individual enclosure and then select **Migrate Storage**.
- 3 Review the prerequisites on the Welcome panel and then click **Next**.
- 4 On the Storage Component Selection panel, select **Migrate data from an array** and then from the drop-down list choose the source array enclosure.
- 5 Click **Next** to navigate to the Enclosure Disk Selection panel.

The disks may lie on different disk groups. Select the number of required source disks either by manually selecting one or more source disks for migration or by selecting the **Select all disks** option.

Note: Only SFW dynamic disks are available for selection as source disks.

After selecting the source disks click **Next** to continue.

- 6 On the Enclosure Disk Target Selection panel, select the target enclosure, as well as map source disk to the destination disk.



Two options for target disk selection mode are available: **Automatic** and **Manual**.

- In the **Automatic** mode, the source, and target disks are mapped automatically.
- In the **Manual** mode, you can map source disk to target disk manually. Here there is a 1:1 mapping from source to target disk.

Click **Next** to continue after selecting the desired option.

7 The Summary panel shows the summary of the mapping.

The **Save Summary** button on the Summary panel lets you save the selection summary in an HTML file.

The following details are given in the summary:

- **Automatic** mode: Source Enclosure, Target Enclosure, and Source Disks: Target Disks. Here the disk grouping by automatic disk selection is shown.
- **Manual** mode: Source Enclosure, Target Enclosure, and Source Disks: Target Disks. Here the 1:1 mapping of source and target disks is shown.

The default summary file name is

StorageMigration_Summary_YYYY_MM_DD_HH_MM_SS.html.

The default location of the file is %allusersprofile%\veritas.

- 8 Review the summary and click **Finish** to begin the enclosure migration operation.
- 9 Now check the storage migration status (whether successful or failed) by completing the following on the VEA:
 - Click the **Console** tab on the bottom panel of the VEA.
 - Verify that separate Subdisk move tasks are created per subdisk.
 - Click the **Tasks** tab on the VEA to check the tasks progress in the bottom panel of the console.
 - Wait for all Subdisk move tasks to complete.
 - From the **Disk View** tab verify that all selected volumes are now migrated to the selected destination disks.

Converting your existing Hyper-V configuration to live migration supported configuration

Through the Storage Migration Wizard, it is possible to convert your existing Hyper-V Virtual Machine (VM) configuration to a live migration supported configuration (one Hyper-V Virtual Machine per disk group).

Before trying to convert or change your existing Hyper-V Virtual Machine configuration to a Live Migration supported configuration, it is presumed here that you have completed the following procedures already:

- Creating disk groups and dynamic volumes
- Copying Virtual Hard Disks (VHDs) to volumes
- Creating at least two Hyper-V Virtual Machines (VMs)

To convert a Hyper-V Virtual Machine configuration to a live migration supported configuration, you need to perform the following steps:

To convert a Hyper-V configuration to a live migration supported configuration

- 1 To migrate volumes that belong to a Hyper-V Virtual machine to empty disks, launch the Storage Migration Wizard from the VEA GUI by right clicking on **localhost** node. Select the **Migrate Storage** option from the menu.
- 2 On the Source Volume page, select a virtual machine from the **Hyper-V Virtual Machines** drop-down list created by you.

Select all the volumes that belong to the Hyper-V Virtual Machine.

Click **Next**.

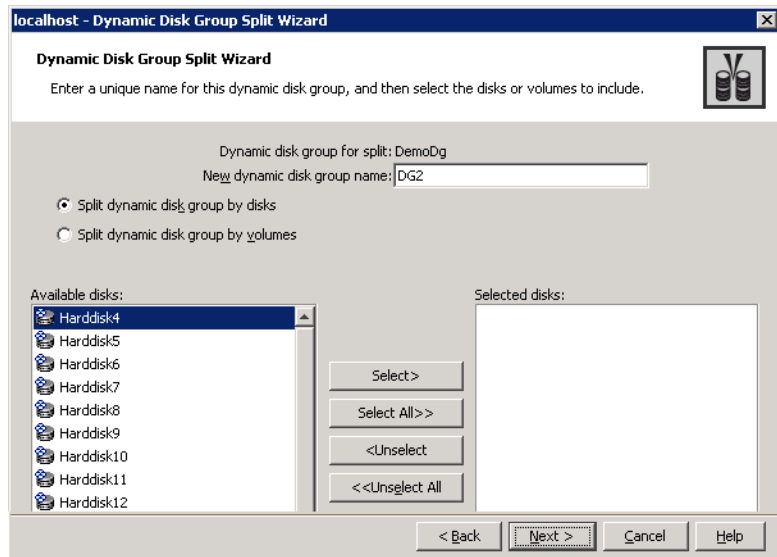
- 3 On the Target Disks selection page, target disks selection is possible for each individual volume or for all volumes.

Note that multiple volume selection is not possible.

Complete the following on this page:

- To assign target disks for all the selected volumes, select **Select All Volumes** check box.
 - To assign target disks for each individual volume, select individual volume under **Source Volumes**
Wizard shows all possible target disks based on volume(s) selection.
 - Select all the empty disks as target disks.
 - Selected destination disks for all the source volumes are displayed in the **Selected Disks** table.
Click **Next**.
- 4 Validate information displayed on the Summary page and click **Finish** to exit the wizard.
 - 5 Now check storage migration status (whether successful or failed) by completing the following on the VEA:
 - Click the **Console** tab on the bottom panel of the VEA.
Verify that separate Subdisk move tasks are created per subdisk.
 - Click the **Tasks** tab on the VEA to check the tasks progress in the bottom panel of the console.
Wait for all Subdisk move tasks to complete.
 - From the **Disk View** tab verify that all selected volumes are now migrated to the selected destination disks.
 - 6 After storage migration completes successfully, split the disk group into two disk groups by selecting the dynamic disk group created by you already in the preceding sections. Right-click the disk group to launch the Split Dynamic Disk Group Wizard
OR
On the VEA, right-click a disk group to select the Split Dynamic Disk Group option.
Click **Next**.

7 Specify the **New dynamic disk group name** (Say **DG2**).



Select **Split dynamic disk group by disks** option.

Select the disks to which Hyper-V volumes are migrated as shown in step 3.

Click **Next**.

8 The Summary page is displayed. Click **Finish** to exit the Wizard.

Now the configuration is changed to one virtual machine per disk group.

Optional Storage Foundation for Windows features for Hyper-V environments

This chapter includes the following topics:

- [About using optional Storage Foundation for Windows features in the Hyper-V parent](#)
- [Dynamic Multi-Pathing for the virtual environment](#)
- [Replicating virtual machines](#)
- [Virtual machine volume snapshots](#)
- [Campus clusters](#)

About using optional Storage Foundation for Windows features in the Hyper-V parent

Running Storage Foundation for Windows (SFW) in the Hyper-V parent partition offers benefits for virtual machines (VMs) that would otherwise be unavailable at the guest level.

SFW also offers advanced features and functionality, such as multi-pathing, replication, and snapshots, as product options. More information is available on how to use the following features and on the benefits they provide when running SFW in the Hyper-V parent:

- Using Dynamic Multi-pathing (DMP) to provide failover and load-balancing to the LUNs that host the VMs in the child partition (DMP Device Specific Modules option)
See “[Dynamic Multi-Pathing for the virtual environment](#)” on page 58.
- Replicating VMs between sites (Veritas Volume Replicator option)
See “[Replicating virtual machines](#)” on page 59.
- Maintaining Quick Recovery snapshots of the VMs (FlashSnap option)
See “[Virtual machine volume snapshots](#)” on page 60.

In addition, running SFW in the parent partition facilitates implementing campus clusters (stretched clusters) in the Hyper-V environment.

See “[Campus clusters](#)” on page 61.

Dynamic Multi-Pathing for the virtual environment

Veritas Storage for Windows (SFW) offers the Dynamic Multi-pathing (DMP) feature. DMP provides an advanced multi-pathing solution for Hyper-V environments.

Multi-pathing software provides the intelligence necessary to manage multiple I/O paths between a server and a storage subsystem. This becomes even more important in virtual environments, where a single physical server hosts multiple operating system instances and applications. Loss of access to storage due to an HBA, cable or controller failure can cause widespread impact, with potential for greater disruption than would be the case with a single physical system.

In virtual environments, a hypervisor separates the VMs from the physical hardware, making it difficult to directly map virtual host bus adapters (HBAs) in the guest to the physical HBAs that service the physical host, and to therefore have a true multi-pathing solution in the guest. Installing DMP in the parent partition allows for true multi-pathing, providing failover and load-balancing to the LUNs that host the VMs in the child partition.

Most multi-pathing solutions are developed by storage vendors for their specific brand of storage and, as such, generally cannot service the many types of storage that can exist in today’s heterogeneous SAN environments. In contrast, DMP is a truly heterogeneous solution. DMP fully integrates with the Microsoft Multipath I/O (MPIO) architecture. DMP includes several Device Specific Modules (DSMs) which provide array-specific support for a wide variety of the most popular array families available.

Replicating virtual machines

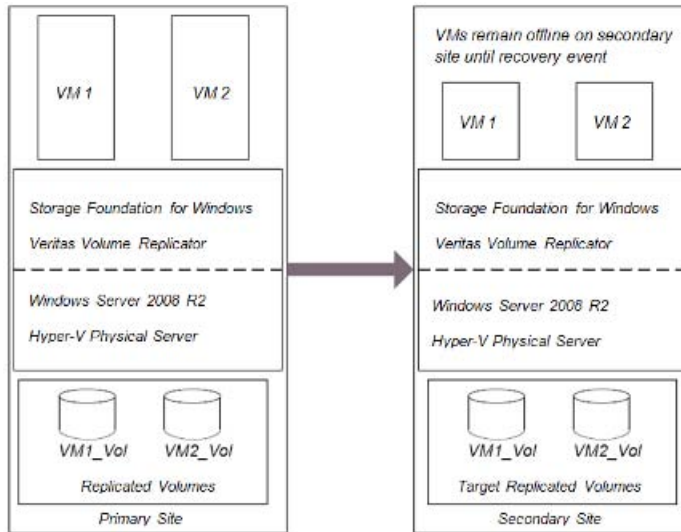
The Veritas Volume Replicator (VVR) option offered by Veritas Storage for Windows (SFW) provides a means to replicate virtual machine (VM) data.

While there are host-based technologies for replicating data across distances, they're usually expensive, requiring not only more storage, but also exactly the same hardware at both ends. They can also be limited in their ability to provide a solution that accounts not only for the data, but also for the applications that access it.

VVR runs at the host level, making it possible to replicate data volumes across distances to provide a means to extended disaster recovery without requiring that the hardware be exactly the same at both ends, and generally requiring less storage. But, while this works fine to protect application data in the guest, allowing recovery at the DR site, running in the guest does nothing to protect the VM.

By installing SFW and VVR in the parent partition, volumes that contain VHD files used for VMs and/or application data can be selectively replicated to DR sites, either synchronously or asynchronously, over an IP network. VVR uses a replicator log to store all writes to the volumes grouped together in what is known as a replicated volume group in the correct order and replicate them to the DR site, maintaining write order fidelity, and thereby, consistency. The replicated VMs remain offline at the DR site until required to be brought online, either due to an intentional migration of services from the primary site, or due to an outage at the primary site, requiring the DR site to take over services.

Figure 4-1 VVR in the parent partition



For planned outages at the primary site, perhaps for a maintenance window, the primary role can be migrated to a DR site, allowing the VMs that have been replicated to be brought online and applications that run on them to access data that has also been replicated to the DR site.

For unplanned outages at the primary site, operations can be moved to a DR site via a takeover operation, which turns the DR site into an active primary, allowing VMs to be brought online. Depending on the mode of replication, they can either be completely up to date or behind the previous primary. In either event, consistency will be maintained and applications will be able to successfully attach to their data. The primary can be migrated back to the original site when it becomes available.

Virtual machine volume snapshots

Snapshot technology is available with the Veritas Storage for Windows (SFW) FlashSnap option. Running in the guest, snapshot copies of data volumes can be created, allowing for quickly recovering from a disaster, or for off-host operations, which can occur to another virtual machine (VM) or to a physical server connected to storage shared with the guest when it uses pass-through disks. SFW also supports Microsoft's VSS framework for creating consistent snapshots.

However, this offers no protection against possible disaster that can occur at the VM level. If the VHD file that holds a VM is corrupted, the volume that hosts the

VHD file(s) is lost, or the LUN hosting the volume used by the VHD file fails, snapshots in the guest will be useless until the VM can be rebuilt and the application(s) reinstalled.

By running SFW in the parent partition, you have the advantage of being able to create snapshots of the volumes containing the VHDs. These snapshots can be used to quickly recover the entire VM in the event of a disaster. They can also be moved to another server and brought online to be backed up or used for other operations such as testing.

Campus clusters

As a host-based volume manager, Veritas Storage Foundation for Windows (SFW) provides the ability to mirror volumes across arrays. Clusters which rely on shared storage can be stretched beyond the confines of a single datacenter to a datacenter located at a DR site, as long as the distance between the two datacenters doesn't exceed fiber channel latency limitations. These stretched clusters, also known as campus clusters, provide a level of high availability that can withstand a complete site failure.

SFW running in the parent partition can facilitate stretching of the failover cluster to another site by providing support for dynamic disks through its Volume Manager Disk Group cluster resource. With dynamic disks now available in the cluster, volumes can be mirrored across arrays which are located in different datacenters and are seen and shared by cluster nodes located in those respective datacenters. If the active site should experience a failure, virtual machines that were running at that site can be failed over to the cluster node at the other datacenter, and applications running on them can be brought back online.

Veritas Cluster Server for Windows (VCS)

- [Chapter 5. Overview of the Disaster Recovery for Hyper-V solution](#)
- [Chapter 6. Deploying Hyper-V disaster recovery](#)
- [Chapter 7. Hyper-V DR agent](#)

Overview of the Disaster Recovery for Hyper-V solution

This chapter includes the following topics:

- [About wide-area disaster recovery for Microsoft Hyper-V](#)
- [Advantages of Disaster Recovery Manager for Microsoft Hyper-V](#)
- [About the Disaster Recovery Manager for Microsoft Hyper-V configuration](#)
- [How disaster recovery with Disaster Recovery Manager works](#)

About wide-area disaster recovery for Microsoft Hyper-V

Veritas Cluster Server (VCS) for Windows introduces the Disaster Recovery Manager for Microsoft Hyper-V option to support wide-area disaster recovery for a Microsoft Hyper-V cluster.

Wide-area disaster recovery maintains data and critical services if a disaster affects a local area or metropolitan area. Data and critical services can be failed over to a site that is located hundreds or thousands of miles away.

A Microsoft failover cluster can provide high availability within a primary site. This includes live migrating virtual machines (VMs) between local cluster nodes if a node fails. However, current disaster recovery solutions for VMs have limitations that reduce their effectiveness.

Wide-area disaster recovery with the Disaster Recovery Manager for Microsoft Hyper-V option provides new capabilities for a robust disaster recovery solution. A separate Management VM is set up on the primary and DR site to implement, monitor, and manage disaster recovery and replication.

Advantages of Disaster Recovery Manager for Microsoft Hyper-V

Wide-area disaster recovery with the Veritas Cluster Server (VCS) Disaster Recovery Manager for Microsoft Hyper-V option provides the following advantages:

- This easy-to-deploy solution embraces your existing Hyper-V environment and extends its capabilities with a true disaster recovery solution. High availability for the application virtual machines (VMs) is provided by Windows Failover Cluster, while disaster recovery is provided by the Disaster Recovery Manager for Hyper-V solution.
- Disaster recovery communication between sites is easily configured with the Disaster Recovery Configuration Wizard for Microsoft Hyper-V and monitored by a separate Management VM at each site.
- The primary and DR site can be on different subnets. The Management VM enables automated configuration of network settings at both sites. The DR site's network settings are implemented automatically, upon failover, by the Symantec Network Updater Service.
- The Management VMs monitor heartbeat communications between the primary and DR cluster.
- The Management VMs monitor the state of the application VMs and replication. No heartbeat is required between the Management VM and the application VMs.
- If the primary site fails, the Management VMs handle the failover between sites, providing options for quickly bringing all application VMs online.

About the Disaster Recovery Manager for Microsoft Hyper-V configuration

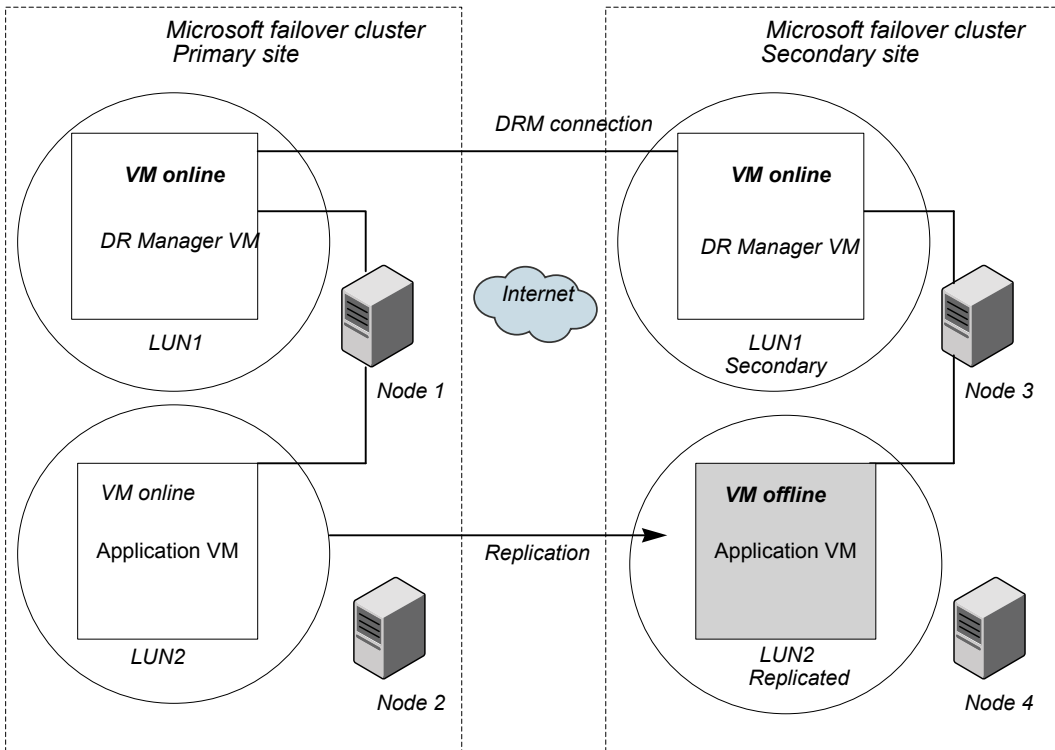
Wide-area disaster recovery for VMs with Disaster Recovery Manager uses the following configuration components:

- Separate Microsoft clusters are set up on the primary and DR site.

- The application VMs are configured for high availability within the Microsoft cluster on the primary site.
Disaster Recovery Manager includes support for configuring application VMs on storage managed as Storage Foundation for Windows (SFW) dynamic disk groups and volumes. SFW has added support for live migration and storage migration for Hyper-V VMs.
See [“Advantages of running Storage Foundation for Windows in the Hyper-V parent”](#) on page 13.
- Hardware replication is set up to replicate the application VMs and data LUNs from the primary to the DR site.
Disaster Recovery Manager supports Hitachi TrueCopy or EMC SRDF replication.
- A separate VM for the Disaster Recovery Manager is set up on the primary and DR sites on separate non-replicated storage. The Disaster Recovery Manager is installed on both VMs.
- Using the wizard provided, a connection is configured between the DR Manager VMs on the primary and DR site, and other settings are configured for disaster recovery support.

As shown in the illustration, the application VMs and the Management VM are online on node 1 of the Microsoft cluster on the primary site. On the DR site the Management VM is online on node 3 but the application VMs remain offline. If node 1 becomes unavailable, the VMs can fail over to node 2 on the primary site. If both node 1 and node 2 go down on the primary site, the Management VM on the DR site enables you to quickly bring the application VMs online on node 3.

Figure 5-1 VCS Hyper-V disaster recovery



How disaster recovery with Disaster Recovery Manager works

The Disaster Recovery Manager monitors the application virtual machines (VMs) in the Microsoft failover cluster. All the VMs that are being monitored by the Disaster Recovery Manager must fail to trigger the DR failover. Failover occurs at the site level, not at the individual VM level.

Optionally, the Disaster Recovery Manager can automate updating of network settings when the application VMs are failed over between sites. The primary and DR site can be on different subnets. The application VMs at each site can therefore have different network settings. To implement automatic updating of network settings, you configure a network settings file for each site using the wizard. The network settings are implemented automatically, upon failover, by the Symantec Network Updater Service, which is copied to the application VMs during disaster recovery configuration.

As part of the disaster recovery configuration process, you export the application VM configurations using the wizard. Bringing the VMs online on the DR site imports the application VM configurations on the DR site and updates their network settings.

If the primary site fails, you use the Disaster Recovery Manager to bring the VMs online at the disaster recovery site. You can also manually initiate a failover at any time to test the disaster recovery solution.

See [“Bringing the DR site up if the primary site fails”](#) on page 92.

Deploying Hyper-V disaster recovery

This chapter includes the following topics:

- [Requirements for Disaster Recovery Manager](#)
- [Ports used by Disaster Recovery Manager](#)
- [Workflow for deploying Hyper-V disaster recovery](#)
- [Configuration requirements for the DR site disaster recovery cluster](#)
- [Setting up the hardware replication for disaster recovery](#)
- [Preparing a VM for Hyper-V DR installation](#)
- [Installing the Disaster Recovery Manager for Hyper-V](#)
- [Configuring disaster recovery using the Disaster Recovery Configuration Wizard for Microsoft Hyper-V](#)
- [Connecting to the Disaster Recovery Manager with the Java Console](#)
- [Manually failing over the VMs between sites](#)
- [Bringing the DR site up if the primary site fails](#)

Requirements for Disaster Recovery Manager

Disaster Recovery Manager has the following requirements:

Table 6-1 Requirements for Disaster Recovery Manager

Requirement	Description
Hardware	Hardware must meet the requirements specified in the Veritas Cluster Server for Windows 6.0.2 hardware compatibility list: http://www.symantec.com/docs/TECH152806
Operating system in Hyper-V parent	Windows Server 2012 must be installed on the Hyper-V parents in the Microsoft cluster on the primary and DR sites.

Table 6-1 Requirements for Disaster Recovery Manager (*continued*)

Requirement	Description
<p>Application virtual machines (VMs)</p>	<p>Application VMs can run the following Windows operating systems:</p> <ul style="list-style-type: none"> ■ Windows Server 2008 ■ Windows Server 2008 R2 ■ Windows Server 2012 ■ Windows Vista ■ Windows 7 <p>For details on supported versions, see the Veritas Cluster Server for Windows 6.0.2 software compatibility list: http://www.symantec.com/docs/TECH201485</p> <p>Application VMs can be configured on CSV (cluster shared volume) or physical disk resource in the Microsoft failover cluster.</p> <p>If Storage Foundation for Windows (SFW) is installed in the parent, application VMs can be configured on SFW dynamic disk groups and volumes configured as Volume Manager Disk Group (VMDg) resources</p> <p>See “Advantages of running Storage Foundation for Windows in the Hyper-V parent” on page 13.</p> <p>The application VMs are configured on the primary site only.</p> <p>Note: Disaster Recovery Manager provides an optional feature to automate updating of network settings. To ensure that the automatic update works consistently, install following Microsoft hotfixes as appropriate.</p> <p>For VMs running Vista or Windows Server 2008, see Microsoft KB 950134: http://support.microsoft.com/kb/950134/</p> <p>For VMs running Windows 7 or Windows Server 2008 R2, see Microsoft KB 2530185: http://support.microsoft.com/kb/2530185</p>

Table 6-1 Requirements for Disaster Recovery Manager (*continued*)

Requirement	Description
MAC address range	<p>If the application VMs use dynamically generated MAC addresses, ensure that all Hyper-V hosts in the primary and DR clusters use the same MAC address range. You can verify or configure the MAC address range setting with the Virtual Network Manager, available from the Microsoft Hyper-V Manager. In the Virtual Network Manager window, in the Global Network Settings section, select the MAC Address Range setting. Note that changing this setting does not affect network adapters that have already been created. To apply to existing adapters, recreate the adapter.</p>
VMs for Disaster Recovery Manager	<p>A separate VM, running Windows Server 2012, must be configured for the Disaster Recovery Manager on both the primary and DR site.</p> <p>The Disaster Recovery Manager software is installed within this VM.</p> <p>There are additional requirements for setting up the Management VM.</p>
Hardware replication	<p>The sites must support Hitachi TrueCopy or EMC SRDF array-based replication for the application VMs.</p> <p>See “Setting up the hardware replication for disaster recovery” on page 78.</p>
Disk space required	<p>Approximately 824 MB of disk space is required for the Disaster Recovery Manager installation.</p> <p>If you configure the feature to update network settings automatically, the Symantec Network Updater Service is copied to the application VMs by the Disaster Recovery Manager during disaster recovery configuration. It requires a small amount of disk space.</p>
Permissions and rights	<p>Installing and configuring Disaster Recovery Manager requires local administrator permission for the Microsoft cluster systems, the Management VM, and the application VMs.</p> <p>If configuring the feature to update network settings automatically, the application VMs must allow a connection to the administrative share to allow copying the Symantec Network Updater Service binaries to the application VMs.</p>

Table 6-1 Requirements for Disaster Recovery Manager (*continued*)

Requirement	Description
Required services	The following services must be running inside the application VMs to support disaster recovery using Disaster Recovery Manager: <ul style="list-style-type: none"> ■ Hyper-V Integration Services (typically installed during Hyper-V installation) ■ Symantec Network Updater Service (optional) Installed on the application VMs during configuration of disaster recovery only if you want to implement the feature to update network settings automatically.
Required static IPv4 address	Each Management VM must be configured with at least one static IPv4 address.
Firewall settings and ports	Ensure that the firewall settings allow access to the ports used by Disaster Recovery Manager. See “Ports used by Disaster Recovery Manager” on page 75.

Ports used by Disaster Recovery Manager

Ensure that the firewall settings allow access to the following ports that may be used by Disaster Recovery Manager.

Table 6-2 Ports used by Disaster Recovery Manager

Port number	Protocol	Description	Process
14150	TCP	Veritas Command Server	CmdServer.exe
14141	TCP	Veritas High Availability Engine Veritas Cluster Manager (Java Console) (ClusterManager.exe) VCS Agent driver (VCSAgDriver.exe)	had.exe
7419	TCP	Symantec Plugin Host Service	pluginHost.exe

Table 6-2 Ports used by Disaster Recovery Manager (*continued*)

Port number	Protocol	Description	Process
14149	TCP/UDP	VCS Authentication Service	vcsauthserver.exe
14144	TCP/UDP	VCS Notification	Notifier.exe
14155	TCP/UDP	VCS Global Cluster Option (GCO)	wac.exe

Workflow for deploying Hyper-V disaster recovery

To deploy Hyper-V disaster recovery for application virtual machines (VMs) in a Microsoft cluster using Disaster Recovery Manager, perform the following tasks in the sequence shown.

Table 6-3 Process for deploying Hyper-V disaster recovery

Action	Description
Review the requirements	<p>Make sure that you understand the planned configuration and that your environment meets the hardware and software requirements.</p> <p>See “About the Disaster Recovery Manager for Microsoft Hyper-V configuration” on page 66.</p>
On the primary site, set up the application VMs for high availability in a Microsoft cluster	<p>See the Microsoft documentation.</p> <p>Information is also available on configuring VMs on Storage Foundation for Windows (SFW) for live migration.</p>
On a DR site, set up a Microsoft cluster for Hyper-V	<p>On a DR site, set up the desired number of nodes as part of a separate Microsoft cluster.</p> <p>More information is available on the requirements for the DR cluster.</p>
Set up hardware replication between both sites	<p>Disaster Recovery Manager supports EMC SRDF or Hitachi TrueCopy replication.</p> <p>See “Setting up the hardware replication for disaster recovery” on page 78.</p>

Table 6-3 Process for deploying Hyper-V disaster recovery (*continued*)

Action	Description
On both sites, configure the VM to be used for the Disaster Recovery Manager	
On both sites, install the Disaster Recovery Manager	See “Installing the Disaster Recovery Manager for Hyper-V” on page 82.
On the primary site, run the Disaster Recovery Configuration Wizard for Microsoft Hyper-V	See “Configuring disaster recovery using the Disaster Recovery Configuration Wizard for Microsoft Hyper-V” on page 83.
Verify that the DR configuration is successful using the VCS Java Console	Use the VCS Java Console to connect to the Disaster Recovery Manager and to verify that configuration was successful See “Connecting to the Disaster Recovery Manager with the Java Console” on page 86.
(Optional) Manually switch the application VMs between sites	Optionally, manually switch the application VMs between sites using the VCS Java Console. See “Manually failing over the VMs between sites” on page 89.
If the primary site goes down, bring the DR site up	Use the VCS Java Console to bring the DR site up if the primary site fails. Note: In addition, only for VMs configured on Storage Foundation for Windows storage, configure a dependency in the Microsoft failover cluster on the DR site between the VM resources and the VMDg resources. See “Bringing the DR site up if the primary site fails” on page 92.
To add or remove application VMs or modify the DR configuration, run the Disaster Recovery Configuration Wizard for Microsoft Hyper-V.	See “Configuring disaster recovery using the Disaster Recovery Configuration Wizard for Microsoft Hyper-V” on page 83.

Configuration requirements for the DR site disaster recovery cluster

Disaster recovery for Hyper-V VMs using Disaster Recovery Manager requires setting up a separate Microsoft failover cluster on a DR site. Note the following requirements for the DR cluster:

- Ensure that the DR cluster is in the same domain as the primary cluster.
- Ensure that the systems that form the cluster nodes of the DR cluster meet the same hardware and software requirements as the primary cluster.
- When configuring Hyper-V on the DR cluster, set up the same virtual network switches as on the primary site.
- If the application VMs use dynamically generated MAC addresses, ensure that all Hyper-V hosts use the same MAC address range as on the primary cluster. To verify or configure the MAC address range setting, use Virtual Network Manager, available from the Hyper-V Manager. In the **Virtual Network Manager** window, in the **Global Network Settings** section, select the **MAC Address Range** setting.
- Create cluster shared volumes (CSVs) or physical disk resources on the DR site cluster to match what was set up for the application VMs on the primary cluster.

If the application VMs are configured on SFW storage, create the appropriate Volume Manager Disk Group (VMDg) resources on the DR site cluster to match what was set up for the application VMs on the primary cluster.

See [“Adding the Volume Manager Disk Group \(VMDg\) resource”](#) on page 38.

Note: CSV or SFW (VMDg) resources created on the DR site cluster will remain in a FAILED state until a failover occurs and replication roles reverse.

Setting up the hardware replication for disaster recovery

As part of the workflow for setting up disaster recovery for Hyper-V VMs using Disaster Recovery Manager, configure hardware replication for the application VMs and their associated storage.

Note: Configure replication only for the application VMs and their associated storage. Do not configure replication for the Management VM and its associated storage.

The hardware replication must be set up before you can configure disaster recovery with the Disaster Recovery Manager.

The Disaster Recovery Manager supports EMC SRDF replication and Hitachi TrueCopy replication. It uses agents to monitor the replication.

The replication agents do not support specialized software solutions that the array vendor may have developed for cluster shared volumes (CSV).

The following topics cover the requirements for configuring replication to work with the Disaster Recovery Manager.

See [“Prerequisites for EMC SRDF replication with Disaster Recovery Manager”](#) on page 79.

See [“Prerequisites for Hitachi TrueCopy replication with Disaster Recovery Manager”](#) on page 80.

Prerequisites for EMC SRDF replication with Disaster Recovery Manager

The Disaster Recovery Manager supports EMC SRDF replication. Before you configure the Disaster Recovery Manager, you must configure the replication for the application VMs and their associated storage.

Note: Do not configure replication for the Management VM and its associated storage.

The Disaster Recovery Manager includes agent software that supports SYMCLI versions that EMC recommends for the firmware on the array. The agent supports SRDF on all microcode levels on all Symmetrix arrays, provided that the host/HBA/array combination is in EMC's hardware compatibility list.

To enable the Disaster Recovery Manager to configure the required agent settings for EMC SRDF, ensure that the following software requirements are met:

- The EMC Solutions Enabler is installed on all cluster nodes.
- The SYMCLI version that is installed supports the generation of XML output.
- The SYMCLI version and the microcode level support dynamic swapping.

The agent has the following requirements for EMC SRDF replication:

- The device group must not span more than one array (no composite device groups).
- Dynamic swap must be enabled on both sites.
- On the primary site:
 - All devices must be RDF1 and part of an RDF1 device group.
 - Devices must have write access.
- On the DR site:
 - All devices must be RDF2 and part of an RDF2 device group.
 - Write access must be disabled.
- Device group configuration must be the same on all nodes of the cluster.

Prerequisites for Hitachi TrueCopy replication with Disaster Recovery Manager

The Disaster Recovery Manager supports Hitachi TrueCopy replication. Before you configure the Disaster Recovery Manager, you must configure the replication for the application VMs and their associated storage.

Note: Do not configure replication for the Management VM and its associated storage.

Make sure that the Hitachi RAID Manager version installed is 01-19-03/04 or later.

Ensure that the following requirements are met before configuring the Disaster Recovery Manager:

- RAID Manager is installed on the Management VMs.
- The horcm files are named horcmnn.conf (where nn is a positive number without a leading zero, for example, horcm1.conf but not horcm01.conf).
- All configured instances are running.
- The device group does not span more than one array.
- At the primary site, all devices are of the type P-VOL.
- At the DR site, all devices are of the type S-VOL.
- All device groups at the primary site are paired to an IP address which must be online on the DR node.

- Device group and device names include only alphanumeric characters or the underscore character.

Preparing a VM for Hyper-V DR installation

Prepare a Hyper-V virtual machine (VM) for the Disaster Recovery Manager as follows:

- Set up a separate Management VM in the Microsoft Hyper-V cluster on both the primary and DR site.
- The Management VMs must be in the same Active Directory domain as the Microsoft failover clusters.
- For the Management VM operating system, install Windows Server 2012.
- Install the Microsoft failover cluster management tools on the Management VMs.
- Install the Remote Server Administration Tools on the Management VMs on the DR site.
- Configure at least one static IPv4 address on each Management VM.
- Enable PowerShell Remoting on the Management VMs.
- Configure the storage for the Management VM on a separate (non-replicated) LUN from the application VMs.
- Install and configure the array replication management software on the Management VMs.
- Configure the arrays to give the Management VMs access to the array hardware gateway (command) devices. This is required during Disaster Recovery Manager configuration.
- On the primary site, set up the Management VM for failover within the primary Microsoft failover cluster.
- When configuring a Virtual Machine role in the failover cluster, make sure that its name matches the host name of the Management VM. Otherwise, the Disaster Recovery Configuration Wizard for Microsoft Hyper-V might incorrectly list the Management VM along with the application VMs that are available for disaster recovery.
If the user selects the Management VM in the wizard, the disaster recovery configuration will not succeed.

Installing the Disaster Recovery Manager for Hyper-V

Install the Veritas Cluster Server for Windows (VCS) Disaster Recovery Manager for Hyper-V on the DR Manager virtual machines (VMs). Before installing, ensure that you have met the requirements.

To install the Disaster Recovery Manager for Hyper-V

- 1 From any system that can connect to the Management VMs, insert the disc containing the installation software into the system's disk drive or download the VCS 6.0.2 for Windows package from the Symantec Web site.
<https://fileconnect.symantec.com>
- 2 Allow the autorun feature to start the installation or double-click **Setup.exe**. The CD browser appears.
- 3 Select **Veritas Cluster Server 6.0.2**.
- 4 On the Welcome panel, review the list of prerequisites and click **Next**.
- 5 On the License panel, read the license terms, select **I accept the terms of License Agreement**, and then click **Next**.
- 6 On the System Selection panel, make the following selections:
 - Select both Management VM(s) for installation. You can install on multiple systems.
 - By default the wizard uses `%ProgramFiles%\Veritas` as the installation directory. To customize the installation directory, click **Browse** and select the desired location, then click **OK**.
 - Select the required license type from the **License key** drop-down list. The default license type is "Keyless".
If you select the "Keyless," all the available product options are displayed and are selected by default.
If you select "User entered license key," the License Details panel appears by default. On the License Details panel, enter the license key and then click **Add**.
The wizard validates the entered license keys. After the validation is complete, click **OK**.
 - From the product options list, select the **Disaster Recovery Manager for Microsoft Hyper-V** option to install on both VMs.

- 7 When you have completed all selections on the System Selection panel, click **Next**.

If the validation checks have failed, review the details and rectify the issue. Then select the system and click **Re-verify** to re-initiate the validation checks for this system.

- 8 On the Pre-install Summary panel, review the summary and click **Next**.

Note that the **Automatically reboot systems after installer completes operation** check box is selected by default. This will reboot all the selected remote systems immediately after the installation is complete on the respective system. If you do not want the wizard to initiate this auto reboot, clear the selection of **Automatically reboot systems after installer completes operation** check box.

- 9 When installation is complete, click **Next**.

- 10 On the Post-install Summary panel, review the installation result and click **Next**.

If the installation has failed, refer to the log file for details. You may have to reinstall the software.

- 11 On the Finish panel, click **Finish**.

If you had chosen the auto reboot, a confirmation message appears. Click **Yes** to reboot immediately or **No** to reboot later.

If you did not choose the auto reboot, ensure that you manually restart.

Configuring disaster recovery using the Disaster Recovery Configuration Wizard for Microsoft Hyper-V

The Disaster Recovery Configuration Wizard for Microsoft Hyper-V assists you to perform the following tasks:

- Export the application VMs configuration files on the primary site
- Configure the network settings file on the primary site and DR site
- Configure global clustering (GCO)
- Configure replication monitoring between the primary site and DR site

The wizard is launched when you log on to the Management VM after product installation. You can also launch it from the **Start** menu.

To run the wizard you must be a domain user with administrator rights on the Management VMs on the primary site and DR site.

To configure disaster recovery using the wizard

- 1 Launch the Disaster Recovery Configuration Wizard for Microsoft Hyper-V. Click **Start>All Programs > Symantec > Veritas Cluster Server > Disaster Recovery Configuration Wizard for Microsoft Hyper-V**.
- 2 On the Welcome panel, review the list of prerequisites and click **Next**.
- 3 On the Failover Cluster Inputs panel, specify the following information:
 - The Fully Qualified Host Name (FQHN) or the IP address of the Management VM on the DR site.
 - Administrator account credentials for the primary site Microsoft failover cluster.
 - Administrator account credentials for the DR site Microsoft failover cluster (if different from the primary site cluster).Click **Next**.
- 4 On the System Selection panel, select the application VMs to be added or removed from the disaster recovery configuration and click **Next**.

The wizard discovers the application VMs configured on the primary site failover cluster and displays their state and export status.

Caution: If the names of the Virtual Machine role in the failover cluster and the Management VM host do not match, the Management VM might incorrectly appear in the list along with the available application VMs.

Do not select the Management VM. Otherwise, the disaster recovery configuration will not succeed.

- To add an application VM to the disaster recovery configuration, select the VM and click **Next**. The wizard exports the configuration of the VM, that is, it saves the VM configuration file on a shared disk. Only the application VMs with exported configurations are failed over to the DR site.

Note: The selected application VMs are shut down during this operation.

The wizard changes the **Cluster-Controlled Offline Action** setting on the VM cluster resources in the Microsoft cluster from the default of to **Shutdown (Forced)**. This setting change is required for the VMs to properly fail over to the recovery site.

- The VMs configured for disaster recovery are pre-selected. To remove the VM from the disaster recovery configuration, deselect the VM and click **Next**.
- 5 On the Network Settings panel, select an application VM and specify the following network settings at the primary site and DR site:
- IP address
 - Subnet mask
 - Gateway address
 - DNS server address

The Management VM enables automated configuration of network settings when the application virtual machines (VMs) are failed over between sites. The primary and DR site can be on different subnets. The application VMs at each site can therefore have different network settings.

To implement automatic updating of the application VM network settings, the wizard configures two network settings files, one for each site. The file on the primary site contains the network settings used on the primary site for each application VM. The file on the DR site contains the network settings to be used on the DR site for each application VM. The network settings in the file are implemented automatically, upon failover, by the Symantec Network Updater Service, which is copied to the application VMs during disaster recovery configuration.

If the selected VM and the primary site Management VM are in different domains, specify the user account credentials for the selected VM.

- 6 On the Replication Settings panel, the details of the array-based replication method configured between the primary site and the DR site is displayed:
- If Hitachi TrueCopy replication is configured, the following details are displayed:

Device Group Name	Enter the name of the Hitachi TrueCopy device group that contains the LUNs for the selected instance. The same device group name is used on both sites.
Instance ID	Enter the instance number of the device group. Multiple device groups may have the same instance number. The same instance number is used on both sites.

In case the RAID Manager bin path is different from the path displayed by wizard, replication settings may not be discovered. In such situations, modify the base directory (RAID Manager bin path) and configuration

directory (HORCM files location) as per your environment and click **Refresh**.

- If EMC SRDF replication is configured, the name of the Symmetrix Array ID (SID) and device group are displayed.

To add a hardware resource to the control group, use the **Add** button. To remove a hardware resource from the control group, use the **Remove** button.

- 7 On the Summary panel, review the summary and click **Next**.
- 8 On the Implementation panel, review the progress of the implementation and click **Next** after the implementation is complete.
- 9 On the Finish panel, review the post-configuration summary and click **Finish** to exit the wizard.

Modifying the application VM configuration after configuring disaster recovery

While configuring disaster recovery using the wizard, the application VM configuration is exported. If you modify the application VM configuration after configuring disaster recovery, the modified configuration is different from the exported configuration. In this case, if the application VM fails over to the DR site, it may not come online on the DR site.

Perform the following steps to avoid this issue.

To modify the application VM configuration after configuring disaster recovery

- 1 Run the Disaster Recovery Configuration Wizard for Microsoft Hyper-V to remove the application VM from the disaster recovery configuration.
- 2 Navigate to the location where the application VM configuration is stored.
- 3 Delete the Export folder.
- 4 Re-run the Disaster Recovery Configuration Wizard for Microsoft Hyper-V to add the application VM to the disaster recovery configuration.

Connecting to the Disaster Recovery Manager with the Java Console

Use the Veritas Cluster Server (VCS) Cluster Manager Java Console to connect to the Disaster Recovery Manager. You can use the Java Console to do the following:

- Verify that the DR Manager configuration was successful
- Switch application VMs between sites

- Bring the DR site online because of a primary site failure


For a complete description of all Java Console features, see the *Veritas Cluster Server Administrator's Guide*.

To connect to the Disaster Recovery Manager with the Java Console

- 1 If using a firewall, ensure that you have added ports 14141 and 14150 for firewall exceptions.
- 2 Start the Java Console: Click **Start > All Programs > Symantec > Veritas Cluster Server > Veritas Cluster Manager - Java Console**

The first window that appears is Cluster Monitor. The primary and DR Management VMs are each a separate VCS cluster. You will add cluster panels for the DR Manager clusters and then log on to complete the connection.

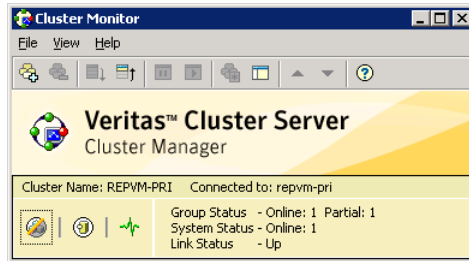
- 3 Add a cluster panel for the primary Management VM cluster, as follows:
 - In Cluster Monitor, click **File > New cluster**.
 - In the **Host Name** field, enter the host name or IP address of the primary Management VM.
 - Click **OK**.
The primary DR Manager cluster panel is added, as shown in the illustration. The panel is inactive (not connected to the cluster) until you log on.



- 4 To add a cluster panel for the DR Management VM cluster, repeat the previous step, but this time specify the host name or IP address of the DR Management VM in the **Host Name** field.
- 5 To log on to a cluster panel:
 - Click on the inactive cluster panel. The log on dialog box is displayed.
 - Enter the credentials of a native user. You can use nis or nis+ accounts or accounts set up on the local system. If you do not enter the name of the domain, VCS assumes the domain is the local system.


- Click **OK**.

The cluster panel changes color and displays connection status.



The Cluster Explorer is launched automatically upon initial logon.

- 6 To verify that the DR configuration was successful, continue with the following step. Otherwise, see one of the following topics for additional tasks that you can perform from the Java Console:
 - See [“Manually failing over the VMs between sites”](#) on page 89.
 - See [“Bringing the DR site up if the primary site fails”](#) on page 92.
- 7 To verify that the DR configuration was successful, log on to the primary site cluster panel and do the following steps from the Cluster Explorer for the primary site:



- By default the Service Groups tab is selected in the left pane configuration tree, and the tree displays the **ControlGroup** node. Select **ControlGroup**.
- In the right pane, check that in the **Status** view, the **Group Status on Member Systems** shows the state of the primary Management VM as **Online**, as shown in the illustration.
- If the state is not **Online**, in the configuration tree, right-click **ControlGroup** and click **Online**.

Manually failing over the VMs between sites

You can use the Veritas Cluster Server (VCS) Cluster Manager Java Console to test the DR failover manually when the primary site is still online. The failover brings the application VMs offline on the primary site cluster and online on the DR site cluster.

If the primary site has failed, use the procedure in the following topic instead of the procedure below:

See [“Bringing the DR site up if the primary site fails”](#) on page 92.

You can test the failover from either the primary or DR site. The instructions in the following procedure demonstrate logging on to the Java Console and Management VM cluster on the DR site.

Before doing this procedure, complete the steps to connect to the Java Console and verify the configuration.


See [“Connecting to the Disaster Recovery Manager with the Java Console”](#) on page 86.

To manually fail over the VMs between sites

- 1** On the DR site, launch the Java Console from the Management VM. From the Start menu, click **Start > All Programs > Symantec > Veritas Cluster Server > Veritas Cluster Manager - Java Console**
- 2** If you have not yet done so, use the Java Console Cluster Monitor to log on to the cluster panel for the DR site Management VM.

See [“Connecting to the Disaster Recovery Manager with the Java Console”](#) on page 86.


- Cluster Explorer for the DR site Management VM is launched. In the configuration tree, select **ControlGroup**.



If you are logged on to the DR cluster, the **Status** view shows the group status for the DR site Management VM as **Offline**, as shown in the illustration. The status for the primary site cluster is listed under **Remote cluster** and shown as **Online**.

- In the configuration tree, right-click **ControlGroup**.
- Click **Switch To**, and click **Remote switch**.

The **Switch global group** dialog box is displayed.



- 6 In the **Switch global group** dialog, the default shown in the **Cluster** field is **Local** (for the DR Manager cluster you are logged on to). If you are logged on to the DR site cluster, leave the selection as **Local** and click **OK**. (Since there is only one Management VM on each site, no selection is needed in the **System** field.)

If you are logged on to the primary site cluster instead, select the name of the DR site DR Manager in the **Cluster** list.
- 7 When prompted, confirm the switch. In the Cluster Explorer **Status** view, verify that the state of the DR site Management VM changes from **Offline** to **Online**.
- 8 Confirm that the application VMs are online on the DR site using the Windows failover cluster management console.
- 9 If you configured automatic update of network settings, verify that the settings are updated. In some cases, upon the first failover, Windows detects configuration changes in the application VM and displays a message prompting you to restart the computer to apply the changes. If settings are not updated, check for this message and restart the VM.
- 10 If VMs are configured on SFW storage, then the first time the VMs are failed over, you must configure the resource dependency between the application VM resource and the VMDg resource.

See [“Setting the dependency of the virtual machine on the VMDg resource”](#) on page 41.

To switch the application VMs back to the primary site

- 1 In the Cluster Explorer configuration tree, right-click **ControlGroup**.
- 2 Click **Switch To**, and click **Remote switch**.
- 3 In the **Switch global group** dialog box, if you are logged on to the DR site Management VM cluster, ensure that the **Cluster** field displays the name of the Management VM for the primary site. Click **OK**.
- 4 In the Cluster Explorer **Status** view, verify that the state of the primary site Management VM cluster changes to **Online**.

Bringing the DR site up if the primary site fails

If the primary site fails, use the Veritas Cluster Server (VCS) Cluster Manager Java Console to connect to the Disaster Recovery Manager and bring the DR site online.

For a complete description of all Java Console features, see the *Veritas Cluster Server Administrator's Guide*.

The following procedure assumes that in a disaster recovery scenario, you are running the VCS Java Console from the Management VM on the DR site, since the primary site is down.

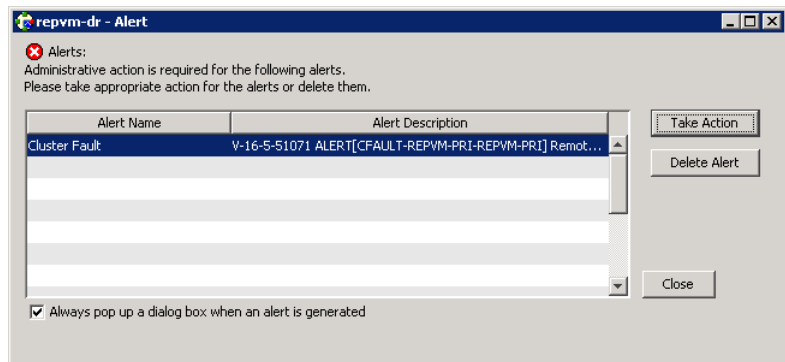
To bring the DR site up if the primary site fails

1 On the DR site, launch the Java Console from the Management VM. From the Start menu, click **Start > All Programs > Symantec > Veritas Cluster Server > Veritas Cluster Manager - Java Console**

2 If you have not yet done so, in the Java Console, log on to the DR site Management VM cluster panel.

See [“Connecting to the Disaster Recovery Manager with the Java Console”](#) on page 86.

3 If the primary site cluster has failed in a disaster recovery scenario, the **Alert** dialog box is displayed by default when you log on to the DR site cluster:



If the Alert dialog box is displayed, continue with the following steps. Otherwise, skip to step 7.

4 On the **Alert** dialog box, click **Take Action**.

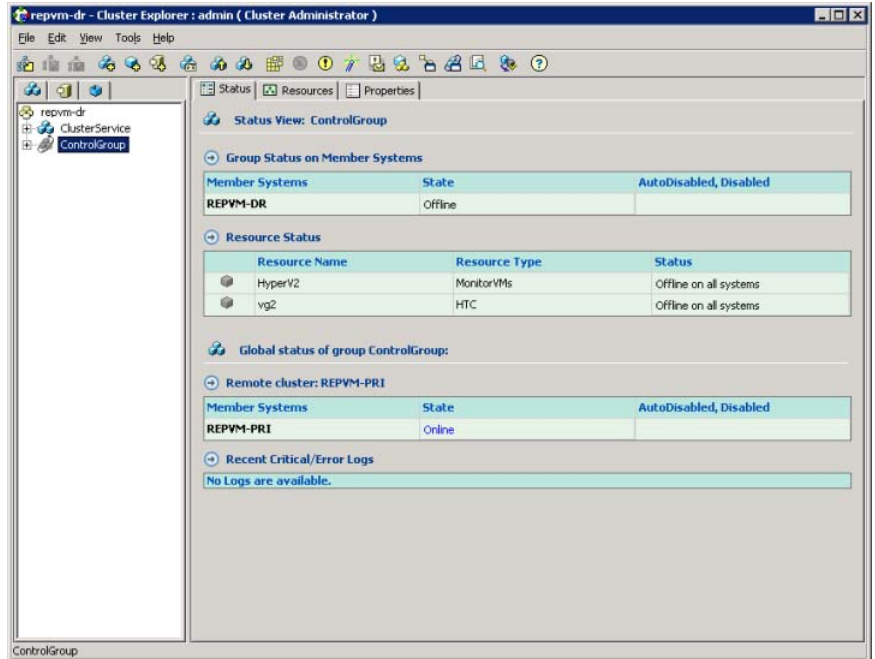
The **Declare Cluster** dialog box is displayed.

5 On the **Declare Cluster** dialog box, in the **Declare** field, select **Disaster**, and then select the name of the DR Management VM cluster and click **OK**.

In Cluster Explorer, verify that the **Status** view shows the DR Management VM cluster as **Online**.

6 Confirm that the application VMs are online on the DR site using the Windows failover cluster management console.

- 7 If the Alert message is not displayed, but the primary site is down, you can use Cluster Explorer to bring the DR Manager cluster online on the DR site, as follows:
 - In the Cluster Explorer configuration tree, select **ControlGroup**
The **Status** view shows the DR site Management VM as **Offline**, as shown in the illustration.



- 8 In the configuration tree, right-click **ControlGroup**, click **Online**, and select the name of the DR site Management VM.
- 9 In Cluster Explorer, verify that the **Status** view shows the DR site Management VM cluster as **Online**.
- 10 Confirm that the application VMs are online on the DR site using the Windows failover cluster management console.

- 11 If you configured automatic update of network settings, verify that the settings are updated. In some cases, upon the very first failover, Windows detects configuration changes in the application VM and displays a message prompting you to restart the computer to apply the changes. If settings are not updated, check for this message and restart the VM.
- 12 If VMs are configured on SFW storage and have not previously been failed over to the DR site, you must configure the resource dependency between the application VM resource and the VMDg resource.

See [“Setting the dependency of the virtual machine on the VMDg resource”](#) on page 41.

Hyper-V DR agent

This chapter includes the following topics:

- [About the Hyper-V DR agents](#)
- [MonitorVMs agent functions](#)
- [MonitorVMs agent state definitions](#)
- [MonitorVMs agent attribute definitions](#)
- [MonitorVMs agent resource type definition](#)

About the Hyper-V DR agents

Disaster Recovery Manager is a special feature of Veritas Cluster Server (VCS) for Windows. The installation includes the following Veritas Cluster Server (VCS) agents:

- Hardware replication agents
- The Monitor VMs agent (MonitorVMs) for Disaster Recovery Manager

The MonitorVMs agent monitors the health of application VMs that are configured for disaster recovery failover. The configuration wizard for the Disaster Recovery Manager creates a VCS global service group on the primary and DR sites. The configuration scripts then create the following resources and dependencies in the VCS global service group:

- A MonitorVM resource is created for all application VMs that you configure for Hyper-V disaster recovery.
- A replication resource is created for each storage resource associated with the application VMs.
- The MonitorVM resource depends on the storage replication resource.

The following information is provided for troubleshooting purposes. For more information on VCS agents and agent configuration, refer to the VCS documentation.

- See [“MonitorVMs agent functions”](#) on page 98.
- See [“MonitorVMs agent attribute definitions”](#) on page 99.
- See [“MonitorVMs agent resource type definition”](#) on page 99.

MonitorVMs agent functions

The MonitorVMs agent can monitor, start, and stop the application VMs and their associated storage by monitoring these resources and bringing them online or offline.

Table 7-1 MonitorVMs agent functions

Function	Description
Monitor	<p>If all VM resources in the Microsoft failover cluster are not in a FAILED state, report the resources as ONLINE.</p> <p>If the Microsoft failover cluster group containing the VM resources is in a failed state, report the resources as FAULTED.</p>
Online	<p>Bring the CSVs or physical disk resources online. If VMs are configured on SFW storage, bring SFW Volume Manager Disk Group (VMDg) resources online.</p> <p>If application VMs are not configured in the Microsoft cluster, then perform the steps necessary to import the VM configuration and create VM resources in the Microsoft cluster.</p> <p>Import the VM configuration using the documented API.</p> <p>Create the required resources in the Microsoft cluster on the DR site, using the add-ClusterVirtualMachineRole cmdlet.</p> <p>Bring online all VMs being monitored.</p>
Offline	<p>Bring offline all VMs and VM configurations that depend on the same CSV volume or physical disk resources.</p> <p>Bring offline the CSV volume or physical disk resources</p>

MonitorVMs agent state definitions

The following are the MonitorVMs agent state definitions:

ONLINE	All virtual machines (VMs) being monitored are running
OFFLINE	All VMs being monitored are not running
UNKNOWN	Some of the VMs are not running or in a failed state

MonitorVMs agent attribute definitions

The following table describes the MonitorVMs agent attribute definitions.

Table 7-2 MonitorVMs agent required attributes

Required attributes	Type and dimension	Definition
FOCClusterName	string-scalar	The name of the Microsoft failover cluster
ClusterNodes	string-vector	The name of the physical nodes in the Microsoft cluster
VMNames	string-association	The names of the application virtual machines
VMNetworkSettingsFile	string-scalar	The full path including file name of the network settings file

MonitorVMs agent resource type definition

The following is the MonitorVMs agent resource type definition:

```
type MonitorVMs (  
    static boolean IntentionalOffline = 1  
    static i18nstr ArgList[] = { ClusterName, ClusterNodes, VMNames, VMNetworkSettingsFile }  
    i18nstr FOCClusterName  
    i18nstr ClusterNodes[]  
    i18nstr VMNames{}  
    i18nstr VMNetworkSettingsFile
```

