# Symantec™ Storage Foundation and High Availability Solutions 6.1 Virtualization Guide - Linux

July 2014

# Symantec™ Storage Foundation and High Availability Solutions Virtualization Guide

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: 6.1

Document version: 6.1 Rev 4

## Legal Notice

Copyright © 2014 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.

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, whether delivered by Symantec as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

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 | supportsolutions@symantec.com |

## Documentation

Product guides are available on the media in PDF format. Make sure that you are using the current version of the documentation. The document version appears on page 2 of each guide. The latest product documentation is available on the Symantec website.

https://sort.symantec.com/documents

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

## 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

# Contents

Section 1

# Overview of SFHA Solutions used in Linux virtualization

Chapter 1

# Overview of supported products and technologies

This chapter includes the following topics:



## Overview of the Symantec Storage Foundation and High Availability Solutions Virtualization Guide

Virtualization technologies use software partitioning to provide a means of virtualizing operating system services. Partitioning enables the creation of isolated virtual machine environments for running applications. This isolation prevents processes running in one virtual machine from affecting processes running in other virtual machines. The virtualized computing environment is abstracted from all physical devices, enabling you to consolidate and centrally manage your workloads on a system.

This document provides information about Symantec Storage Foundation and High Availability (SFHA) Solutions support for Linux virtualization technologies. It contains:

- High-level conceptual information for SFHA Solutions and how they function in Linux virtual environments.
- High level implementation information for setting up SFHA products in Linux virtual environments.
- Use case chapters with examples of how SFHA Solutions can improve performance outcomes for common Linux virtualization use cases.

The information in this guide supplements rather than replaces SFHA Solutions product guides. It assumes you are a skilled user of Symantec products and knowledgeable concerning virtualization technologies.

See "Symantec Storage Foundation and High Availability Solutions product documentation" on page 164.

See "Linux virtualization documentation" on page 165.

# About Storage Foundation and High Availability Solutions support for Linux virtualization environments

Symantec Storage Foundation and High Availability (SFHA) Solutions products support the following virtualization technologies in Linux environments:

- Kernel-based Virtual Machine (KVM) technology for Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server (SLES)
- Red Hat Enterprise Virtualization (RHEV) environment
- Oracle Virtual Machine (OVM) environment
- Microsoft Hyper-V environment
- Linux guests in VMware ESXi environments

**Table 1-1** Supported Linux virtualization technologies

| Symantec product | KVM | RHEV | OVM | Microsoft Hyper-V | Linux in VMware ESXi |
|---|---|---|---|---|---|
| Symantec Dynamic Multi-Pathing (DMP) | Y | N | N | N | Y |

**Table 1-1** Supported Linux virtualization technologies *(continued)*

| Symantec product | KVM | RHEV | OVM | Microsoft Hyper-V | Linux in VMware ESXi |
|---|---|---|---|---|---|
| Symantec Storage Foundation (SF) | Y | N | N | N | Virtual machine only |
| Symantec Cluster Server (VCS) | Y | Y | Y | Virtual machine only | Virtual machine only |
| Symantec Storage Foundation and High Availability (SFHA) | Y | N | N | N | Virtual machine only |
| Symantec Storage Foundation Cluster File System High Availability (SFCFSHA) | Y | N | N | N | Virtual machine only |
| Symantec Replicator Option | Virtual machine only | N | N | N | Virtual machine only |
| Symantec ApplicationHA | Virtual machine on RHEL only | N | N | N | Virtual machine only |

For configuring SFHA Solutions in VMware guest environments, see the *Symantec Storage Foundation and High Availability Solutions Virtualization Guide for ESXi*.

For configuring DMP in VMware environments, see the *Symantec Dynamic Multi-Pathing Administrator's Guide for ESXi*.

For configuring Symantec ApplicationHA in VMware environments, see the *Symantec ApplicationHA User's Guide*

# About Kernel-based Virtual Machine (KVM) technology

The Symantec Storage Foundation and High Availability (SFHA) solutions can be used in Kernel-based Virtual Machine-based virtualization environments to provide advanced storage management, mission-critical clustering, fail-over, and migration capabilities.

Linux Kernel-based Virtual Machine (KVM) is released by Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server (SLES) as a full virtualization solution. KVM differs from other popular alternatives like Xen and VMware in terms of operation, performance and flexibility. KVM comes as a kernel module, with a set of user space utilities to create and manage virtual machines (VM).

Kernel-based Virtual Machine technology includes the following:

- A full virtualization solution for Linux on AMD64 & Intel 64 hardware.
- Each KVM virtualized guest or "VM guest" is run as a single Linux process.
- A hypervisor-independent virtualization API, `libvirt`, which provides a common generic and stable layer to securely manage VM guests on a host.
- A command line tool `virsh` used to manage the VM guests.
- A graphical user interface (GUI) `virt-manager` for managing the VM guests.
- Configuration of each VM guest stored in an XML file.

**Figure 1-1** KVM process

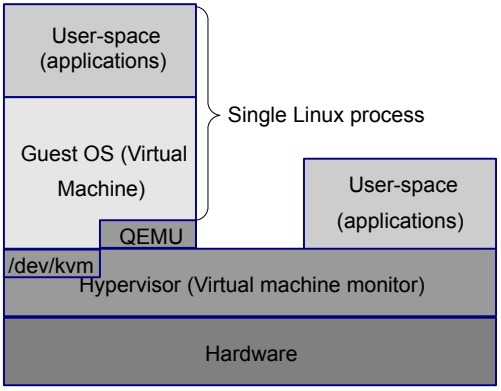


This guide illustrates some reference configurations which can be customized to fit most implementations. An assumption is made that the reader understands the Linux operating system, including its architecture, as well as how to configure and manage KVM virtual machines using the management software already provided by Linux. There is also an expectation that the user is familiar with the basic Symantec Storage Foundation and High Availability Solutions software and is well versed with its administration and management utilities. Additional details regarding Linux and Symantec Storage Foundation and High Availability Solutions software are available in the Additional documentation section.

See “Linux virtualization documentation” on page 165.

# Kernel-based Virtual Machine Terminology

**Table 1-2** KVM terminology used in this document

| Term | Definition |
|---|---|
| KVM | Kernel-based Virtual Machine |
| KVMGuest | VCS agent for managing virtual machines in a KVM or RHEV environment. |
| VM, KVM guest | Virtual machine, also referred to as a KVM virtualized guest. |
| Host | The physical host on which KVM is installed. |
| PM | The physical machine running VCS. |
| VM-VM | VCS-supported configuration in which a cluster is formed between VM guests running inside of the same or different hosts. |
| VM-PM | VCS-supported configuration in which a cluster is formed between VM guests and physical machines. |
| PM-PM | VCS-supported configuration in which a cluster is formed between hosts, and which is mainly used to manage VM guests running inside them. |
| Bridge | A device bound to a physical network interface on the host which enables any number of VM guests to connect to the local network on the host. It is mapped to a physical NIC which acts as a switch to VM guests. |
| VirtIO | VirtIO is an abstraction layer for paravirtualized hypervisors in Kernel-based Virtual Machine (VM) technology. |

# VirtIO disk drives

VirtIO is an abstraction layer for paravirtualized hypervisors in Kernel-based Virtual Machine (VM) technology. Unlike full virtualization, VirtIO requires special paravirtualized drivers running in each VM guest. VirtIO provides support for many devices including network devices and block (disk) devices. Using the VirtIO to export block devices to a host allows files, VxVM volumes, DMP meta-nodes, SCSI devices or any other type of block device residing on host to be presented to the VM guest. When SCSI devices are presented to a VM guest using VirtIO, in addition to simple reads and writes, SCSI commands such as SCSI inquiry commands can be performed allowing VxVM to perform deep device discovery. Running VxVM and DMP in the host and the VM guest provides for consistent naming of SCSI devices from the array, to the host through to the VM guest.

Symantec Storage Foundation and High Availability Solutions 6.1 supports VirtIO SCSI devices and VirtIO block devices with Linux KVM. virtio-scsi is a new virtual SCSI HBA interface. It is the foundation of an alternative storage implementation for virtual machines, replacing virtio-blk on Red Hat Enterprise Linux (RHEL) with improved scalability and providing standard SCSI command set support.

VirtIO features:

- Dynamically adding devices:
  VirtIO disk devices can be both added and removed from a running VM guest dynamically, without the need of a reboot.

VirtIO limitations:

- Disk caching:
  When disks are exported to the VM guest with the cache enabled, the VxVM configuration changes may get cached on the KVM host and not be applied to the disks. When disks are shared between more than one VM guest, such a configuration change is not visble from other VM guest systems than the one which made the change. To avoid potential configuration conflict, caching the host must be disabled (cache=no) while exporting the disks.
- SCSI Commands:
  SCSI devices that are presented as VirtIO devices to a VM guest support a limited subset of the SCSI command set. The KVM hypervisor blocks the restricted commands.
- PGR SCSI-3 Reservations:
  PGR SCSI-3 reservations are not supported on VirtIO block devices. To use SCSI-3 PR operations inside the KVM guest operating system, Symantec recommends that you use virtio-scsi to export SCSI devices to the guest. This limitation is applicable to releases prior to RHEL 6.4.
- DMP Fast Recovery with SCSI devices:
  DMP Fast Recovery bypasses the normal VirtIO read/write mechanism, performing SCSI commands directly against the device. If DMP Fast Recovery is used within the VM guest, caching in the host must be disabled (cache=none), to avoid data integrity issues.
- Thin Reclamation:
  Thin reclamation is not supported on VirtIO devices. The ‘WRITE-SAME’ command is blocked by the hypervisor. This limitation may be removed in future releases of Linux.
- Resizing devices:
  Linux does not support online disk resizing of VirtIO devices. To re-size a VirtIO device the VM guest must be fully shut down and re-started. Support for online re-sizing of block devices is under evaluation for Linux.

- Maximum number of devices:
  virtio-blk currently has a per-guest limitation of 32 devices. This device limitation includes all VirtIO devices, such as network interfaces and block devices. The device limitation is a result of the current VirtIO implementation where each device acts as a separate PCI device. virtio-scsi solves this limitation by multiplexing numerous storage devices on a single controller. Each device on a virtio-scsi controller is represented as a logical unit, or LUN. The LUNs are grouped into targets. The device limit per target is much larger; each device can have a maximum of 256 targets per controller and 16,384 logical units per target. You can use virtio-scsi instead of virtio-blk to use more than 32(28) disk devices inside the KVM guest.
- VxFS:
  In a KVM environment under heavy I/O load, data corruption may occur on VxFS file systems created on LUNs attached as VirtIO block devices. Please refer Red Hat Support Case #00945974 for more details:
  https://access.redhat.com/support/cases/00945974

# About the RHEV environment

Red Hat Enterprise Virtualization consists of the following components:

- Red Hat Enterprise Virtualization Hypervisor:
  This is a thin hypervisor layer, which is based on Kernel-based Virtual Machine (KVM). As KVM forms a core part of the Linux kernel, it proves to be a very efficient virtualization option.
- Agents and tools:
  These include bundled as well as application-specific agents, and Virtual Desktop Server Manager (VDSM) that runs in the hypervisor. Together, the agents and tools help you administer the virtual machines and the related network and storage.
- Red Hat Enterprise Virtualization platform management infrastructure:
  This provides the interface to view and manage all the system components, machines and images. This management infrastructure provides powerful search capabilities, resource management, live migration, and provisioning.

# RHEV terminology

**Table 1-3** RHEV terminology used in this document

| Term | Definition |
|---|---|
| KVM | Kernel-based Virtual Machine. |
| KVMGuest | VCS agent for managing virtual machines in a KVM or RHEV environment. |
| VM | Virtual machine created in a KVM or RHEV environment. |
| Host | The physical host on which the virtual machine is created or running. |
| PM | The physical machine running VCS. |
| PM-PM | VCS-supported configuration in which a cluster is formed between hosts, and which is mainly used to manage VM guests running inside them. |
| RHEV | Red Hat Enterprise Virtualization. |
| RHEV-M | Red Hat Enterprise Virtualization Manager is a centralized management console for managing the RHEV environment. |
| RHEL-H | Red Hat Enterprise Linux (RHEL) host that runs a complete version of RHEL, and is managed by RHEV-M. |
| RHEV-H | Red Hat Enterprise Virtualization - Hypervisor is a minimal installation of Red Hat Enterprise Linux, which supports the creation and operaiton of virtual machines. |
| VDSM | Virtual Desktop Server Manager. The VDSM service is used by RHEV-M to manage the RHEV-H and RHEL hosts. |
| REST API | Representational state transfer (REST) API. |
| Datacenter | A datacenter is a logical entity in a RHEV-M that defines the set of physical and logical resources used in a managed virtual environment such as clusters of hosts, virtual machines, storage and networks. |
| Cluster | This is a cluster in RHEV-M. A cluster is a collection of physical hosts that share the same storage domains and have the same type of CPU. |
| Storage Domain | This is the storage infrastructure in RHEV for creating and running virtual machines. |

**Table 1-3** RHEV terminology used in this document *(continued)*

| Term | Definition |
|---|---|
| Data Domain | A type of storage domain that holds the disk image of all the virtual machines running in the system, operating system images, and data disks. |
| ISO Domain | This domain stores ISO files (or logical CDs) used to install and boot operating systems and applications for the virtual machines. |

For more information on Red Hat Enterprise Virtualization, see Red Hat Enterprise Virtualization documentation.

# About virtual-to-virtual (in-guest) clustering and failover

When you run Symantec Cluster Server (VCS) in multiple guest virtual machines, you can create guest-to-guest (also called virtual-to-virtual) clusters. You can use VCS to monitor individual applications running inside each guest. In case of application failure, you can fail over the application to another guest virtual machine in the virtual-to-virtual cluster.

The following figure illustrates a sample in-guest VCS deployment in one virtual machine each across two physical hosts.

**Figure 1-2** VCS in-guest clustering

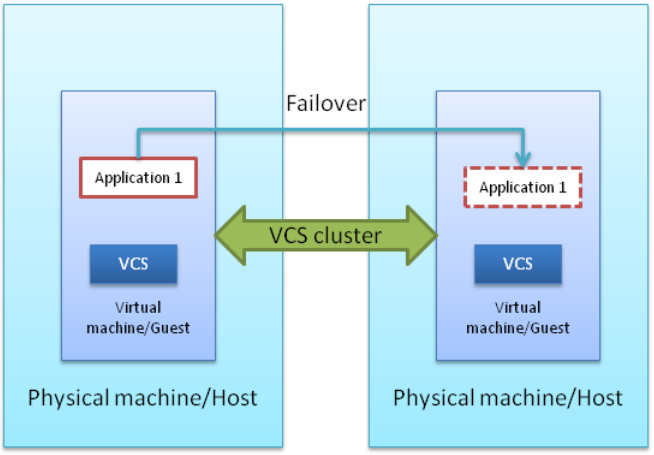

The virtual machines in the cluster can either be on the same physical host or on different physical hosts. VCS is installed in the virtual machines and creates a cluster. This is just like the cluster that VCS creates among physical systems. The cluster monitors the applications and services that run inside the virtual machines. Any faulted application or service is failed over to another virtual machine in the cluster.

To ensure application failover, application data must reside on storage shared by member virtual machines within the cluster.

---

**Note:** In this configuration, since VCS runs inside a virtual machine, VCS cannot fail over the virtual machine itself.

---

VCS can be deployed inside guest virtual machines (in-guest support) in the following virtualization environments:

- Microsoft Hyper-V
- Red Hat Enterprise Virtualization (RHEV)
- Oracle Virtual Machine (Oracle VM)
- Kernel-based Virtual Machine (KVM) technology for Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server (SLES)
- Linux guests in VMware ESXi environments

# About the Symantec Storage Foundation and High Availability Solutions products

Symantec Storage Foundation and High Availability (SFHA) Solutions is a set of products that provide storage administration and management in a heterogeneous storage environment.

This section can help you determine which product you need.

Table 1-4 shows the benefits of each product and its components.

**Table 1-4** SFHA Solutions product comparisons

| Product | Components | Benefits |
|---|---|---|
| Symantec Cluster Server (VCS) connects multiple, independent systems into a management framework for increased availability. Each system, or node, runs its own operating system and cooperates at the software level to form a cluster. VCS links commodity hardware with intelligent software to provide application failover and control. When a node or a monitored application fails, other nodes can take predefined actions to take over and bring up services elsewhere in the cluster. | VCS | ■ Minimizes downtime<br>■ Facilitates the consolidation and the failover of servers<br>■ Effectively manages a wide range of applications in heterogeneous environments<br>■ Provides data integrity protection through I/O fencing<br>■ Provides High Availability of applications |
| Dynamic Multi-Pathing (DMP) provides multi-pathing functionality for the devices configured on the system. The product creates DMP metadevices (also known as DMP nodes) to represent all the device paths to the same physical LUN. | DMP | ■ Extends DMP metadevices to support OS native logical volume managers (LVM)<br>■ Provides improved storage I/O performance with load balancing<br>■ Provides storage path failure protection and fast failover<br>■ Centralizes storage path management regardless of operating system or storage hardware |
| Symantec Replicator Option enables cost-effective replication of data over IP networks, giving organizations an extremely flexible, storage hardware independent alternative to traditional array-based replication architectures. | VVR<br>VFR | Volume Replicator (VVR)<br>■ Provides block-based continuous replication<br>■ Provides effective bandwidth management<br>■ Supports cross-platform replication, and replication in a Portable Data Container (PDC) environment<br>File Replicator (VFR)<br>■ Provides file-based periodic replication<br>■ Supports reversible data transfer<br>■ Deduplication<br>■ Supports protection of the target file system from accidental writes |

**Table 1-4** SFHA Solutions product comparisons *(continued)*

| Product | Components | Benefits |
|---|---|---|
| Storage Foundation (SF) is a storage management offering that consists of Veritas Volume Manager (VxVM), Veritas File System (VxFS), and DMP.<br><br>Veritas Volume Manager is a storage management subsystem that enables you to manage physical disks and logical unit numbers (LUNs) as logical devices called volumes.<br><br>Veritas File System is an extent-based, intent logging file system. | DMP, VxVM, VxFS | ■ Increased storage utilization across heterogeneous environments<br>■ Deduplication and compression<br>■ Automated storage tiering<br>■ Centralized storage management<br>■ Easy OS and storage migration with minimum downtime<br>■ All benefits of DMP |
| Storage Foundation High Availability (SFHA) includes all the functionality of SF plus the high availability of VCS. | DMP, VxVM, VxFS, VCS | ■ All benefits of DMP<br>■ All benefits of SF<br>■ All benefits of VCS |
| Storage Foundation Cluster File System High Availability (SFCFSHA) extends Symantec Storage Foundation to support shared data in a storage area network (SAN) environment. Multiple servers can concurrently access shared storage and files transparently to applications.<br><br>Cluster Volume Manager (CVM) extends VxVM to support shared disk groups. Cluster File System (CFS) extends VxFS to support parallel clusters. | DMP, VxVM, VxFS, VCS, CVM, SFCFSHA | ■ All benefits of DMP<br>■ All benefits of SF<br>■ All benefits of VCS<br>■ Increased automation and intelligent management of availability and performance across shared storage |

**Table 1-4** SFHA Solutions product comparisons *(continued)*

| Product | Components | Benefits |
|---|---|---|
| Symantec ApplicationHA provides monitoring capabilities for applications running inside virtual machines in the virtualization environment. Symantec ApplicationHA adds a layer of application awareness to the core high availability functionality that is offered by Symantec Cluster Server (VCS) in the physical host. Symantec ApplicationHA is based on VCS, and uses similar concepts such as agents, resources, and service groups. However, Symantec ApplicationHA has a lightweight server footprint that enables faster installation and configuration in virtualization environments. | Symantec ApplicationHA, VCS | ■ Out of the box integration with VCS<br>■ Full visibility and control over applications with the ability to start, stop, and monitor applications running inside virtual machines<br>■ High availability of the application as well as the virtual machine inside which the application runs<br>■ A series of incremental fault management remedies that include application restart, graceful virtual machine reboot, forceful virtual machine reboot, and vMotion. ApplicationHA tries the first two remedies and upon failure prompts VMwareHA to try the next two remedies.<br>■ Standardized way to manage applications using a single interface that is integrated with the Veritas Operations Manager console<br>■ Specialized Application Maintenance mode, in which Symantec ApplicationHA enables you to intentionally take an application out of its purview for maintenance or troubleshooting |
| Veritas Operations Manager provides a centralized management console for Symantec Storage Foundation and High Availability products. You can use Veritas Operations Manager to monitor, visualize, and manage storage resources and generate reports. | All | ■ Centralized, standardized way to manage the various features in the Storage Foundation products<br>■ Visual interface for managing individual hosts and their storage<br>■ Visibility into all instances of Storage Foundation that are running in the datacenter, across multiple operating systems |

**Table 1-4** SFHA Solutions product comparisons *(continued)*

| Product | Components | Benefits |
|---|---|---|
| Symantec Cluster Server (VCS) agents provide high availability for specific resources and applications. Each agent manages resources of a particular type. Typically, agents start, stop, and monitor resources and report state changes.<br><br>In addition to the agents that are provided in this release, other agents are available through an independent Symantec offering called the Symantec High Availability Agent Pack. The agent pack includes the currently shipping agents and is re-released quarterly to add the new agents that are now under development.<br><br>You can download the latest agents from the Symantec Operations Readiness (SORT) website at:<br><br>https://sort.symantec.com/agents | VCS, Symantec ApplicationHA | ■ All benefits of VCS<br>■ All benefits of Symantec ApplicationHA |

# Virtualization use cases addressed by Storage Foundation and High Availability Solutions

Storage Foundation and High Availability (SFHA) Solutions products support the following virtualization environment use cases:

**Table 1-5** Virtualization use cases addressed by SFHA Solutions in a Linux environment

| Virtualization use case | Recommended SFHA Solution | Virtualization technology supported | Implementation details |
|---|---|---|---|
| Server consolidation | SFHA or SFCFSHA in the guest | Red Hat Enterprise Linux (RHEL) KVM<br><br>SUSE Linux Enterprise Server (SLES) KVM | How to run virtual machines as physical servers.<br><br>See "Server consolidation" on page 92. |

**Table 1-5** Virtualization use cases addressed by SFHA Solutions in a Linux environment *(continued)*

| Virtualization use case | Recommended SFHA Solution | Virtualization technology supported | Implementation details |
|---|---|---|---|
| Physical to virtual migration | SF in the host<br>SFHA or SFCFSHA | RHEL KVM<br>SLES KVM | How to migrate data from physical to virtual environments safely and easily.<br>See "Physical to virtual migration" on page 95. |
| Simplified management | SFHA or SFCFSHA in the host | RHEL KVM<br>SLES KVM | How to manage virtual machines using the same command set, storage namespace, and environment as in a non-virtual environment.<br>See "Simplified management" on page 99. |
| Application monitoring | ApplicationHA in the guest | RHEL kVM<br>Linux on VMware ESXi | How to manage application monitoring on virtual machines.<br>See "About application monitoring using Symantec ApplicationHA" on page 107.<br>See ApplicationHA documentation. |
| Application failover | VCS or SFHA in the guest | RHEL KVM<br>Red Hat Enterprise Virtualization (RHEV)<br>SLES KVM<br>Linux on VMware ESXi | How to manage application monitoring on virtual machines.<br>See ApplicationHA documentation.<br>How to manage application failover on virtual machines.<br>See "Symantec Cluster Server In a KVM Environment Architecture Summary" on page 116. |

**Table 1-5** Virtualization use cases addressed by SFHA Solutions in a Linux environment *(continued)*

| Virtualization use case | Recommended SFHA Solution | Virtualization technology supported | Implementation details |
|---|---|---|---|
| Virtual-to-virtual (in-guest) clustering | VCS in the guest | RHEL KVM<br>RHEV<br>SLES KVM<br>Microsoft Hyper-V<br>Linux on VMware ESXi<br>Oracle Virtual Machine (OVM) | How to configure VCS for virtual-to-virtual clustering.<br>See "Installing and configuring Symantec Cluster Server for Red Hat Enterprise Virtualization (RHEV) virtual-to-virtual clustering" on page 128.<br>See "Installing and configuring Symantec Cluster Server for Red Hat Enterprise Virtualization (RHEV) virtual-to-virtual clustering" on page 128.<br>See " Installing and configuring Symantec Cluster Server with Microsoft Hyper-V virtual-to-virtual clustering" on page 131.<br>See "Installing and configuring Symantec Cluster Server for Oracle Virtual Machine (OVM) virtual-to-virtual clustering" on page 133. |
| Virtual machine availability | VCS in the host | RHEL KVM<br>RHEV<br>SLES KVM | How to manage virtual machine failover.<br>See "VCS in host monitoring the Virtual Machine as a resource" on page 121. |
| Virtual machine Live Migration | SFCFSHA in the host | RHEL KVM<br>SLES KVM | How to use features such as instant snapshots to contain boot images and manage them from a central location in the host.<br>See "About live migration" on page 123. |
| Virtual machine Live Migration | VCS in the host | RHEV | How to use features such as instant snapshots to contain boot images and manage them from a central location in the host.<br>See "About live migration" on page 123. |

**Table 1-5** Virtualization use cases addressed by SFHA Solutions in a Linux environment *(continued)*

| Virtualization use case | Recommended SFHA Solution | Virtualization technology supported | Implementation details |
|---|---|---|---|
| Disaster recovery (DR) in the virtual environment | VCS in the host | RHEV | How to configure virtual machines for disaster recovery.<br>See "About disaster recovery for Red Hat Enterprise Virtualization virtual machines" on page 136. |
| Application to storage visibility | Configuration for Veritas Operations Manager (VOM) use case | RHEL KVM<br>SLES KVM<br>Linux on VMware ESXi<br>Microsoft Hyper-V | How to configure for storage to application visibility.<br>See "About storage to application visibility using Veritas Operations Manager" on page 88. |
| Multi-tier Business service support | VOM, Virtual Business Service (VBS) | RHEL KVM<br>SLES KVM | How to discover and configure devices for multi-tier application.<br>See "About Virtual Business Services" on page 144. |

**Note:** Symantec ApplicationHA is supported in the RHEL KVM environment only.

Section 2

# Implementing a basic KVM environment

Chapter 2

# Getting started with basic KVM

This chapter includes the following topics:



## Storage Foundation and High Availability Solutions configuration options for the kernel-based virtual machines environment

Symantec Storage Foundation and High Availability Solutions (SFHA Solutions) products support the configurations listed in Table 2-1. The configurations profiled in the table below are the minimum required to achieve the storage and availability

objectives listed. You can mix and match the use of SFHA Solutions products as needed to achieve the desired level of storage visibility, management, replication support using VVR, availability, and cluster failover for your kernel-based virtual machines (KVM) hosts and guest virtual machines.

**Table 2-1** Storage Foundation and High Availability Solutions supported configuration options in the KVM environment

| Objective | Recommended SFHA Solutions product configuration |
|---|---|
| Storage visibility for KVM guest virtual machines | Dynamic Multi-Pathing (DMP) in the KVM guest virtual machines |
| Storage visibility for KVM hosts | DMP in the KVM hosts |
| Storage management features and replication support using VVR for KVM guest virtual machines | Storage Foundation (SF) in the KVM guest virtual machines<br>See "Symantec Storage Foundation in the virtualized guest machine" on page 37. |
| Advanced storage management features and replication support using VVR for KVM hosts | Storage Foundation Cluster File System (SFCFSHA) in the KVM hosts |
| End-to-end storage visibility in KVM hosts and guest virtual machines | DMP in the KVM host and guest virtual machines |
| Storage management features and replication support using VVR in the KVM guest virtual machines and storage visibility in in the KVM host | DMP in the KVM host and SF in the KVM guest virtual machines<br>See "Symantec Dynamic Multi-Pathing in the KVM host and Symantec Storage Foundation HA in the KVM guest virtual machine" on page 40. |
| Application monitoring and availability for KVM guest virtual machines | Symantec ApplicationHA in the KVM guest virtual machines<br>See "Symantec ApplicationHA in the KVM virtualized guest machine" on page 41. |
| Virtual machine monitoring, migration, and failover for KVM hosts | Symantec Cluster Server (VCS) in the KVM hosts<br>See "Symantec Cluster Server in the KVM host" on page 42. |

**Table 2-1** Storage Foundation and High Availability Solutions supported configuration options in the KVM environment *(continued)*

| Objective | Recommended SFHA Solutions product configuration |
|---|---|
| Application failover for KVM guest virtual machines | VCS in the KVM guest virtual machines<br>See "Symantec Cluster Server in the guest" on page 43. |
| Application availability and virtual machine availability | Symantec ApplicationHA in the KVM guest virtual machines and VCS in the KVM host<br>See "Symantec ApplicationHA in the guest and Symantec Cluster Server in the host" on page 44. |
| Application failover across KVM guest virtual machines and physical hosts | VCS in KVM guest virtual machines and KVM physical host machines<br>See "Symantec Cluster Server in a cluster across virtual machine guests and physical machines" on page 45. |

**Note:** Symantec ApplicationHA is supported in the Red Hat Enterprise Linux (RHEL) KVM environment only.

See "About setting up KVM with Symantec Storage Foundation and High Availability Solutions" on page 46.

See "Virtualization use cases addressed by Storage Foundation and High Availability Solutions" on page 28.

## Symantec Dynamic Multi-Pathing in the KVM guest virtualized machine

Use Symantec Dynamic Multi-Pathing (DMP) to provide storage visibility in KVM guest virtualized machines. DMP in the KVM guest virtualized machine provides:

- Multi-pathing functionality for the operating system devices configured in the guest
- DMP metadevices (also known as DMP nodes) to represent all the device paths to the same physical LUN
- Support for enclosure-based naming
- Support for standard array types

**Figure 2-1** Symantec Dynamic Multi-Pathing in the guest

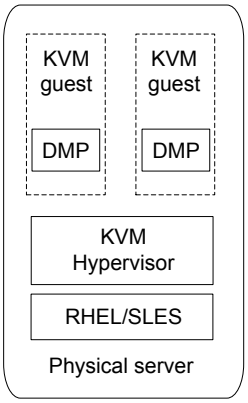


For more information on DMP features, see the *Symantec Dynamic Multi-Pathing Administrator's Guide*.

## Symantec Dynamic Multi-Pathing in the KVM host

Use Symantec Dynamic Multi-Pathing (DMP) to provide storage visibility in the KVM hosts. Using DMP in the KVM host enables:

- Centralized multi-pathing functionality
- Enables active/passive array high performance failover
- Centralized storage path management
- Fast proactive failover
- Event notification

**Figure 2-2** Symantec Dynamic Multi-Pathing in the KVM host

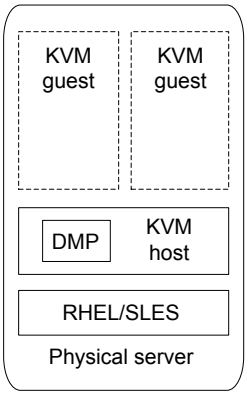


For more information on DMP features, see the *Symantec Dynamic Multi-Pathing Administrator's Guide*.

## Symantec Storage Foundation in the virtualized guest machine

Use Symantec Storage Foundation (SF) in the guest to provide storage management functionality for KVM guest virtual machine resources. Symantec Storage Foundation enables you to manage KVM guest storage resources more easily by providing:

- Enhanced database performance
- Point-in-time copy features for data back-up, recovery, and processing
- Options for setting policies to optimize storage
- Methods for migrating data easily and reliably
- Replication support

**Figure 2-3** Symantec Storage Foundation in the virtualized guest machine

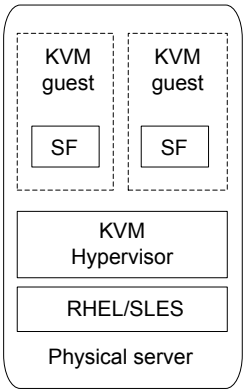


For more information on Symantec Storage Foundation features, see the *Symantec Storage™ Foundation Administrator's Guide*.

## Symantec Storage Foundation Cluster File System High Availability in the KVM host

Use Symantec Storage Foundation Cluster File System High Availability (SFCFSHA) to provide advanced storage management functionality for the KVM host. SFCFSHA enables you to manage your KVM host storage resources more easily by providing:

- Enhanced database performance
- Point-in-time copy features for data back-up, recovery, and processing
- Options for setting policies to optimize storage
- Methods for migrating data easily and reliably
- Replication support
- High availability for virtual machines
- High availability and disaster recovery for virtual machines
- Simplified management of virtual machines

**Figure 2-4** Symantec Storage Foundation Cluster File System High Availability in the KVM host

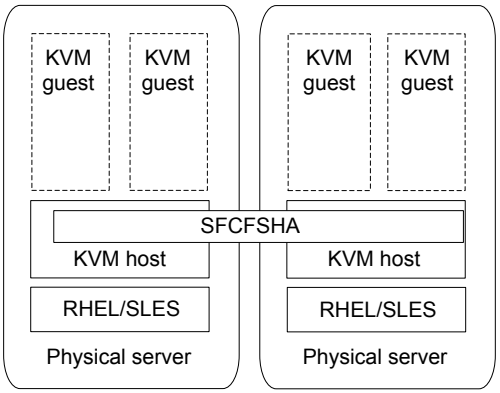


For more information on Storage Foundation features, see the *Symantec Storage Foundation™ Cluster File System High Availability Administrator's Guide*.

# Symantec Dynamic Multi-Pathing in the KVM host and guest virtual machine

Use Symantec Dynamic Multi-Pathing (DMP) to provide end-to-end storage visibility across both the KVM host and guest virtual machine. Using DMP in the KVM guest virtualized machine provides:

- Multi-pathing functionality for the operating system devices configured in the guest
- DMP metadevices (also known as DMP nodes) to represent all the device paths to the same physical LUN
- Support for enclosure-based naming
- Support for standard array types

Using DMP in the KVM host enables:

- Centralized multi-pathing functionality
- Enables active/passive array high performance failover
- Centralized storage path management
- Fast proactive failover
- Event notification

**Figure 2-5** Symantec Dynamic Multi-Pathing in the KVM virtualized guest and the KVM host

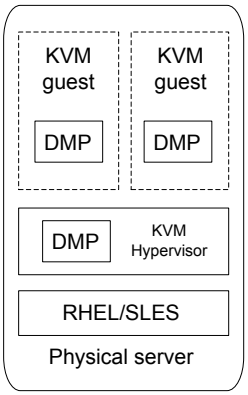


For more information on DMP features, see the *Symantec Dynamic Multi-Pathing Administrator's Guide*.

## Symantec Dynamic Multi-Pathing in the KVM host and Symantec Storage Foundation HA in the KVM guest virtual machine

Use Symantec Storage Foundation and High Availability (SFHA) in the guest in combination with Dynamic Multi-Pathing (DMP) in the KVM host to combine storage management functionality for KVM guest virtual machine resources and storage visibility in the KVM host.

Using SFHA in the KVM guest provides:

- Enhanced database performance
- Point-in-time copy features for data back-up, recovery, and processing
- Options for setting policies to optimize storage
- Methods for migrating data easily and reliably
- Replication support
- High availability for applications running inside virtual machines

Using DMP in the host provides:

- Centralized multi-pathing functionality
- Fast proactive failover.
- Event notification

**Figure 2-6** Symantec Storage Foundation HA in the KVM guest virtual machine and DMP in the KVM host

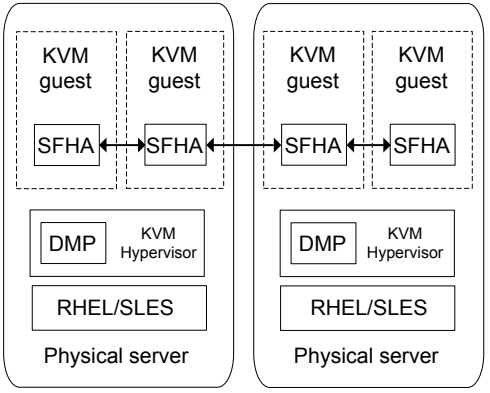


For more information on SFHA features, see the *Symantec Storage Foundation Cluster File System High Availability Administrator's Guide*.

For more information on DMP features, see the *Symantec Dynamic Multi-Pathing Administrator's Guide*.

## Symantec ApplicationHA in the KVM virtualized guest machine

Use Symantec ApplicationHA to enable configuration of KVM virtualized guest resources for application failover. ApplicationHA provides the following for KVM virtualized guest machines:

- Full visibility and control over applications with the ability to start, stop, and monitor applications running inside virtual machines.
- Graded application fault-management responses such as:
  - Application restart
  - ApplicationHA-initiated, internal or soft reboot of a Virtual Machine
- Standardized way to manage applications using a single interface that is integrated with the Veritas Operations Manager (VOM) dashboard
- Specialized Application Maintenance mode, in which ApplicationHA enables you to intentionally take an application out of its purview for maintenance or troubleshooting

**Figure 2-7** Symantec ApplicationHA in the virtualized guest machine

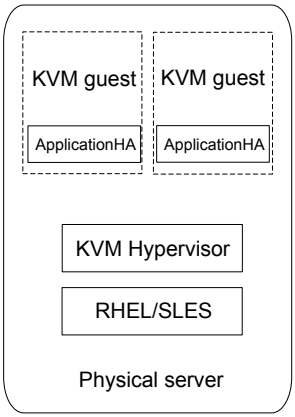


**Note:** Symantec ApplicationHA is supported only in the Red Hat Enterprise Linux (RHEL) KVM environment.

For more information on Symantec ApplicationHA features, see the *Symantec™ ApplicationHA User's Guide*.

## Symantec Cluster Server in the KVM host

Use Symantec Cluster Server (VCS) to provide virtual machine monitoring and failover to another KVM host. VCS enables the following for KVM hosts:

- Connects multiple, independent systems into a management framework for increased availability.
- Enables nodes to cooperate at the software level to form a cluster.
- Links commodity hardware with intelligent software to provide application failover and control.
- Enables other nodes to take predefined actions when a monitored application fails, for instance to take over and bring up applications elsewhere in the cluster.

**Figure 2-8** Symantec Cluster Server in the KVM host

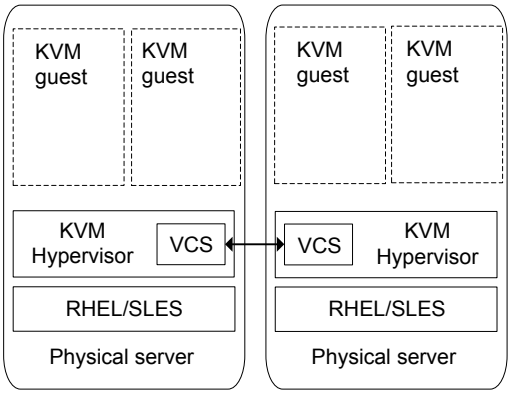


For more information on Symantec Cluster Server features, see the *Symantec Cluster Server Administrator's Guide*.

## Symantec Cluster Server in the guest

Use Symantec Cluster Server (VCS) to provide application monitoring and failover to another KVM guest.

- Connects multiple, independent systems into a management framework for increased availability
- Enables nodes to cooperate at the software level to form a cluster
- Links commodity hardware with intelligent software to provide application failover and control
- Enables other nodes to take predefined actions when a monitored application fails, for instance to take over and bring up applications elsewhere in the cluster

**Figure 2-9** Symantec Cluster Server in the guest

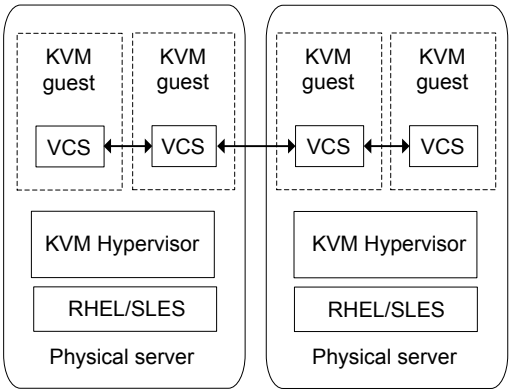


For more information on Symantec Cluster Server features, see the *Symantec Cluster Server Administrator's Guide*.

## Symantec ApplicationHA in the guest and Symantec Cluster Server in the host

Use Symantec ApplicationHA in the KVM virtualized guest in combination with Symantec Cluster Server (VCS) in the KVM host to provide the following:

- Full visibility and control over applications with the ability to start, stop, and monitor applications running inside virtual machines.
- High availability of the application as well as the virtual machine on which the application runs.
- Graded application fault-management responses such as:
  - Application restart
  - ApplicationHA-initiated, internal or soft reboot of a KVM virtualized guest machine
- VCS-initiated or hard reboot of virtual machine or failover of the KVM virtual machine to another physical host
- Standardized way to manage applications using a single interface that is integrated with the Veritas Operations Manager (VOM) dashboard
- Specialized Application Maintenance mode, in which ApplicationHA enables you to intentionally take an application out of its purview for maintenance or troubleshooting
- VCS in the host enables virtual machine availability

- ApplicationHA monitors the applications running inside the guest
- ApplicationHA configured in the guest restarts the application in case of application fault
- ApplicationHA can notify VCS running in the host to trigger a virtual machine failover

**Figure 2-10** Symantec ApplicationHA in the guest and Symantec Cluster Server in the host

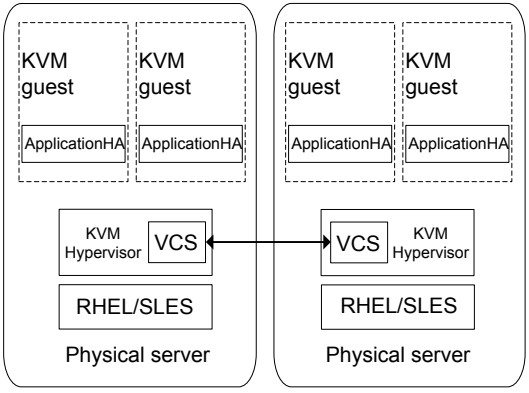


**Note:** Symantec ApplicationHA is supported only in the Red Hat Enterprise Linux (RHEL) KVM environment.

For more information on Symantec ApplicationHA features, see the *Symantec ApplicationHA User's Guide*.For more information on Symantec Cluster Server features, see the *Symantec Cluster Server Administrator's Guide*.

## Symantec Cluster Server in a cluster across virtual machine guests and physical machines

Use Symantec Cluster Server (VCS) in both the guest and host to enable an integrated solution for resource management across virtual machines and physical hosts. You can create a physical to virtual cluster combining VCS in a KVM guest together with VCS running on another physical host, enabling VCS to:

- Monitor applications running within the guest
- Failover applications to another physical host
- Failover an application running on a physical host to a VM virtualized guest machine

**Figure 2-11** Symantec Cluster Server in a cluster across guests and physical machines

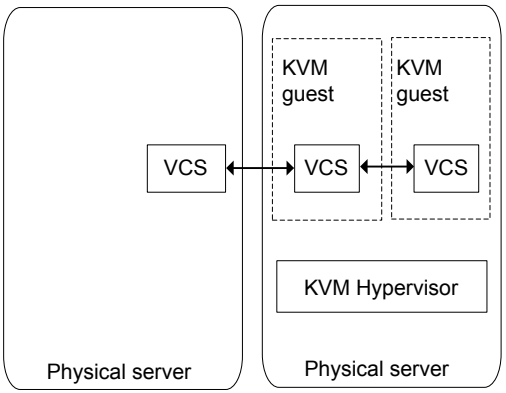


For more information on Storage Foundation features, see the *Symantec Cluster Server Administrator's Guide*.

# About setting up KVM with Symantec Storage Foundation and High Availability Solutions

Before setting up your virtual environment, verify that your planned configuration will meet the system requirements, licensing and other considerations for installation with Symantec Storage Foundation and High Availability (SFHA) Solutions products.

- Licensing: customers running Symantec Storage Foundation (SF) or Symantec Storage Foundation Cluster File System High Availability (SFCFSHA) in a kernel-based virtual machine (KVM) environment are entitled to use an unlimited number of guests on each licensed server or CPU.
- Red Hat and SUSE system requirements: see Table 2-2
- Symantec product requirements: see Table 2-3
- *Release Notes*: each Symantec product contains last minute news and important details for each product, including updates to system requirements and supported software. Review the *Release Notes* for the latest information before you start installing the product.
  The product documentation is available on the Web at the following location:
  https://sort.symantec.com/documents

**Table 2-2** Red Hat and SUSE system requirements

| | Red Hat Enterprise Linux (RHEL) | SUSE Linux Enterprise Server (SLES) |
|---|---|---|
| Supported architecture | ■ Intel 64<br>■ AMD 64 | ■ Intel 64<br>■ AMD 64 |
| Minimum system requirement | ■ 6 GB free disk space<br>■ 2 GB of RAM | ■ 6 GB free disk space<br>■ 2 GB of RAM |
| Recommended system requirement | ■ 6 GB plus the required disk space recommended by the guest operating system per guest. For most operating systems more than 6 GB of disk space is recommended<br>■ One processor core or hyper-thread for each virtualized CPU and one for the host<br>■ 2 GB of RAM plus additional RAM for virtualized guests | ■ 6 GB plus the required disk space recommended by the guest operating system per guest. For most operating systems more than 6 GB of disk space is recommended<br>■ One processor core or hyper-thread for each virtualized CPU and one for the host<br>■ 2 GB of RAM plus additional RAM for virtualized guests |
| Hardware requirement | Full virtualization-enabled CPU | Full virtualization-enabled CPU |
| Symantec Storage Foundation and High Availability Solutions version | 6.1 | 6.1 |
| Supported OS version in the host | RHEL 6 Update 3, Update 4 | SLES11 SP2, SP3 |
| Supported OS version in the virtual machine | RHEL 5 Update 4, Update5, Update 6, Update 7, Update 8, Update 9<br>RHEL 6 Update 3, Update 4 | SLES11 SP2, SP3 |

**Table 2-3** Symantec product requirements

| | |
|---|---|
| Hardware | http://www.symantec.com/docs/TECH211575 |

**Table 2-3** Symantec product requirements *(continued)*

| | |
|---|---|
| Software | ■ Symantec Dynamic Multi-Pathing 6.1<br>Used for storage visibility on KVM hosts and guest virtual machines<br>■ Symantec Storage Foundation 6.1<br>Used for storage management on KVM hosts and guest virtual machines<br>■ Symantec Storage Foundation HA 6.1<br>Used for storage management and clustering on KVM hosts and guest virtual machines<br>■ Storage Foundation Cluster File System High Availability 6.1<br>Used for storage management and clustering multiple KVM hosts to enable live migration of guest virtual machines<br>■ Symantec Cluster Server 6.1<br>Used for virtual machine monitoring, migration, and failover<br>■ Symantec ApplicationHA 6.0<br>Used for application monitoring and availability<br>■ Veritas Operations Manger 5.0<br>Used for application visibility and virtual host management |
| Storage | ■ Shared storage for holding the guest image. (VM failover)<br>■ Shared storage for holding the application data. (Application failover) |
| Networking | ■ Configure the guest for communication over the public network<br>■ Setup virtual interfaces for private communication. |
| Documentation: see the product release notes to for the most current system requirements, limitations, and known issues: | ■ *Symantec Dynamic Multi-Pathing Release Notes*<br>■ *Symantec Storage Foundation Release Notes*<br>■ *Symantec Storage Foundation High Availability Release Notes*<br>■ *Symantec Storage Foundation Cluster File System High Availability Release Notes*<br>■ *Symantec Cluster Server Release Notes*<br>■ *Symantec ApplicationHA Release Notes*<br>■ Symantec Operations Readiness Tools:<br>https://sort.symantec.com/documents<br>■ Storage Foundation DocCentral Site:<br>http://sfdoccentral.symantec.com/ |

**Table 2-4** VCS system requirements for KVM-supported Red Hat Enterprise Linux configurations

| | |
|---|---|
| VCS version | 6.1 |

**Table 2-4** VCS system requirements for KVM-supported Red Hat Enterprise Linux configurations *(continued)*

| | |
|---|---|
| Supported OS version in host | RHEL 6 Update 3, Update 4 |
| Supported OS in VM guest | RHEL 5 Update 4, Update5, Update 6, Update 7, Update 8, Update 9<br>RHEL 6 Update 3, Update 4 |
| Hardware requirement | Full virtualization-enabled CPU |

### Limitations and unsupported kernel-based virtual machine features

The DiskReservation agent cannot work with disks exported over a VirtIO bus.

For more information on limitations and known issues, see the *Symantec Cluster Server Release Notes* for Linux.

For KVM related limitations, see the Virtualization technology provider (RHEL or SLES) release notes.

See "Linux virtualization documentation" on page 165.

# Creating and launching a kernel-based virtual machine (KVM) host

KVM is available as part of Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server (SLES). Management for RHEL KVM is provided through separate RPMs that can be downloaded into the standard RHEL installation. Management for SLES KVM is provided through SLES or through separate RPMs that can be downloaded into the standard SLES installation.

The `virt-manager` tool provides a very simple, easy-to-use and intuitive GUI interface for all virtual machine operations, along with `virt-viewer`. A command line alternative, `virsh`, also provides a shell that can be used to create and manage virtual machines using a rich set of commands. The features provided by these tools include taking snapshots of virtual machines, creating virtual networks and live migration of virtual machines to another KVM host.

Once you have configured the required hardware setup:

- Install KVM on the target systems.
  See "Linux virtualization documentation" on page 165.
- Create and launch the required KVM virtual machines.
  See "Setting up a KVM guest" on page 50.

- Proceed to install the required SFHA product on the guest or host:
  See "Installing Storage Foundation and High Availability Solutions in the kernel-based virtual machine environment" on page 51.
  See "Installing and configuring Symantec Cluster Server in a kernel-based virtual machine (KVM) environment" on page 52.

# RHEL-based KVM installation and usage

You can list the available groups for virtualization from all yum repos with using the following `yum` command:

```
# yum grouplist|grep -i virtualization
```

This command lists the package group that has 'virtualization' as a substring in the group name among the list of all group names and does not install the virtualization RPM.

Subsequently, you can install the virtualization RPM with the following command:

```
# yum groupinstall "Virtualization"
```

# Setting up a KVM guest

The following is a high-level overview of the steps required for setting up KVM. For detailed instructions, refer to the applicable Linux documentation.

1. Before creating KVM guests, ensure that CPU and memory resources are available to create KVM guests on all nodes in the cluster.
2. Make sure that the required KVM packages are installed on the hosts.
3. Make sure that the service libvirtd is running on the hosts where KVM guests are to be created.
4. Create KVM guests. For network configuration, refer to the *Network configuration for VM-VM cluster* in Appendix A..
5. Install the operating system in the KVM guests.
6. Repeat the above steps for all KVM guests that you want to be a part of the cluster.
7. Install VCS on all the KVM guests. For information about installing VCS, refer to the *Symantec Cluster Server Installation Guide.*
8. Configure the VCS resources that you want VCS to manage. For more information, refer to the VCS documentation.

See "Network configuration for VM-VM cluster" on page 66.

# Installing Storage Foundation and High Availability Solutions in the kernel-based virtual machine environment

To set up a guest in a kernel-based virtual machine (KVM) environment with Storage Foundation and High Availability (SFHA) Solutions after installing KVM:

**Table 2-5** Tasks for installing SFHA Solutions in the KVM guest

| Task | Information |
|---|---|
| Set up the KVM host as needed. Create the KVM guests as needed. | See "Creating and launching a kernel-based virtual machine (KVM) host" on page 49.<br>See "Setting up a KVM guest" on page 50. |
| Install the SFHA Solutions product on the required KVM guest virtual machines. | For SFHA Solutions installation information, see the product installation guides.<br>See "Symantec Storage Foundation and High Availability Solutions product documentation" on page 164. |
| Configure the SFHA Solutions product on the required KVM guest virtual machines. | For SFHA Solutions configuration information, see the product installation guides.<br>See "Symantec Storage Foundation and High Availability Solutions product documentation" on page 164. |
| Configure resources as required for the KVM guest virtual machines. | See "About kernel-based virtual machine resources" on page 56. |

The tasks above apply to the following configurations:

- Dynamic Multi-Pathing in the guest
  See "Symantec Dynamic Multi-Pathing in the KVM guest virtualized machine" on page 35.
- Storage Foundation in the guest
  See "Symantec Storage Foundation in the virtualized guest machine" on page 37.
- Storage Foundation High Availability in the guest
- Storage Foundation Cluster File System High Availability in the guest

See "Symantec Dynamic Multi-Pathing in the KVM host and Symantec Storage Foundation HA in the KVM guest virtual machine" on page 40.

To set up a host in KVM environment with Storage Foundation and High Availability (SFHA) Solutions after installing KVM:

**Table 2-6** Tasks for installing SFHA Solutions in the KVM host

| Task | Information |
|---|---|
| Configure the KVM host. | See "Creating and launching a kernel-based virtual machine (KVM) host" on page 49. |
| Install the SFHA Solutions product on the KVM host. | For SFHA Solutions installation information, see the product installation guides.<br>See "Symantec Storage Foundation and High Availability Solutions product documentation" on page 164. |
| Configure the SFHA Solutions product on the required KVM hosts. | For SFHA Solutions configuration information, see the product installation guides.<br>See "Symantec Storage Foundation and High Availability Solutions product documentation" on page 164. |
| Create the KVM guests as needed. | See "Setting up a KVM guest" on page 50. |
| Configure resources as required for KVM guest virtual machines. | See "About kernel-based virtual machine resources" on page 56. |

The tasks above apply to the following configurations:

- Dynamic Multi-pathing in the host
  See "Symantec Dynamic Multi-Pathing in the KVM host" on page 36.
- Storage Foundation Cluster File System High Availability in the host
  See "Symantec Storage Foundation Cluster File System High Availability in the KVM host" on page 38.

# Installing and configuring Symantec Cluster Server in a kernel-based virtual machine (KVM) environment

To set up Symantec Cluster Server (VCS) in a KVM environment:

**Table 2-7** Tasks for installing VCS in a KVM environment

| Task | Information |
|---|---|
| Set up the KVM host as needed. Create the KVM guests as needed. | See "Creating and launching a kernel-based virtual machine (KVM) host" on page 49.<br>See "Creating and launching a kernel-based virtual machine (KVM) host" on page 49. |
| Install VCS. | For the: *Symantec Cluster Server Installation Guide*<br>See "Symantec Storage Foundation and High Availability Solutions product documentation" on page 164. |
| Configure VCS.<br>No additional VCS configuration is required to make it work inside the guest, provided the host as well as the network are configured. | For the: *Symantec Cluster Server Installation Guide*<br>See "Symantec Storage Foundation and High Availability Solutions product documentation" on page 164. |
| Configure network as required for KVM guest virtual machines. | See "Network configuration for VM-VM cluster" on page 66. |

The steps above apply for the following configurations:

- VCS in the KVM host
  See "Symantec Cluster Server in the KVM host" on page 42.
- VCS in the KVM guest
  See "Symantec Cluster Server in the guest" on page 43.
- VCS in the KVM host and ApplicationHA in the KVM guest virtual machine
  See "Symantec ApplicationHA in the guest and Symantec Cluster Server in the host" on page 44.
- VCS in a cluster across guests and physical machines
  See "Symantec Cluster Server in a cluster across virtual machine guests and physical machines" on page 45.

# How Symantec Cluster Server (VCS) manages Virtual Machine (VM) guests

High-level overview of how VCS manages VM guests.

- Physical machines form a cluster with VCS installed on them.

For information about installing VCS, see the *Symantec Cluster Server Installation Guide*.

- CPU and memory resources are made available to create VM guests on all nodes in the cluster.
- VCS is installed on all the hosts to manage the VM guest.
- The operating system is installed on the VM guest.

**Note:** The VM guest can be created on an image file or on a shared raw disk, provided the disk names are persistent across all the physical hosts.

- The VM guest is configured as a KVMGuest resource in VCS.

For detailed instructions on creating and configuring a VM guest, see the installation section in the Red Hat Enterprise Linux (RHEL) or SUSE Linux Enterprise Server (SLES) virtualization documentation.

To configure a VM guest for a physical machine to physical machine (PM-PM) configuration, the following conditions apply:

- You must configure a VM guest on one node with operating system installed on a shared storage accessible to all the VCS cluster nodes.
- Ensure that the image file resides on the shared storage so that the virtual machines can fail over across cluster nodes.
- You can configure the first VM guest using the standard installation procedure. See "Installing Storage Foundation and High Availability Solutions in the kernel-based virtual machine environment" on page 51.

Bundled agents are included with VCS for managing many applications. The KVMGuest agent is included and can be used to manage and provide high availability for KVM guests. For information on KVMGuest agent attributes, resource dependency and agent function, refer to the *Symantec Cluster Server Bundled Agents Reference Guide*.

# Installing and configuring ApplicationHA for application availability

To set up Symantec ApplicationHA in KVM environment:

**Table 2-8** Tasks for installing VCS in a KVM environment

| Task | Information |
|---|---|
| Set up the KVM host as needed. Create the KVM guests as needed. | See "Creating and launching a kernel-based virtual machine (KVM) host" on page 49.<br>See "Creating and launching a kernel-based virtual machine (KVM) host" on page 49. |
| Install ApplicationHA. | For the: *Symantec ApplicationHA Installation Guide*<br>See "Symantec Storage Foundation and High Availability Solutions product documentation" on page 164. |
| Configure ApplicationHA. | For the: *Symantec ApplicationHA Installation Guide*<br>See "Symantec Storage Foundation and High Availability Solutions product documentation" on page 164. |

The steps above apply for the following guest configurations:

- ApplicationHA in the KVM guest virtual machine
  See "Symantec ApplicationHA in the KVM virtualized guest machine" on page 41.
- VCS in the KVM host and ApplicationHA in the KVM guest virtual machine
  See "Symantec ApplicationHA in the guest and Symantec Cluster Server in the host" on page 44.

Chapter 3

# Configuring KVM resources

This chapter includes the following topics:



## About kernel-based virtual machine resources

After installing kernel-based virtual machine (KVM) and Storage Foundation and High Availability (SFHA) products and creating the virtual machines, you can configure your KVM resources to optimize your environment. Configuration processes vary depending on the SFHA solution you want to configure:

- If you are using Symantec Dynamic Multi-Pathing (DMP), Symantec Storage Foundation (SF), SFHA, or Symantec Storage Foundation Cluster File System High Availability (SFCFSHA) in your guests or hosts, you can optimize your storage for visibility and convenient management.
  See "Configuring storage" on page 56.
- If you are using Symantec Cluster Server (VCS), SFHA, or SFCFSHA in your guests or hosts, you can optimize your network to make your KVM resources highly available.
  See "Configuring networking" on page 63.

## Configuring storage

Symantec Storage Foundation and High Availability Solutions enable you to map and manage your storage more efficiently whether you have a guest or host solution.

# Consistent storage mapping in the KVM environment

Managing storage in the KVM environment requires consistent mapping. Storage which is presented to the guest either using the para-virtualized VirtIO drivers, or the fully virtualized IDE emulation, needs to be mapped from the host to the guest. Due to the volatile nature of the device naming used in Linux, care must be taken when mapping storage from the host to the guest. In Linux, the device names are based on enumeration order which can change when systems are rebooted.

Consistent mapping can be achieved by using:

- DMP meta-device
- Mapping devices using device ID
- Mapping devices using paths
- Mapping devices using volumes
- Linux `udev` device sym-links.

Avoid using disk labels when mapping storage to a guest. Disk labels can be modified by a guest and are not guaranteed.

In clustered environments, Active-Passive DMP devices cannot be mapped directly to a guest.

# Mapping devices to the guest

Non-persistent mappings can be made using `virsh attach-device`. The non-persistent mappings can be made persistent by redefining the KVM guests using `virsh dumpxml` *`domain`* followed by `virsh define` *`domain`*. Alternatively, persistent mappings can be created when a virtual machine is rebooted, these non-persistent mappings are lost. Persistent mappings can be created on the host using either `virt-manager` or by modifying the guests XML configuration using `virsh edit` *`domain`*.

The device links created in the directory `/dev/disk/by-path` should be consistent or if possible identical across all the physical hosts. Using different device links can cause issues with virtual machine live migration or VCS KVMGuest Agent failover operations.

See “Mapping devices using the virtio-scsi interface” on page 60.

## Mapping DMP meta-devices

Consistent mapping can be achieved from the host to the guest by using the Persistent Naming feature of DMP.

Running DMP in the host has other practical benefits:

- Multi-path device can be exported as a single device. This makes managing mapping easier, and helps alleviate the 32 device limit, imposed by the VirtIO driver.
- Path failover can be managed efficiently in the host, taking full advantage of the Event Source daemon to proactively monitor paths.
- When Symantec Storage Foundation and High Availability Solutions products are installed in the guest, the 'Persistent Naming' feature provides consistent naming of supported devices from the guest through the host to the array. The User Defined Names feature, or UDN, allows DMP virtual devices to have custom assigned names.

**To map a DMP meta-device to a guest**

1 Map the device to the guest. In this example the dmp device *xiv0_8614* is mapped to *guest_1*.

```
# virsh attach-disk guest_1 /dev/vx/dmp/xiv0_8614 vdb
```

2 The mapping can be made persistent by redefining the guest.

```
# virsh dumpxml guest_1 > /tmp/guest_1.xml
# virsh define /tmp/guest_1.xml
```

## Consistent naming across KVM Hosts

While enclosure based naming (EBN) provides persistent naming for a single node, it does not guarantee consistent naming across nodes in a cluster. The User Defined Names (UDN) feature of DMP allows DMP devices to be given both persistent and consistent names across multiple hosts. When using User Defined Names, a template file is created on a host, which maps the serial number of the enclosure and device to unique device name. User Defined Names can be manually selected, which can help make mappings easier to manage.

**To create consistent naming across hosts**

1 Create the User Defined Names template file.

```
# /etc/vx/bin/vxgetdmpnames enclosure=3pardata0 > /tmp/user_defined_names
# cat /tmp/user_defined_names
enclosure vendor=3PARdat product=VV serial=1628 name=3pardata0
 dmpnode serial=2AC00008065C name=3pardata0_1
 dmpnode serial=2AC00002065C name=3pardata0_2
 dmpnode serial=2AC00003065C name=3pardata0_3
 dmpnode serial=2AC00004065C name=3pardata0_4
```

2 If necessary, rename the devices. In this example, the DMP devices are named using the name of the guest they are to be mapped to.

```
# cat /dmp/user_defined_names
enclosure vendor=3PARdat product=VV serial=1628 name=3pardata0
         dmpnode serial=2AC00008065C name=guest1_1
         dmpnode serial=2AC00002065C name=guest1_2
         dmpnode serial=2AC00003065C name=guest2_1
         dmpnode serial=2AC00004065C name=guest2_2
```

3 Apply the User Defined Names file to this node, and all other hosts.

```
# vxddladm assign names file=/tmp/user_defined_names
```

4 Verify the user defined names have been applied.

```
# vxdmpadm getdmpnode enclosure=3pardata0
NAME         STATE     ENCLR-TYPE  PATHS  ENBL  DSBL  ENCLR-NAME
==========================================================================
guest_1_1   ENABLED   3PARDATA    2      2     0     3pardata0
guest_1_2   ENABLED   3PARDATA    2      2     0     3pardata0
guest_2_1   ENABLED   3PARDATA    2      2     0     3pardata0
guest_2_2   ENABLED   3PARDATA    2      2     0     3pardata0
```

## Mapping devices using paths

Mapping can be achieved using device ID: /dev/disk/by-path/

These links use the persistent properties of a path. For fibre channel devices, the sym-link name is composed of the bus identifier, the Worldwide Name (WWN) of the target, followed by the LUN identifier. A device will have an entry for each path to the device. In environments where multi-pathing is to be performed in the guest, make a mapping for each path for the device.

In the following example both paths to device *sdd* are mapped to *guest_3*.

**To map a path to a guest**

1 Identify the devices to map to the guest. Obtain the device IDs.

```
# udevadm info -q symlink --name sdd | cut -d\  -f 3
disk/by-id/scsi-200173800013420cd
```

In multi-path environments the device ID can be used to find all paths to the device.

```
# udevadm info --export-db |grep disk/by-id/scsi-200173800013420cd\  \
| cut -d\  -f 4
/dev/disk/by-path/pci-0000:0b:00.0-fc-0x5001738001340160:0x000000
/dev/disk/by-path/pci-0000:0c:00.0-fc-0x5001738001340161:0x000000
```

2 Map the device to the guest using the path using the device path.

```
# virsh attach-disk guest_3  \
/dev/disk/by-path/pci-0000:0b:00.0-fc-0x5001738001340160:0x000000 vdb
Disk attached successfully
# virsh attach-disk guest_3  \
/dev/disk/by-path/pci-0000:0c:00.0-fc-0x5001738001340161:0x000000 vdc
Disk attached successfully
```

3 Make the mapping persistent by re-defining the guest.

```
# virsh dumpxml guest_3 > /tmp/guest_3.xml
# virsh define /tmp/guest_3.xml
```

## Mapping devices using volumes

Mapping can be achieved by using Veritas Volume Manager volumes (VxVM volumes).

For more about mapping a VxVM volume to a guest:

See “Simplified management” on page 99.

## Mapping devices using the virtio-scsi interface

In Red Hat Enterprise Linux (RHEL) 6 Update 4, devices can be mapped to the guest through the virtio-scsi interface, replacing the virtio-blk device and providing the following improvements:

- The ability to connect to multiple storage devices

- A standard command set
- Standard device naming to simplify migrations
- Device pass-through

---

**Note:** Mapping using paths is also supported with the virtio-scsi interface.

---

To enable SCSI passthrough and use the exported disks as bare-metal SCSI devices inside the guest, the `<disk>` element's `device` attribute must be set to "lun" instead of "disk". The following disk XML file provides an example of the `device` attribute's value for virtio-scsi:

```
<disk type='block' device='lun' sgio='unfiltered'>
<driver name='qemu' type='raw' cache='none'/>
<source dev='/dev/disk/by-path/pci-0000:07:00.1-fc-0x5001438011393dee-lun-1'/
<target dev='sdd' bus='scsi'/>
<address type='drive' controller='4' bus='0' target='0' unit='0'/>
</disk>
```

**To map one or more devices using virtio-scsi**

1. Create one XML file for each SCSI controller, and enter the following content into the XML files:

   ```
   <controller type='scsi' model='virtio-scsi' index='1'/>
   ```

   The XML file in this example is named `ctlr.xml`.

2. Attach the SCSI controllers to the guest:

   ```
   # virsh attach-device guest1 ctlr.xml --config
   ```

3 Create XML files for the disks, and enter the following content into the XML files:

```
<disk type='block' device='lun' sgio='unfiltered'>
<driver name='qemu' type='raw' cache='none'/>
<source dev='/dev/disk/by-path/pci-0000:07:00.1-fc-0x5001438011393dee-lun
<target dev='sdd' bus='scsi'/>
<address type='drive' controller='1' bus='0' target='0' unit='0'/>
</disk>
```

The XML file in this example is named `disk.xml`.

4 Attach the disk to the existing guest:

```
# virsh attach-device guest1 disk.xml --config
```

# Resizing devices

Red Hat Linux Enterprise (RHEL) 6.3 and 6.4, and SUSE Linux Enterprise Server (SLES) 11 SP2 and SP3 do not support online disk re-sizing of VirtIO devices. To re-size a VirtIO device, the guest must be fully shut down and re-started.

You can use the following methods to resize the devices.

**To grow devices**

1 Grow the storage.

- If the storage device is a VxVM Volume, re-size the volume.
- If the storage device is a LUN from a storage array, re-size the device on the array.

2 Update the size of the disk device in the host.

- Stop all virtual machines using the storage device.
- If the device is a LUN from a storage array, update the size of the device:

  ```
  # blockdev --rereadpt device
  ```

- Restart the virtual machines.

3 Update the size of the storage device in the guest .

- If VxVM is managing the storage in the guest, use the `vxdisk resize` command.
- If VxVM is not managing the storage in the guest, see the appropriate documentation.

**To shrink devices**

1. Update the size of the disk device in the guest.
   - If VxVM is managing the device in the guest, if necessary, first use the vxresize utility to shrink any file systems and volumes which are using the device. Use the vxresize utility to update the size of the public region of the device:

     ```
     # vxdisk resize access_name length=new_size
     ```

   - If VxVM is not managing the storage in the guest, see the appropriate documentation.
2. Shrink the storage in the guest.
   - If the device is a VxVM volume, shrink the volume with the vxassist utility.
   - If the device is a LUN from a storage array, shrink the device on storage array.
3. Update the size of the disk device in the host.
   - Stop the guests which are using the devices.
   - If the device is a LUN from a storage array, use the following command:

     ```
     # blockdev --rereadpt device
     ```

4. Start the guests.

# Configuring networking

You must configure a network for the host and KVM guest to enable Symantec Storage Foundation and High Availability Solutions to provide:

- Application failover
- Virtual machine availability

## Bridge network configuration

The bridge network configuration can be performed in two parts:

- Configuring host network
- Configuring guest network

## Host network configuration

The libvirtd service creates a default bridge virbr0 which is a NAT'ed private network. It allocates private IPs from the network 192.168.122.0, to the guests using virbr0 for networking. If the guests are required to communicate on the public network of the host machines, then a bridge must be configured. This bridge can be created using the following steps:

1. Create a new interface file with the name `ifcfg-br0` in `/etc/sysconfig/network-scripts/` location where all the other interface configuration files are present. Its contents are as follows:

   ```
   DEVICE=br0
   Type=Bridge
   BOOTPROTO=dhcp
   ONBOOT=yes
   ```

2. Add the physical interface to the bridge using the following command.

   ```
   # brctl addif eth0 br0
   ```

   This adds the physical interface that the guests shares with the br0 bridge created in the previous step.

3. Verify that your eth0 was added to the br0 bridge using the `brctl show` command.

   ```
   # brctl show
   ```

   The output must look similar to the following:

   ```
   bridge name     bridge id               STP enabled     interfaces
   virbr0          8000.000000000000       yes
   br0             8000.0019b97ec863       yes             eth0
   ```

4. The eth0 network configuration must be changed. The ifcfg-eth0 script is already present.

5. Edit the file and add a line **BRIDGE=br0**, so that the contents of the configuration file look like the following example:

   ```
   DEVICE=eth0
   BRIDGE=br0
   BOOTPROTO=none
   HWADDR=00:19:b9:7e:c8:63
   ONBOOT=yes
   TYPE=Ethernet
   ```

```
USERCTL=no
IPV6INIT=no
PEERDNS=yes
NM_CONTROLLED=no
```

6. Restart the network services to bring all the network configuration changes into effect.

### Configuring guest network

Refer to the virtualiztion-related Linux documentation for instructions on configuring guest network.

## Network configuration for VCS cluster across physical machines (PM-PM)

The network configuration and storage of the hosts is similar to the VCS cluster configurations. For configuration-related information, refer to the *Symantec Cluster Server Installation Guide*. However, you must set up a private link and a shared storage between the physical hosts on which the VM guests are configured.

**Figure 3-1**

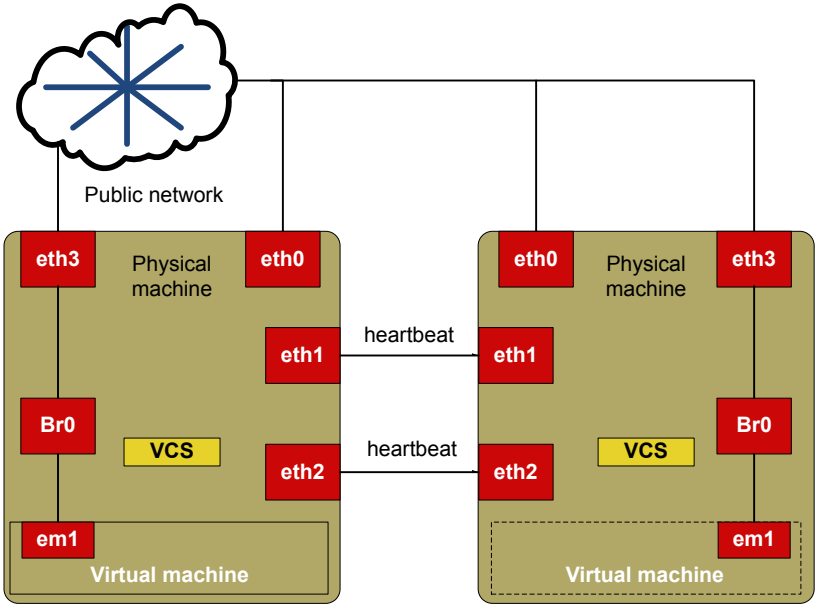

# Standard bridge configuration

The standard bridge configuration is a generic network configuration for bridge networking.

**Figure 3-2** Standard bridge configuration

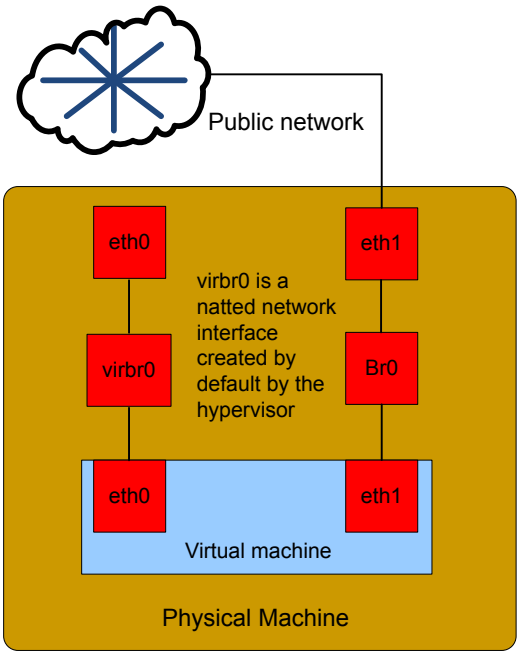


Standard bridge

# Network configuration for VM-VM cluster

To configure the VCS cluster between the virtual machines, you must configure the network and storage for the cluster. The setup details for network and storage configurations are explained in the subsequent sections. Figure 3-3 shows a cluster setup between two VM guests running on two different hosts.

**Figure 3-3** Network configuration for VM- VM cluster

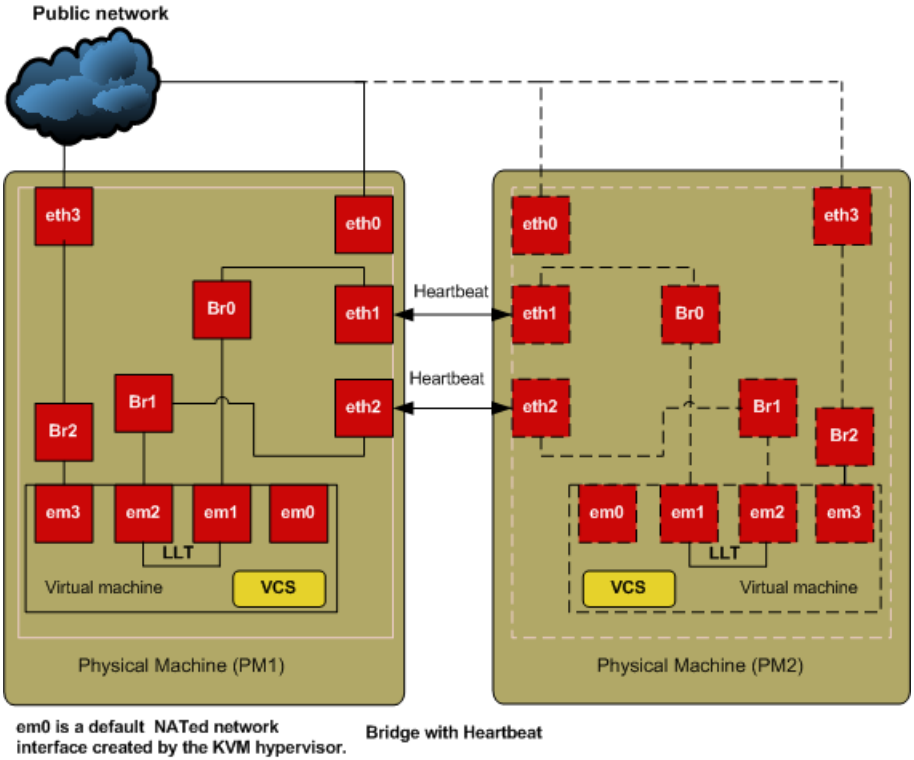


See "Bridge network configuration" on page 63.

Section 3

# Implementing a RedHat Enterprise Virtualization environment

Chapter 4

# Getting started with Red Hat Enterprise Virtualization (RHEV)

This chapter includes the following topics:



## Symantec Cluster Server configuration options for the Red Hat Enterprise Virtualization (RHEV) environment

Symantec Cluster Server (VCS) provides virtual machine monitoring and failover to another host in the Red Hat Enterprise Virtualization (RHEV) environment. VCS enables the following for RHEV hosts:

- Connects multiple, independent systems into a management framework for increased availability.
- Enables nodes to cooperate at the software level to form a cluster.
- Links commodity hardware with intelligent software to provide application failover and control.

- Enables other nodes to take predefined actions when a monitored application fails, for instance to take over and bring up applications elsewhere in the cluster.

VCS supports the following configurations:

**Table 4-1** VCS supported configuration options in the RHEV environment

| Objective | Recommended VCS configuration |
|---|---|
| Virtual machine monitoring and failover for hosts | VCS in the hosts<br>See "Symantec Cluster Server in the KVM host" on page 42. |
| Disaster recovery in virtualized environment | VCS on the Red Hat Enterprise Linux (RHEL) hypervisor |
| Application failover for guest virtual machines | VCS in the guest virtual machines<br>See "Symantec Cluster Server in the guest" on page 43. |
| Application failover across guest virtual machines and physical hosts | VCS in guest virtual machines and physical host machines<br>See "Symantec Cluster Server in a cluster across virtual machine guests and physical machines" on page 45. |

**Note:** Virtual machine disaster recovery is supported in the RHEV environment only

**Note:** Symantec ApplicationHA is supported in the RHEL KVM environment only.

**Figure 4-1** Symantec Cluster Server in the RHEV host

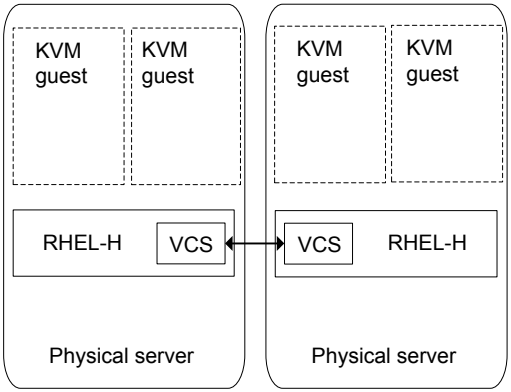

**Figure 4-2** Symantec Cluster Server in the RHEV guest

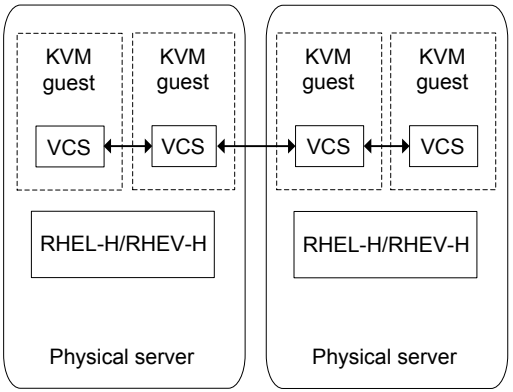


For more information on Symantec Cluster Server features, see the *Symantec Cluster Server Administrator's Guide*.

# About setting up Red Hat Enterprise Virtualization (RHEV) with Symantec Cluster Server

Before setting up RHEV, verify your planned configuration will meet the system requirements, licensing and other considerations for installation with Symantec Cluster Server.

- Licensing: Customers running Symantec Cluster Server in a Linux virtualization environment (KVM and RHEV) are entitled to use an unlimited number of guests on each licensed server or CPU.
- Red Hat system requirements: For the latest on Red Hat Enterprise Linux (RHEL) and RHEV requirements, see Red Hat documentation.
- Symantec product requirements:
  See Table 4-2 on page 72.
- *Release Notes*: Each Symantec product contains last minute news and important details for each product, including updates to system requirements and supported software. Review the Release Notes for the latest information before you start installing the product.
  The product documentation is available on the Web at the following location:
  https://sort.symantec.com/documents

**Table 4-2** Symantec product requirements

| | |
|---|---|
| Software | ■ Symantec Cluster Server 6.1<br>Used for virtual machine monitoring and failover |

For the latest information on supported hardware, visit the following URL:

http://www.symantec.com/docs/TECH211575

**Table 4-3** VCS system requirements for RHEV-supported configurations

| | |
|---|---|
| VCS version | 6.1 |
| Red Hat Enterprise Virtualization | 3.1, 3.2 |
| Supported OS version in host | Red Hat Enterprise Linux 6, Update 3, Update 4 |
| Supported OS in VM guest | RHEL 5 and RHEL 6 |
| Hardware requirement | Full virtualization-enabled CPU |

**Note:** Symantec Storage Foundation is not supported in the Red Hat Enterprise Virtualization (RHEV) environment.

### Limitations in the Red Hat Enterprise Virtualization (RHEV) environment

For more information on VCS limitations and known issues, refer to VCS 6.1 Release Notes for Linux.

For RHEV related limitations, refer to the Red Hat Enterprise Virtualization release notes.

See "Linux virtualization documentation" on page 165.

# Setting up a virtual machine in the Red Hat Enterprise Virtualization (RHEV) environment

Following is a high-level overview of the steps required for setting up virtual machines in Red Hat Enterprise Virtualization (RHEV) environment. For detailed instructions, see the Red Hat Enterprise Virtualization documentation.

**To set up virtual machines in the RHEV environment.**

1. Before creating virtual machines, ensure that CPU and memory resources are available to create virtual machines on all nodes in the cluster.
2. Make sure that the required RHEV RPMs are installed on the hosts.
3. Make sure that the Virtual Desktop Server Manager (VDSM) service is running on the hosts where virtual machines are to be created. Before you create a virtual machine on a host, make sure that the state of the host in RHEV-M is up.
4. Create virtual machines.

   See "Linux virtualization documentation" on page 165.
5. Install the operating system in the virtual machines.

See "Network configuration for VM-VM cluster" on page 66.

Chapter 5

# Configuring VCS to manage virtual machines

This chapter includes the following topics:



## Installing and configuring Symantec Cluster Server for virtual machine and application availability

To set up Symantec Cluster Server (VCS) in Red Hat Enterprise Virtualization (RHEV) environment:

- Install VCS.
- Configure VCS.

### How Symantec Cluster Server (VCS) manages virtual machines

Following is a high-level overview of how VCS manages virtual machines in the Red Hat Enterprise Virtualization (RHEV) environment:

- Physical machines form a cluster with VCS installed on them.
  See the *Symantec Cluster Server Installation Guide* for installation information.

- CPU and memory resources are made available to host virtual machines on all nodes in the cluster.
- VCS is installed on all the hosts to manage the virtual machines.
- The operating system is installed on the virtual machine on any one host.
- The virtual machine is configured as a KVMGuest resource in VCS.

# About the KVMGuest agent

The KVMGuest agent enables Symantec Cluster Server (VCS) to monitor a KVM guest - that is, a virtual machine in the KVM environment or the Red Hat Enterprise Virtualization (RHEV) environment. The agent performs tasks such as bringing virtual machines online and taking them offline. The KVMGuest agent operates in both KVM and RHEV environments. This topic describes its behavior in the RHEV environment.

For details on the KVMGuest agent behavior in open source KVM environment or RHEV environment, see the *Symantec Cluster Server Bundled Agents Reference Guide*.

The KVMGuest agent uses `virsh` commands to manage virtual machines in the KVM environment and Representational State Transfer (REST) APIs to manage virtual machines in RHEV environment by using the REST APIs to determine the state of the virtual machine. The agent determines the resource state, on the basis of the virtual machine state. REST design architecture focuses on resources and their representations for some specific service. REST APIs help software developers and administrators integrate the functionality of the RHEV environment with custom scripts or with external applications which access the API by means of HTTP.

Prerequisites for administering virtual machines in a RHEV environment by using REST APIs:

- A networked installation of Red Hat Enterprise Virtualization Manager, which includes the REST API
- A client or programming library that initiates and receives HTTP requests from the REST API

The following table lists various states of a virtual machine in RHEV environment and the corresponding VCS resource state:

**Table 5-1**

| Virtual machine state | VCS resource state | Resource confidence level |
|---|---|---|
| wait_for_launch | ONLINE | 10 |

**Table 5-1** *(continued)*

| Virtual machine state | VCS resource state | Resource confidence level |
|---|---|---|
| powering_up | ONLINE | 60 |
| up | ONLINE | 100 |
| powering_down | ONLINE | 40 |
| paused | ONLINE | 20 |
| down | OFFLINE | – |
| saving_state | INTENTIONAL OFFLINE | – |
| suspended | INTENTIONAL OFFLINE | – |
| restoring_state | ONLINE | 50 |
| migrating | INTENTIONAL OFFLINE | – |
| reboot_in_progress | INTENTIONAL OFFLINE | – |
| image_locked | UNKNOWN | – |
| unknown | UNKNOWN | – |

**Table 5-2** KVMGuest agent functions

| Function | Tasks |
|---|---|
| Online | KVM environment: Agent uses the `virsh start` command to start the guest virtual machine. When the resource is configured to define the guest configuration, agent uses the `virsh define` command to define the virtual machine while bringing it online.<br><br>RHEV environment: Agent uses the REST APIs to start the virtual machine. If the DROpts attribute is set to configure the guest network, the agent also sets the payload as a cdrom. This payload contains networking parameters to be set within the guest after a DR failover.<br><br>The agent waits for a certain time period after initiating the virtual machine start. You can specify this wait period by using the “DelayAfterGuestOnline” attribute.<br><br>The agent also checks whether the virtual machine is configured for disaster recovery by checking the `DROpts` attribute. If this attribute is set correctly, the agent inserts a virtual CDROM into the virtual machine's configuration. This CDROM contains a file that contains the site-specific network parameters to be applied at this site for the virtual machine. When the virtual machine boots, the `vcs-net-reconfig` service installed inside the guest checks for the CDROM and the disaster recovery information. If the `vcs-net-reconfig` service finds the disaster recovery information, the service applies the networking parameters to the virtual machine. |
| Offline | The Offline function initiates a graceful shutdown of the virtual machine.<br><br>KVM environment: Agent uses the `virsh shutdown` command to shutdown the guest virtual machine. If the SyncDir attribute is configured to synchronize the guest virtual machine configuration file, then the configuration file is copied to the location configured as a SyncDir attribute.<br><br>RHEV environment: Agent uses the REST APIs to shutdown the virtual machine.<br><br>The agents wait for a certain time period after initiating the shutdown for the virtual machine to shut down completely. You can specify this wait period by using the DelayAfterGuestOffline attribute. |

**Table 5-2** KVMGuest agent functions *(continued)*

| Function | Tasks |
|---|---|
| Monitor | KVM environment: Agent uses the `virsh domstate` command to determine the status of the guest virtual machine.<br>RHEV environment: Agent uses the REST APIs to get the status of the virtual machine. |
| Clean | KVM environment: Agent uses the `virsh destroy` command to forcefully shutdown the guest virtual machine. If the SyncDir attribute is configured to synchronize the guest virtual machine configuration file, then the configuration file is copied to the location configured as a SyncDir attribute.<br>RHEV environment: Agent uses REST APIs to stop the virtual machine. |
| Migrate | KVM environment: The agent uses the `virsh migrate` command to start virtual machine migration.<br>RHEV environment: The agent uses REST APIs to start virtual machine migration. Additionally, it checks whether the virtual machine migration is allowed or not.<br>**Note:** When a virtual machine is configured for disaster recovery, the virtual machine cannot be migrated across sites. |

The KVMGuest agent recognizes the following resource states:

**Table 5-3**

| Resource state | Description |
|---|---|
| ONLINE | Indicates that the guest virtual machine is running. |
| OFFLINE | Indicates that the guest virtual machine has stopped. |
| FAULTED | Indicates that the guest virtual machine has failed to start or has unexpectedly stopped. |
| UNKNOWN | Indicates that a problem exists with the configuration or with the ability to monitor the resource. |
| INTENTIONAL OFFLINE | Indicates that the virtual machine has either migrated to another physical host or the administrator intentionally suspended it. |

The Symantec Cluster Server agent for managing virtual machines in a KVM or RHEV environment, is represented by the KVMGuest resource type definition:

```
type KVMGuest (
        static int IntentionalOffline = 1
        static boolean AEPTimeout = 1
        static int MigrateTimeout = 300
        static int MigrateWaitLimit = 2
        static keylist SupportedActions = { "guestmigrated", "vmconfigsync", "DevScan" }
        static keylist SupportedOperations = { "migrate"}
        static keylist RegList = { "GuestName", "DelayAfterGuestOnline", "DelayAfterGuestOffline
        static str ArgList[] = { GuestName, DelayAfterGuestOnline, DelayAfterGuestOffline, SyncD
        str CEInfo{} = { Enabled=0, CESystem=NONE, FaultOnHBLoss=1 }
        str RHEVMInfo{} = { Enabled=0, URL=NONE, User=NONE, Password=NONE, Cluster=NONE, UseManu
        str GuestName
        int DelayAfterGuestOnline = 5
        int DelayAfterGuestOffline = 30
        str SyncDir
        str GuestConfigFilePath
        boolean ResyncVMCfg = 0
        str DROpts{} = { ConfigureNetwork=0, IPAddress=NONE, Netmask=NONE, Gateway=NONE, DNSServ
)
```

The `RHEVMInfo` attribute enables the KVMGuest attribute configuration to support the Red Hat Enterprise Virtualization environment. RHEVMInfo specifies the following information about the RHEV environment:

| Attribute value | Description |
|---|---|
| Enabled | Specifies whether the virtualization environment is a KVM environment or a Red Hat Enterprise Virtualization (RHEV) environment.<br>0 indicates the KVM environment.<br>1 indicates the RHEV environment.<br>The default value is 0. |
| URL | Specifies the RHEV-M URL, that the KVMGuest agent can use for REST API communication. The API can only communicate with the secure port (SSL). For example:<br>`https://rhevm-server.example.com:443` |

| Attribute value | Description |
|---|---|
| User | Specifies the RHEV-M user name that the agent must use for REST API communication. For example:<br>`admin@internal`<br>`rhevadmin@example.com` |
| Password | Specifies the encrypted password associated with the RHEVM user profile. The password should be encrypted using “vcsencrypt” command.<br>To generate the encrypted password, run the following command:<br>`# /opt/VRTSvcs/bin/vcsencrypt -agent plain_text_password` |
| Cluster | Specifies the name of the RHEV-M cluster of which the VCS host is a member. |
| UseManualRHEVMFencing | Specifies if the use of manual RHEV-M fencing is enabled in the event that the physical host on which virtual machine is running crashes.<br>0 indicates that manual RHEV-M fencing is disabled.<br>1 indicates that manual RHEV-M fencing is enabled.<br>The default value is 0. |

The `DROpts` attribute enables the virtual machine for disaster recovery. The attribute contains site-specific network parameters for the virtual machine. The value of this attribute consists of the following keys that define the disaster recovery options for the virtual machine:

| Attribute keys | Description |
|---|---|
| DNSSearchPath | The domain search path used by the virtual machine in this site. The value of this key must contain a list of DNS domain names that are used for the DNS lookup of a hostname in case the domain name of the hostname is not specified. Use spaces to separate the domain names. |
| DNSServers | The list of DNS servers used by the virtual machine in this site. The value of this key must contain a list of IP addresses of DNS servers that are used for the DNS lookup of a hostname. Use spaces to separate the IP addresses. |
| Gateway | The default gateway used by the virtual machine in this site. |

| Attribute keys | Description |
|---|---|
| Device | The Network Interface Card (NIC) that is dedicated to the exclusive IP address of the virtual machine in this site. If this key is not specified, the agent automatically selects the first dedicated NIC for the assignment of the IP address, if specified. Example: eth0. |
| IPAddress | The IP address to be assigned to the virtual machine in this site after a cross-site failover. |
| Netmask | The netmask to be used by the virtual machine in this site after a cross-site failover. |
| ConfigureNetwork | The `DROpts` attribute value is applied to the virtual machine only if this key is set to 1. Type and dimension: string-association. |

**Note:** For information on other attributes associated with the KVMGuest agent, see the *Symantec Cluster Server Bundled Agents Reference Guide*.

# Validating the virtualization environment

The KVMGuest agent validates the virtualization environment with the help of a standalone utility `havirtverify`.

The agent invokes this utility in `open` entry point and `attr_changed` entry point. The utility validates the configured virtualization environment for a resource based on its configuration.

For RHEV, the utility:

- Validates the configured URL and user credentials.
- Verifies whether RHEV HA for a configured virtual machine is disabled or not.
- Verifies the `DROpts` attribute

For KVM, the utility checks whether `libvirtd` is running or not.

Once the validation is passed, the agent can start monitoring the resource. If validation fails for a particular resource, its state is reported as UNKNOWN. This validation is also triggered if value of either of the following attributes changes:`RHEVMInfo`, `GuestName`.

You can also run this utility manually for verifying the environment.

**To validate the RHEV environment**

- Run:

  ```
  # /opt/VRTSvcs/bin/KVMGuest/havirtverify resource_name
  ```

  If validation passes, the following message displays:

  ```
  #/opt/VRTSvcs/bin/KVMGuest/havirtverify resource_name
  Red Hat Enterprise Virtualization Environment validation successfully
  completed for resource resource_name
  ```

  If validation fails, the following message displays:

  ```
  # /opt/VRTSvcs/bin/KVMGuest/havirtverify resource_name
  Virtualization environment validation failed for resource resource_name
  ```

All the log messages of this utility are sent to the engine log file.

# Configuring a resource in a RHEV environment

Before you configure a resource in a RHEV environment, you must:

- Ensure that RHEV-HA is disabled for the virtual machine which you want to configure monitoring with Symantec Cluster Server (VCS).
- Configure the virtual machine to run on a specific host and the virtual machine image must be available to all the hosts in the VCS cluster.
- Configure the firewall settings to allow REST API communication.

**To configure a KVMGuest resource**

1. Validate the virtualization environment.

   See "Validating the virtualization environment" on page 81.

2. Specify the name of the virtual machine that VCS must manage, as the value of the GuestName attribute.
3. Configure the DelayAfterGuestOnline and DelayAfterGuestOffline attributes.

   **Note:** The default value of DelayAfterGuestOnline is 5 and DelayAfterGuestOffline is 30.

4 Validate the RHEV-M URL, valid RHEV-M user (name), and password.

5 To configure the RHEVMInfo attribute, specify the appropriate value of each key. The following table lists each key and its related instruction:

| Key | Instruction |
|---|---|
| Enabled | Set the value to 1. |
| URL | Specify the RHEV-M URL. |
| User | Specify a valid user name.<br>For example: admin@internal<br>rhevadmin@example.com |
| Password | Specify the encrypted password associated with RHEV-M User profile.<br>To generate the encrypted password, run the following command: # **/opt/VRTSvcs/bin/vcsencrypt -agent *plain_text_password*** |
| Cluster | Specify the RHEV-M cluster name. |
| UseManualRHEVMFencing | Enable the use of manual RHEV-M fencing in the event that the physical host on which virtual machine is running crashes.<br>For example: # **UseManualRHEVMFencing=1** |

# Configuring multiple KVMGuest resources

If a VCS service group has more than one KVMGuest resource monitoring virtual machines and one of the virtual machines is migrated to another host, then a service group level concurrency violation occurs as the service group state goes into PARTIAL state on multiple nodes.

Symantec recommends configuring only one KVMGuest resource in a Service group. See the sample configurations below for reference.

Configuration 1:

```
group rhev_grp1 (

SystemList = { sys1 = 0, sys2 = 1 }
```

```
)

KVMGuest kvmres1 (

RHEVMInfo = { Enabled = 1,

URL = "https://rhevm-server.example.com:443",

User = "admin@internal"

Password = bncNfnOnkNphChdHe,

Cluster = dc2_cluster1,

UseManualRHEVMFencing=1 }

GuestName = rhevvm1

DelayAfterGuestOnline = 20

DelayAfterGuestOffline = 35

)
```

Configuration 2:

```
group rhev_grp1 (

SystemList = { sys1 = 0, sys2 = 1 }
)

KVMGuest kvmres1 (

RHEVMInfo = { Enabled = 1,

URL = "https://rhevm-server.example.com:443",

User = "admin@internal"

Password = bncNfnOnkNphChdHe,

Cluster = dc2_cluster1,

UseManualRHEVMFencing=0 }
```

```
GuestName = rhevvm1

DelayAfterGuestOnline = 20

DelayAfterGuestOffline = 35

)

group rhev_grp2 (

SystemList = { sys1 = 0, sys2 = 1 }
)

KVMGuest kvmres2 (

RHEVMInfo = { Enabled = 1,

URL = "https://rhevm-server.example.com:443",

User = "admin@internal"

Password = bncNfnOnkNphChdHe,

Cluster = dc2_cluster1,

UseManualRHEVMFencing=0 }

GuestName = rhevvm2

DelayAfterGuestOnline = 20

DelayAfterGuestOffline = 35

)
```

Section 4

# Implementing Linux virtualization use cases

Chapter 6

# Application visibility and device discovery

This chapter includes the following topics:



## About storage to application visibility using Veritas Operations Manager

Datacenters adopt virtualization technology to effectively use the IT-infrastructure and substantially reduce the capital and operational expenditures. If you have adopted virtualization technology in your datacenter, Veritas Operations Manager provides you an efficient way of discovering and managing your virtual storage and infrastructure assets.

In your datacenter, Veritas Operations Manager helps you view the following relationships:

- Applications in your datacenter that Veritas Operations Manager manages and the virtual hosts on which they are running.
- Physical storage in your datacenter that is exported to the virtual machines.

Veritas Operations Manager supports the following virtualization technologies:

- VMware
- Microsoft Hyper-V
- Kernel-based Virtual Machine (KVM)
- RedHat Enterprise Virtualization (RHEV)

In the VMware virtualization technology, a designated Control Host discovers the VMware vCenter Server in the datacenter. This discovery displays those ESXi servers that VMware vCenter Server manages, and the virtual machines that are configured on the ESXi servers.

For more information, see the *Storage Foundation and High Availability Solutions Virtualization Guide for Linux on ESXi*

For Microsoft Hyper-V, Veritas Operations Manager discovers Hyper-V virtual machines and their correlation with the Hyper-V server. It also discovers the storage that is provisioned to the guests, and its correlation with the virtual machine and Hyper-V server. The Hyper-V guest (with or without `VRTSsfmh` RPM), when added to Veritas Operations Manager Management Server domain, provides storage mapping discovery.

For Kernel-based Virtual Machine (KVM), Veritas Operations Manager discovers KVM virtual machines on the Linux host if the KVM modules are installed, and configured. Veritas Operations Manager discovers basic information about only running virtual machines. For example, virtual machine name, CPU, and so on.

See "About Microsoft Hyper-V virtualization discovery" on page 90.

For more information, see the Veritas Operations Manager documentation.

# About Kernel-based Virtual Machine (KVM) virtualization discovery in Veritas Operations Manager

Kernel-based Virtual Machine (KVM) is a full virtualization solution for Linux on x86 hardware containing virtualization extensions (Intel VT or AMD-V). Veritas Operations Manager discovers KVM virtual machines on the Linux host if the KVM modules are installed, and configured. Veritas Operations Manager discovers basic information about only running virtual machines. For example, virtual machine name, CPU, and so on. Veritas Operations Manager uses `virsh` commands to discover KVM-related information.

Kernel-based Virtual Machine (KVM) discovery pre-requisites are as follows:

- `VRTSsfmh` package must be present on the Linux host.
- KVM modules must be installed and configured.

Kernel-based Virtual Machine (KVM) discovery limitations are as follows:

- Veritas Operations Manager discovers only running virtual machines.
- Exported storage discovery, and storage correlation is not supported.

# About Microsoft Hyper-V virtualization discovery

Hyper-V is a hypervisor-based virtualization technology from Microsoft for x86-64 systems. You can use Veritas Operations Manager to discover Hyper-V host and virtual machine-related information if the Hyper-V role is enabled on the managed host. Veritas Operations Manager uses the Hyper-V WMI API for the discovery.

Hyper-V discovery can be grouped into the following categories:

- Virtual machine discovery: Hyper-V virtual machine discovery by Veritas Operations Manager and its correlation with the Hyper-V server.
- Exported storage discovery: Discovery of storage that is provisioned to the guests and its correlation with the virtual machine and Hyper-V server.

See “Virtual machine discovery in Microsoft Hyper-V ” on page 90.

See “Storage mapping discovery in Microsoft Hyper-V” on page 91.

# Virtual machine discovery in Microsoft Hyper-V

Veritas Operations Manager lets you discover information about Hyper-V virtual machines. For example, the name of the virtual machine, allocated memory, CPU, state, and the storage exported (virtual hard disks and pass through disks) from Hyper-V server to Hyper-V guest. Veritas Operations Manager discovers all virtual machines including the virtual machines without the guest operating system installed.

Agent and agentless discoveries of Hyper-V virtual machines are supported. However, for the agentless method, the discovered information is limited. To discover more information about the configured virtual machines, the agent discovery method should be used. It provides detailed information about the virtual machines.

For more information on agent and agentless discovery, see the *Veritas Operations Manager Management Server Administrator's Guide*

Virtual machine discovery prerequisites are as follows:

- The `VRTSsfmh` package should be installed on the Hyper-V server (parent partition).
- The Hyper-V role should be enabled.
- The Windows Management Instrumentation (WMI) service should be running.

A limitation of virtual machine discovery is listed below:

- Hyper-V discovery is not supported on an agentless Hyper-V Server (parent partition) to which the Hyper-V virtual machines are associated.

See “About Microsoft Hyper-V virtualization discovery” on page 90.

# Storage mapping discovery in Microsoft Hyper-V

Veritas Operations Manager discovers the storage provisioned to the guests from the host's local storage, or storage area network (SAN). The Hyper-V guest (with or without `VRTSsfmh` package), when added to the Veritas Operations Manager Management Server domain, provides storage mapping discovery.

Additional storage attributes are also displayed on the page. For example, size, type of the storage (VHD or passthrough disk), and the storage container (volume on the host where virtual storage is provisioned). The storage device handles on the guest will be mapped to the corresponding VHD or passthrough disk provisioned from host. Veritas Operations Manager also discovers the snapshot disks provisioned to the VMS.

The storage mapping discovery prerequisites are as follows:

- The Hyper-V server must be running Microsoft Windows 2008 R2 or later operating system.
- Windows Management Instrumentation (WMI) should be running on the guest.

The storage mapping discovery limitation is as follows:

- Storage correlation is not supported for Linux guests.

For more information on storage mapping and storage correlation, see the *Veritas Operations Manager Management Server Administrator's Guide*.

See “About Microsoft Hyper-V virtualization discovery” on page 90.

Chapter 7

# Server consolidation

This chapter includes the following topics:



## Server consolidation

Storage Foundation and High Availability Solutions products can be used in many combinations. The configurations listed are the minimum required to accomplish the objectives of the respective use cases.

Server consolidation enables you to run multiple virtual machines, each with the functionality equivalent to a physical server, combining the multiple applications and their workloads onto a single server for better server utilization and reduced datacenter server sprawl.

**Figure 7-1** Server consolidation

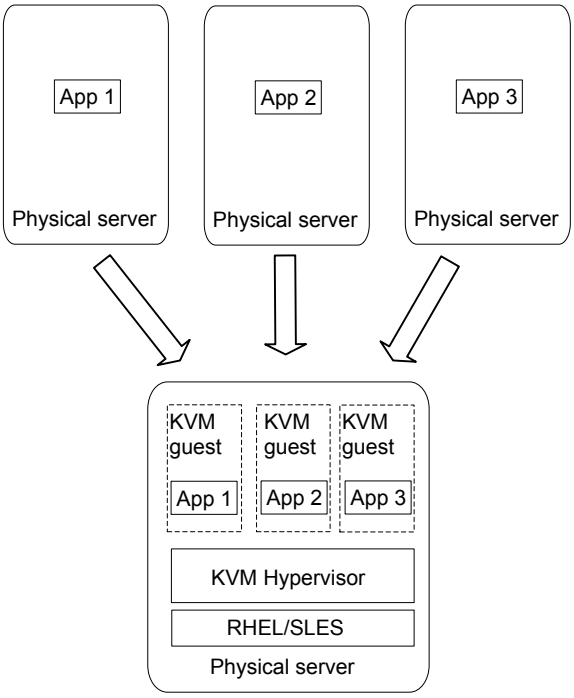


The server consolidation use case is supported for the following Linux virtualization technologies:

- Red Hat Enterprise Linux (RHEL)
- SUSE Linux Enterprise Server (SLES)

# Implementing server consolidation for a simple workload

This solution for a single server with Storage Foundation HA illustrates the migration of a single workload into a KVM Guest.

**Figure 7-2** Server consolidation for a simple workload

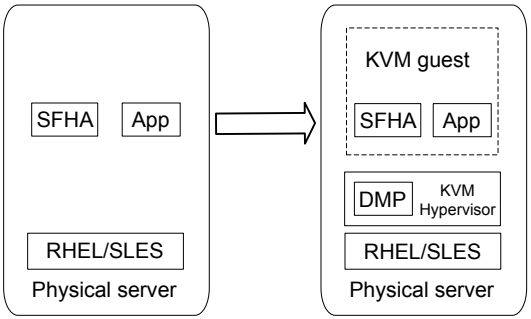


**To implement server consolidation for a simple workload**

1. Install SFHA in the virtual machine.

   See “Installing Storage Foundation and High Availability Solutions in the kernel-based virtual machine environment” on page 51.

2. Map the storage from the array to the host.
3. Map the storage from the array to the guest.

   See “Mapping devices to the guest” on page 57.

4. Go into the guest and make sure you can import disk groups.

Chapter 8

# Physical to virtual migration

This chapter includes the following topics:



## Physical to virtual migration

Migrating data from physical servers to virtual machines can be painful. Symantec Storage Foundation and High Availability Solutions products can make painful migrations of data from physical to virtual environments easier and safer to execute.

With Symantec Storage Foundation and High Availability Solutions, there is no need to copy any data from source to destination, but rather the administrator reassigns the same storage or a copy of the storage for a test migration, to the virtual environment. Data migration with Storage Foundation (SF), Storage Foundation HA (SFHA), or Storage Foundation Cluster File System High Availability (SFCFSHA) can be executed in a central location, migrating all storage from an array utilized by Storage Foundation managed hosts.

Physical to virtual migration (P2V) requires migrating data from a physical server to a virtualized guest. The LUNs are first physically connected to the host, and then the LUNs are mapped in KVM from the host to the guest.

Without SF, SFHA, or SFCFSHA in the host, you must identify which storage devices with mapping to the guest. Putting SF, SFHA, or SFCFSHA in the host enables quick and reliable identification of storage devices to be mapped. If you are running DMP in the host, you can map the DMP devices directly. Symantec Storage Foundation and High Availability Solutions products add manageability and ease of use to an otherwise tedious and time-consuming process.

The physical to virtual migration use case is supported for the following Linux virtualization technologies:

- Red Hat Enterprise Linux (RHEL) KVM
- SUSE Linux Enterprise Server (SLES) KVM

# How to implement physical to virtual migration (P2V)

Migrating data from a physical server to a virtualized guest, the LUNs are first physically connected to the host, and then the LUNs are mapped in KVM from the host to the guest.

This use case procedure is very similar to the server consolidation use case and the procedures are quite similar. Physical to virtual migration is the process used to achieve server consolidation.

This use case requires Symantec Storage Foundation HA or Symantec Storage Foundation Cluster File System HA in the KVM host and Symantec Storage Foundation in the KVM guest. For setup information:

See "Installing Storage Foundation and High Availability Solutions in the kernel-based virtual machine environment" on page 51.

There are two options:

- If SFHA Solutions products are installed on both the physical server and the virtual host, identifying the LUNs which need mapping is made easy. Once the LUNs are connected to the virtual host, 'vxdisk –o alldgs list' can be used to identify the devices in the disk group which require mapping.
- If Symantec Storage Foundation and High Availability Solutions (SFHA Solutions) products are not installed on the virtual host and the physical server is a Linux system, the devices which need mapping can be identified by using the device IDs on the physical server.

**To implement physical to virtual migration with Storage Foundation in the host and guest**

1 Find the Linux device IDs of the devices which need mapping.

```
# vxdg list diskgroup
```

2 For each disk in the disk group:

```
# vxdmpadm getsubpaths dmpnodename=device
# ls -al /dev/disk/by-id/* | grep subpath
```

If Storage Foundation is not installed on the host, before decommissioning the physical server, identify the LUNs which require mapping by using the devices serial

numbers. The LUNs can be mapped to the guest using the persistent "by-path" device links.

**To implement physical to virtual migration if Storage Foundation is not installed in the host**

1. On the physical server, identify the LUNs which must be mapped on the KVM host using the `udevadm` command.
2. Map the LUNs to the virtualization host.

   The udev database can be used to identify the devices on the host which need to be mapped.

   ```
   # udevadm info --export-db | grep '/dev/disk/by-path' | \
                   cut -d' ' -f4

   /dev/disk/by-path/pci-0000:05:00.0-fc-0x5006016239a01884-lun-1
   /dev/disk/by-path/pci-0000:05:00.0-fc-0x5006016239a01884-lun-2
   ```

   Map the LUNs to the guest. As there are multiple paths in this example, the paths sym-link can be used to ensure consistent device mapping for all four paths.

   ```
   # virsh attach-disk guest1 \
   /dev/disk/by-path/pci-0000:05:00.0-fc-0x5006016239a01884-lun-1 \
         vdb
   # virsh attach-disk guest1 \
   /dev/disk/by-path/pci-0000:05:00.0-fc-0x5006016239a01884-lun-2 \
         vdc
   ```

3. Verify that the devices are correctly mapped to the guest. The configuration changes can be made persistent by redefining the guest.

   ```
   # virsh dumpxml guest1 > /tmp/guest1.xml
   # virsh define /tmp/guest1.xm
   ```

**To implement physical to virtual migration with Storage Foundation in the guest and host**

1. Map the LUNs to the virtualization host.
2. On the virtualization host, identify the devices which require mapping. For example, the devices with the disk group *data_dg* are mapped to *guest1*.

   ```
   # vxdisk -o alldgs list |grep data_dg
   3pardata0_1  auto:cdsdisk    -            (data_dg)    online
   3pardata0_2  auto:cdsdisk    -            (data_dg)    online
   ```

3 Map the devices to the guest.

```
# virsh attach-disk guest1 /dev/vx/dmp/3pardata0_1 vdb
Disk attached successfully

# virsh attach-disk guest1 /dev/vx/dmp/3pardata0_2 vdc
Disk attached successfully
```

4 In the guest, verify that all devices are correctly mapped and that the disk group is available.

```
# vxdisk scandisks
# vxdisk -o alldgs list |grep data_dg
3pardata0_1  auto:cdsdisk    -            (data_dg)    online
3pardata0_2  auto:cdsdisk    -            (data_dg)    online
```

5 In the virtualization host make the mapping persistent by redefining the guest:

```
# virsh dumpxml guest1 > /tmp/guest1.xml
# virsh define /tmp/guest1.xml
```

**To use a Veritas Volume Manager volume as a boot device when configuring a new virtual machine**

1 Follow the recommended steps in your Linux virtualization documentation to install and boot a VM guest.

When requested to select managed or existing storage for the boot device, use the full path to the VxVM storage volume block device, for example */dev/vx/dsk/boot_dg/bootdisk-vol*.

2 If using the `virt-install` utility, enter the full path to the VxVM volume block device with the --disk parameter, for example, *--disk path=/dev/vx/dsk/boot_dg/bootdisk-vol*.

Chapter 9

# Simplified management

This chapter includes the following topics:



## Simplified management

Independent of how an operating system is hosted, consistent storage management tools save an administrator time and reduce the complexity of the environment. Symantec Storage Foundation and High Availability Solutions products in the guest provide the same command set, storage namespace, and environment as in a non-virtual environment.

This use case requires Symantec Storage Foundation HA or Symantec Storage Foundation Cluster File System HA in the KVM host. For setup information:

See "Installing Storage Foundation and High Availability Solutions in the kernel-based virtual machine environment" on page 51.

The simplified management use case is supported for the following Linux virtualization technologies:

- Red Hat Enterprise Linux (RHEL) KVM
- SUSE Linux Enterprise Server (SLES) KVM

## Provisioning storage for a guest virtual machine

A volume can be provisioned within a VM guest as a data disk or a boot disk.

- Data disk: provides the advantage of mirroring data across arrays.

- Boot disk: provides the ability to migrate across arrays.

Adding a VxVM storage volume as a data disk to a running guest virtual machine can be done in the following ways:

- Using the `virt-manager` console.
- Using the `virsh` command line.

## Provisioning Veritas Volume Manager volumes as data disks for VM guests

The following procedure uses Veritas Volume Manager (VxVM) volumes as data disks (virtual disks) for VM guests. The example host is *sys1* and the VM guest is *guest1*. The prompts in each step show in which domain to run the command.

**To provision Veritas Volume Manager volumes as data disks**

1. Create a VxVM disk group (*mydatadg* in this example) with some disks allocated to it:

```
sys1# vxdg init mydatadg TagmaStore-USP0_29 TagmaStore-USP0_30
```

2. Create a VxVM volume of the desired layout (in this example, creating a simple volume):

```
sys1# vxassist -g mydatadg make datavol1 500m
```

3. Map the volume *datavol1* to the VM guest:

```
sys1# virsh attach-disk guest1 /dev/vx/dsk/mydatadg/datavol1 vdb
```

4. To make the mapping persistent, redefine the VM guest.

```
sys1# virsh dumpxml guest1 > /tmp/guest1.xml

sys1# virsh define /tmp/guest1.xml
```

## Provisioning Veritas Volume Manager volumes as boot disks for guest virtual machines

The following procedure outlines how to provision a Veritas Volume Manager (VxVM) volume as a boot disk for guest virtual machines.

The example host is *sys1* the VM guest is *guest1*. The prompts in each step show in which domain to run the command.

**To provision Veritas Volume Manager volumes as boot disks for guest virtual machines**

1. On the host, create a VxVM volume. Use the size that is recommended by your Linux documentation. In this example, a 16GB volume is created:

```
sys1# vxassist -g boot_dg make bootdisk-vol 16g
```

2. Follow the recommended steps in your Linux documentation to install and boot a VM guest, and use the virtual disk as the boot disk.

# Boot image management

With the ever-growing application workload needs of datacenters comes the requirement to dynamically create virtual environments. This creates a need for the ability to provision and customize virtual machines on-the-fly. Every virtual machine created needs to be provisioned with a CPU, memory, network and I/O resources.

As the number of guest virtual machines increase on the physical host, it becomes increasingly important to have an automatic, space-optimizing provisioning mechanism. Space-savings can be achieved as all the guest virtual machines can be installed with the same operating system, i.e., boot volume. Hence, rather than allocate a full boot volume for each guest, it is sufficient to create single boot volume and use space-optimized snapshots of that “Golden Boot Volume” as boot images for other virtual machines.

The primary I/O resource needed is a boot image, which is an operating system environment that consists of: the following

- A bootable virtual disk with the guest operating system installed
- A bootable, a guest file system
- A custom or generic software stack

For boot image management, Storage Foundation and High Availability (SFHA) Solutions products enable you to manage and instantly deploy virtual machines based on templates and snapshot-based boot images (snapshots may be full or space optimized). For effective boot image management in KVM based virtual environments, deploy the SFHA Solutions products in the combined host and guest configuration.

Benefits of boot image management:

- Eliminates the installation, configuration and maintenance costs associated with installing the operating system and complex stacks of software

- Infrastructure cost savings due to increased efficiency and reduced operational costs.
- Reduced storage space costs due to shared master or gold image as well as space-optimized boot images for the various virtual machines
- Enables high availability of individual guest machines with Symantec Cluster Server (running on the host) monitoring the VM guests and their boot images
- Ability to create and deploy virtual machines across any remote node in the cluster

## Creating the boot disk group

Once Storage Foundation HA (SFHA) is installed on the Linux server using the combined host and VM guest configuration, the next step is to create a disk-group in which the Golden Boot Volume and all the various space-optimized snapshots (VM boot images) will reside. For a single-node environment, the disk-group is local or private to the host. For a clustered environment (recommended for live migration of VMs), Symantec recommends creating a shared disk-group so that the Golden Boot Volume can be shared across multiple physical nodes.

It is possible to monitor the disk-group containing the Guest VM boot image(s) and the guest VMs themselves under VCS so that they can be monitored for any faults. However it must be kept in mind that since the boot images are in the same disk-group, a fault in any one of the disks backing the snapshot volumes containing the boot disks can cause all the guest VMs housed on this node to failover to another physical server in the Storage Foundation Cluster File System High Availability (SFCFS HA) cluster. To increase the fault tolerance for this disk-group, mirror all volumes across multiple enclosures making the volumes redundant and less susceptible to disk errors.

**To create a shared boot disk group**

1. Create a disk group, for example *boot_dg*.

```
$ vxdg -s init boot_dg  device_name_1
```

2. Repeat to add multiple devices.

```
$ vxdg -g boot_dg adddisk device_name_2
```

# Creating and configuring the golden image

The basic idea is to create a point-in-time image based on a master or gold image. The image will serve as the basis for all boot images once it is set up. Hence, first set up a complete virtual machine boot volume as a golden boot volume.

**To create the golden image**

1. In the selected disk group, create a VxVM volume. Use the size that is recommended by your Linux documentation. For example, the disk group is *boot_dg*, the golden boot volume is *gold-boot-disk-vol*, the volume size is 16GB.

   ```
   sys1# vxassist -g boot_dg make gold-boot-disk-vol 16g
   ```

2. Follow the recommended steps in your Linux documentation to install and boot a VM guest.

   When requested to select managed or existing storage for the boot device, use the full path to the VxVM storage volume block device.

   For example: */dev/vx/dsk/boot_dg/gold-boot-disk-vol*.

3. If using the `virt-install` utility, enter the full path to the VxVM volume block device with the --disk parameter.

   For example: *--disk path=/dev/vx/dsk/boot_dg/gold-boot-disk-vol*.

4. After the virtual machine is created, install any guest operating system with the boot volume and the virtual machine configured exactly as required.

5. After the virtual machine is created and configured, shut it down.

You can now use the boot image as a image (hence called a golden image) for provisioning additional virtual machines that are based on snapshots of the Golden Boot Volume. These snapshots can be full copies (mirror images) or they can be space-optimized snapshots. Using space-optimized snapshots greatly reduces the storage required to host the boot disks of identical multiple virtual machines. Note that since both, the full and space-optimized snapshots, are instantly available (no need to wait for the disk copy operation), provisioning of new virtual machines can now be instantaneous as well.

# Rapid Provisioning of virtual machines using the golden image

As mentioned above, for rapid provisioning of new virtual machines based on the golden image, we need to have full or space-optimized snapshots of the Golden Boot Volume. These snapshots can then be used as boot images for the new virtual machines. The process to create these snapshots is outlined below in the procedures below.

Creating Instant, Full Snapshots of Golden Boot Volume for Rapid Virtual Machine Provisioning

**To create instant, full snapshots of the golden boot volume for rapid virtual machine provisioning**

1. Prepare the volume for an instant full snapshot. In the example, the disk group is *boot_dg* and the golden boot volume is *gold-boot-disk-vol*.

   ```
   $ vxsnap -g boot_dg prepare gold-boot-disk-vol
   ```

2. Create a new volume which will be used as the boot volume for the new provisioned guest. The size of the guests boot volume must match the size of the golden boot volume.

   ```
   $ vxassist -g boot_dg make guest1-boot-disk-vol 16g layout=mirror
   ```

3. Prepare the new boot volume so it can be used as a snapshot volume.

   ```
   $ vxsnap -g boot_dg prepare guest1-boot-disk-vol
   ```

4. Create the full instant snapshot of the golden boot volume.

   ```
   $ vxsnap -g boot_dg make source=gold-boot-disk-vol/snapvol=\
     guest1-boot-disk-vol/syncing=off
   ```

5. Create a new virtual machine, using the snapshot *guest1-boot-disk-vol* as an "existing disk image."

**To create instant, space-optimized snapshots of the golden boot volume for rapid virtual machine provisioning**

1. Prepare the volume for an instant snapshot. In the example, the disk group is*boot_dg* and the golden boot volume is *gold-boot-disk-vol*.

   ```
   $ vxsnap -g boot_dg prepare gold-boot-disk-vol
   ```

2. Use the `vxassist` command to create the volume that is to be used for the cache volume. The cache volume will be used to store writes made to the space-optimized instant snapshots.

   ```
   $ vxassist -g boot_dg make cache_vol 5g  layout=mirror init=active
   ```

3. Use the `vxmake cache` command to create a cache object on top of the cache volume which you created in the previous step.

   ```
   $ vxmake -g boot_dg cache cache_obj cachevolname=cache_vol autogrow=on
   ```

4 Start the cache object:

```
$ vxcache -g boot_dg start cache_obj
```

5 Create a space-optimized instant snapshot of the golden boot image:

```
$ vxsnap -g boot_dg make source=\
gold-boot-disk-vol/newvol=guest1-boot-disk-vol/cache=cache_obj
```

6 Create a new virtual machine, using the snapshot of the golden image as an existing disk image.

## Storage Savings from space-optimized snapshots

With the large number of virtual machines housed per physical server, the number of boot images used on a single server is also significant. A single bare-metal Linux boot image needs around 3 GB of space at a minimum. Installing software stacks and application binaries on top of that requires additional space typically resulting in using around 6 GB of space for each virtual machine that houses a database application.

When a user provisions a new virtual machine, the boot image can be a full copy or a space-optimized snapshot. Using a full copy results in highly inefficient use of storage. Not only is storage consumed to house identical boot images, storage is also consumed in making the boot images highly available (mirror across enclosures) as well in their backup.This large amount of highly available, high performance storage is very expensive, and likely to eliminate the cost advantages that server virtualization would otherwise provide. To add to it, backup and recovery of such capacity is also an expensive task.

In order to address the above issue, Symantec recommends the use of space-optimized snapshots of the gold image as boot images of the various VM guests. Space-optimized snapshots do not make a full copy of the data in the gold image, rather they work on the copy-on-write principle where only the changed blocks are stored locally. This set of changed blocks is called a Cache Object and it is stored in a repository for all such space-optimized snapshots, called the Cache Object Store, which is backed by physical storage. The Cache Object offers a significant storage space reduction, typically occupying a 5-20% storage footprint, relative to the parent volume (the gold image volume in this case). The same Cache Object Store can be used to store changed blocks for multiple snapshot volumes.

Each Snapshot held in the Cache Object Store contains only changes made to the gold image to support that installation's boot environment. Hence, to achieve the best possible storage reduction, install software on data disks rather than root file

systems and limit as many changes as possible to the gold image operating files (i.e., system, hosts, passwd, etc.).

Chapter 10

# Application monitoring using Symantec ApplicationHA

This chapter includes the following topics:



## About application monitoring using Symantec ApplicationHA

You can use Symantec Cluster Server (VCS) or another Storage Foundation and High Availability (SFHA) Solutions product such as Storage Foundation Cluster File System High Availability (SFCFSHA) to monitor the health of the applications, as well as provide clustering and failover capabilities. However, you may not need the full feature set of VCS, SFHA, or SFCFSHA simply to monitor an application in a virtual machine. In a virtual environment, the size and resource requirements for an application are a serious consideration.

Symantec provides a lightweight, simple, application availability solution for virtualized workloads that combines seamlessly with clustering and disaster recovery solutions from Symantec and other vendors. You can use Symantec ApplicationHA. Symantec ApplicationHA provides an easy GUI- and wizard-based method for

installing and configuring Symantec High Availability products and administering application monitoring on virtual machines.

For lightweight, easy-to-use application monitoring capability, use Symantec ApplicationHA in the kernel-based virtual machine (KVM) guest.

# What is Symantec ApplicationHA

Symantec ApplicationHA provides monitoring capabilities for applications running inside guest virtual machines in the KVM virtualization environment. ApplicationHA adds a layer of application awareness to the core high availability (HA) functionality offered by Symantec™ Cluster Server (VCS) in the physical host.

ApplicationHA is based on VCS, and uses similar concepts such as agents, resources, and service groups. However, ApplicationHA has a lightweight server footprint that allows faster installation and configuration in virtualization environments.

Key benefits include the following:

- Out of the box integration with VCS
- Full visibility and control over applications, including the ability to start, stop, and monitor applications running inside virtual machines
- High availability of the application as well as the virtual machine on which the application runs
- Graded application fault-management responses such as:
  - Application restart
  - ApplicationHA-initiated, internal reboot, or soft reboot of a virtual system
  - VCS-initiated, external reboot or hard reboot of a virtual system
  - Failover of the virtual system to another VCS node
- Standardized way to manage applications using a single interface that is integrated with the Veritas Operations Manager (VOM) console
- Specialized Application Maintenance mode, in which ApplicationHA lets you intentionally take an application out of its purview for maintenance or troubleshooting

## How ApplicationHA is deployed in the KVM environment

Kernel-based Virtual Machine (KVM) is a full virtualization solution for Linux on AMD64 and Intel64 hardware. KVM lets you create and manage multiple virtual machines on a single physical host.

ApplicationHA provides high availability of applications running on virtual machines. Symantec Cluster Server (VCS) provides high availability of the virtual machines that run on a physical host.

The following figure illustrates how ApplicationHA and VCS are deployed in a typical KVM virtualization environment.

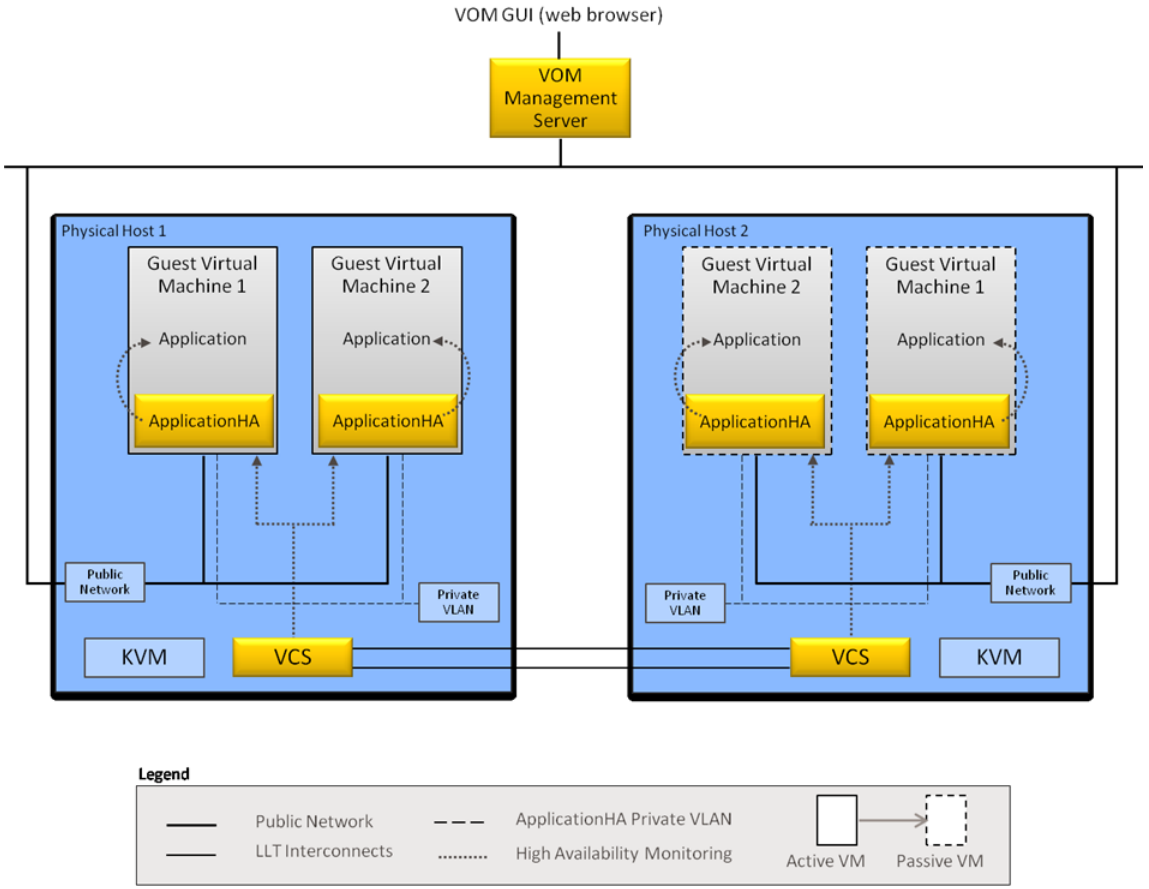


ApplicationHA is installed on the virtual machine, and provides high availability to a configured application running on the virtual machine. VCS is installed on the physical host, as part of a Storage Foundation Cluster File System High Availability (SFCFSHA) stack installation. VCS provides high availability to the virtual machine where the configured application runs.

You must enable VCS to support ApplicationHA to ensure application-aware monitoring of virtual machines.

For more information, see the *Symantec ApplicationHA User's Guide*.

When you enable VCS to support ApplicationHA, a private VLAN is created between monitored virtual machines and the VCS node (physical host). The private VLAN facilitates heartbeat communication between VCS in the physical host and ApplicationHA in the virtual machines.

Veritas Operations Manager (VOM) provides you with a centralized management console (GUI) to administer application monitoring with ApplicationHA.

# Symantec ApplicationHA agents

Agents are application-specific modules that plug into the ApplicationHA framework that manages applications and resources of predefined resource types configured for applications and components on a system. The agents are installed when you install Symantec ApplicationHA guest components. These agents start, stop, and monitor the resources configured for the applications and report state changes. If an application or its components fail, ApplicationHA restarts the application and its resources on a virtual system.

Symantec ApplicationHA agents are classified as follows:

- Infrastructure agents
  Agents such as NIC, IP, and Mount are classified as infrastructure agents. Infrastructure agents are automatically installed as part of an ApplicationHA installation on KVM guests.
  For more details about the infrastructure agents, refer to the operating system-specific *Symantec Cluster Server Bundled Agents Reference Guide*.
- Application agents
  Application agents are used to monitor third party applications such as Oracle. These agents are packaged separately and are available in the form of an agent pack that gets installed when you install Symantec ApplicationHA guest components.
  The ApplicationHA agent pack is released on a quarterly basis. The agent pack includes support for new applications, as well as fixes and enhancements to existing agents. You can install the agent pack on an existing ApplicationHA guest components installation.
  Refer to the Symantec Operations Readiness Tools (SORT) Web site for information on the latest agent pack availability.
  https://sort.symantec.com/agents
  Refer to the agent-specific configuration guide for more details about the application agents.

# Getting started with ApplicationHA

There are two sets of steps that you can use to get started with ApplicationHA. To monitor high availability of an application running on a virtual machine:

See "Ensuring high availability of applications" on page 111.

To monitor the high availability of the application as well as the virtualization infrastructure on which the virtual machine runs:

See "Ensuring high availability of virtualization infrastructure" on page 111.

## Ensuring high availability of applications

You can ensure high availability of applications running inside virtual machines by using ApplicationHA. To provide high availability to the applications, perform the following steps:

- Install Veritas Operations Manager Add-on for ApplicationHA Management on the VOM Management Server.
- Install ApplicationHA on the virtual machine.
- Add the virtual machine as a managed host to Veritas Operations Manager (VOM).
- Configure application monitoring on the virtual machine.

The following figure illustrates the workflow for ensuring high availability of applications with ApplicationHA. The figure also indicates the corresponding document that you must refer for detailed instructions at each step.

## Ensuring high availability of virtualization infrastructure

In addition to high availability of applications using ApplicationHA you can also ensure high availability of the virtualization infrastructure with VCS. By using VCS, you can externally restart virtual machines and fail over the virtual machines in case of application failures or virtual machine failures. To ensure high availability of the virtualization environment, perform the following steps:

- Install Veritas Operations Manager Add-on for ApplicationHA Management on the VOM Management Server.
- Install SFCFSHA on the physical host.
- Enable ApplicationHA capabilities in underlying VCS in the virtual machine.
- Install ApplicationHA on the virtual machine.

- Add virtual machine and physical host as managed hosts to Veritas Operations Manager (VOM).
- Configure application monitoring on the virtual machine.

The following figure illustrates the workflow for ensuring high availability of the applications running inside the virtual machine and the virtualization infrastructure. The figure also indicates the corresponding documents that you must refer for detailed instructions at each step.

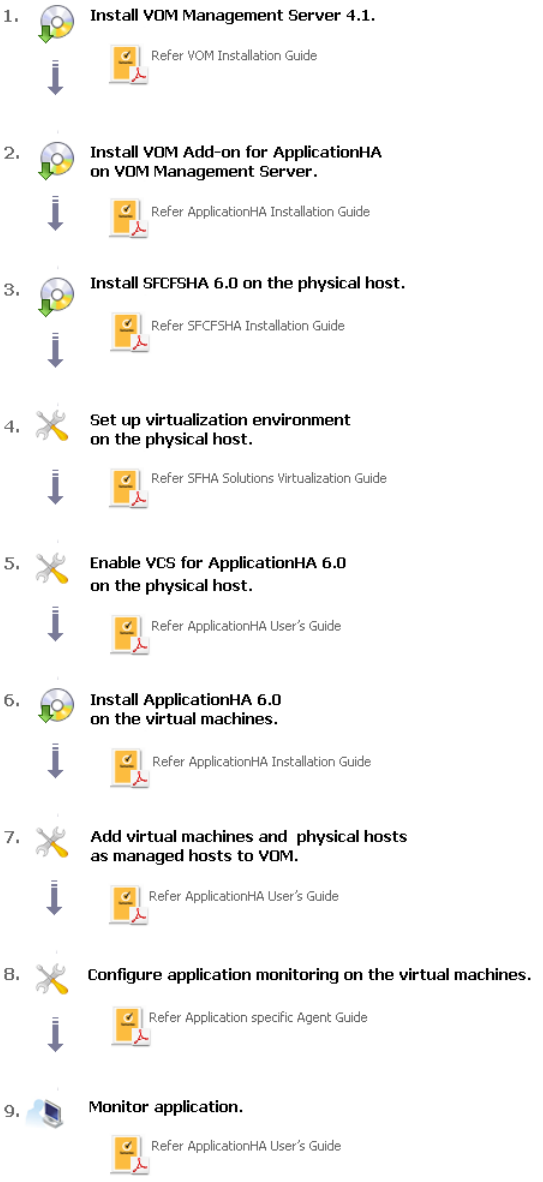
1.
Install VOM Management Server 4.1.
Refer VOM Installation Guide
2.
Install VOM Add-on for ApplicationHA
on VOM Management Server.
Refer ApplicationHA Installation Guide
3.
Install SFCFSHA 6.0 on the physical host.
Refer SFCFSHA Installation Guide
4.
Set up virtualization environment
on the physical host.
Refer SFHA Solutions Virtualization Guide
5.
Enable VCS for ApplicationHA 6.0
on the physical host.
Refer ApplicationHA User's Guide
6.
Install ApplicationHA 6.0
on the virtual machines.
Refer ApplicationHA Installation Guide
7.
Add virtual machines and physical hosts
as managed hosts to VOM.
Refer ApplicationHA User's Guide
8.
Configure application monitoring on the virtual machines.
Refer Application specific Agent Guide
9.
Monitor application.
Refer ApplicationHA User's Guide

Chapter 11

# Application availability using Symantec Cluster Server

This chapter includes the following topics:



## About application availability options

Symantec products can provide the ultimate levels of availability in your KVM environment. In a KVM environment, you can choose a different combination of Symantec High Availability solutions: ApplicationHA and Symantec Cluster Server (VCS).

ApplicationHA by itself provides application monitoring and restart capabilities while providing ultimate visibility and manageability through Veritas Operations Manager. When ApplicationHA is adopted together with Symantec Cluster Server in the host, the two solutions work together to ensure that the applications are monitored and restarted if needed, and virtual machines are restarted if application restarts are not effective. These two solutions work together to provide the ultimate level of availability in your KVM environment.

If your KVM environment requires the same level of application availability provided by a VCS cluster in a physical environment, you can choose to adopt Symantec Cluster Server in the virtual machines. In this configuration, your application enjoys fast failover capability in a VCS cluster in the virtual machines.

**Table 11-1** Comparison of availability options

| Required availability level | Recommended solution | Supported virtualization option |
|---|---|---|
| Application monitoring and restart | ApplicationHA in the virtual machines | Red Hat Enterprise Linux (RHEL) KVM |
| Virtual machine monitoring and restart | VCS cluster in the host monitoring the virtual machines as a resource | Red Hat Enterprise Linux (RHEL) KVM<br>Red Hat Enterprise Virtualization (RHEV)<br>SUSE Linux Enterprise Server (SLES) KVM |
| Combined application and virtual machine availability | ApplicationHA in the virtual machine and VCS cluster in the host | Red Hat Enterprise Linux (RHEL) KVM |
| Application failover to standby node in cluster | VCS cluster in the virtual machines | Red Hat Enterprise Linux (RHEL) KVM<br>SUSE Linux Enterprise Server (SLES) KVM<br>Red Hat Enterprise Virtualization (RHEV)<br>Microsoft Hyper-V<br>Oracle Virtual Machine (OVM) |

**Note:** For application high availability and failover capabilities the application data must be on the shared storage accessible to all the nodes of the VCS cluster.

For setup information for ApplicationHA or VCS:

See "Installing and configuring Symantec Cluster Server in a kernel-based virtual machine (KVM) environment" on page 52.

**Note:** You can also use the cluster functionality of Symantec Storage Foundation HA or Symantec Storage Foundation Cluster File System HA if you need storage management capabilities in addition to application availability for your KVM environment.

# Symantec Cluster Server In a KVM Environment Architecture Summary

VCS in host architecture

- Manages multiple guest virtual machines as a single unit of control
- Provides automatic restart or fail-over of individual guest virtual machines in response to failures
- Provides Start / Stop / Monitor of individual guest virtual machines from a common console across the entire server pool using Veritas Operations Manager (VOM)

VCS in guest architecture

- Manages applications running in the guest virtual machine as a single unit of control
- Provides automatic restart or fail-over of individual applications to other guest virtual machine or physical machine.
- Provides Start / Stop / Monitor of individual applications from a common console across appropriate guest virtual machines in the cluster using Veritas Operations Manager (VOM)

# VCS in host to provide the Virtual Machine high availability and ApplicationHA in guest to provide application high availability

VCS running in the host monitors the virtual machine to provide the VM high availability. ApplicationHA running in the virtual machine (VM guest ensures the application high availability by monitoring the configured application. VCS and ApplicationHA can be combined together to provide the enhanced solution for achieving application and VM high availability.

VCS in host provides the primary VM monitoring. It can start/stop the virtual machine and fail-over it to another node in case of any fault. We then run ApplicationHA within the guest that monitors the application running inside the guest virtual machine. ApplicationHA in guest will not trigger an application fail-over in case of application fault, but it'll try to restart the application on same VM guest. If ApplicationHA fails to start the application, it can notify the VCS running in the host to take corrective action which includes virtual machine restart or virtual machine fail-over to another host.

For detailed information about ApplicationHA and integration of ApplicationHA with VCS, see the *Symantec ApplicationHA User's Guide*.

**Figure 11-1** VCS In host for VM HA and ApplicationHA in guest for application HA

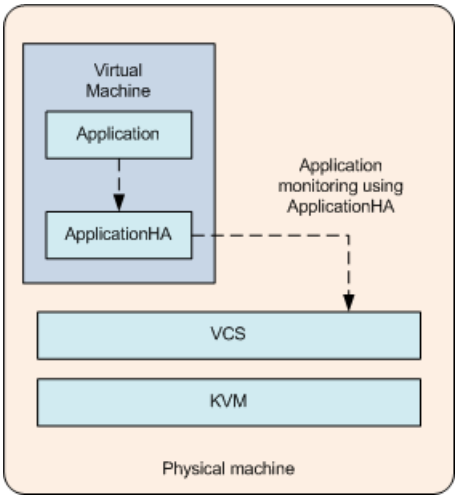


# Virtual to Virtual clustering and failover

Running VCS in multiple guest virtual machines enables guest-to-guest clustering. VCS can then monitor individual applications running within the guest and then fail over the application to another guest in the virtual – virtual cluster.

**Figure 11-2** Clustering between guests for application high availability

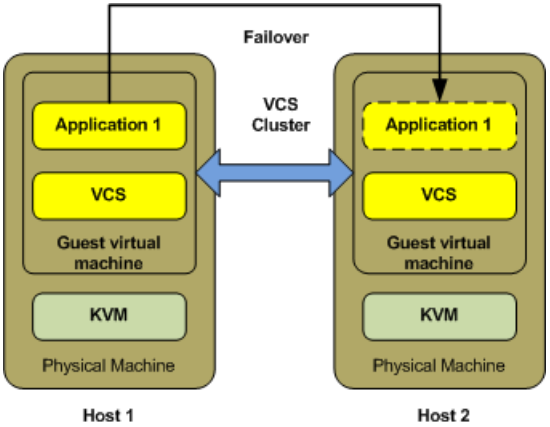

**Note:** I/O fencing support for clustering between guests for application high availability: Non-SCSI3, coordination point (CP) server based fencing is supported. SCSI3 fencing is not supported.

You can run VCS within each guest machine to provide high availability to applications running within the guest.

A VCS cluster is formed among the VM guests in this configuration. The VM guests in the cluster can be either on the same physical host or on different physical hosts. VCS is installed in the VM guests in the cluster. The VCS installation and configuration in a virtual machine is similar to that of VCS in the physical host clusters. This VCS cluster manages and controls the applications and services that run inside the VM guests. Any faulted application or service is failed over to other VM guest in the cluster. This configuration does not take care of the VM guest fail-overs since VCS runs inside the VM guest.

**Figure 11-3** VCS cluster across VM guests on the same physical machine

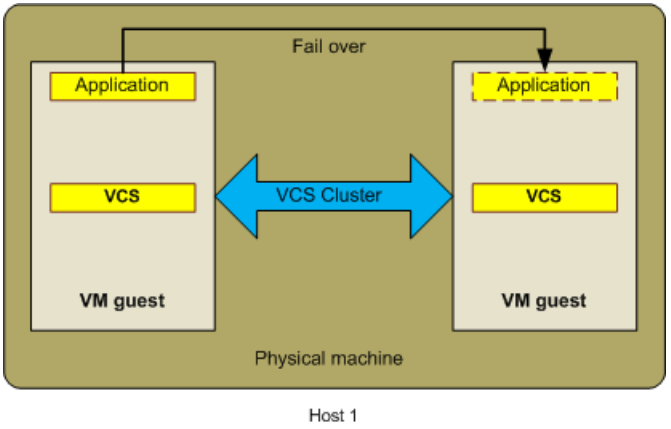


**Note:** I/O fencing support for a VCS cluster across VM guests on the same physical machine: Non-SCSI3, CP server based fencing is supported. SCSI3 fencing is not supported.

# Virtual to Physical clustering and failover

One can also create a physical to virtual cluster by combining VCS inside the virtual machine together with VCS running on any other physical host. This virtual-physical cluster enables VCS to monitor applications running within the guest and then fail

over the application to another host. The reverse flow is also true, thus enabling the fail-over of an application running on a physical host into a VM guest machine.

A VCS cluster is formed among the VM guests and physical machines. VCS is installed on the VM guests and on different physical machines in the cluster. VM guests are connected to physical machines through the network of their VM hosts. In this case, the VM host is a physical machine on which one or more VM guests forming the cluster are hosted.

This VCS cluster manages and monitors the services and applications running on cluster nodes that can either be VM guests or physical machines. Any faulted application on one node fails over to other node that can either be a virtual machine or a physical machine.

See "Standard bridge configuration" on page 66.

**Figure 11-4** VCS cluster across VM guest and physical machine

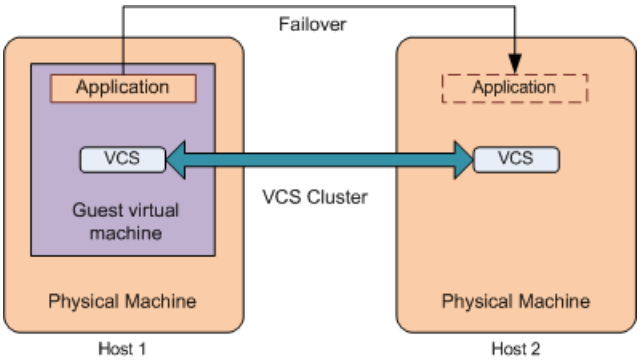


I/O fencing support: Non-SCSI3, CP server based fencing is supported. SCSI3 fencing is not supported.

Chapter 12

# Virtual machine availability

This chapter includes the following topics:



## About virtual machine availability options

While application availability is very important for KVM users, virtual machine availability is equally important. Virtual machine availability can be provided by adopting Symantec Cluster Server (VCS) in the host. VCS in this case monitors the virtual machines as a resource.

See Table 11-1 on page 115.

The virtual machine availability use case is supported for the following Linux virtualization technologies:

- Red Hat Enterprise Linux (RHEL) KVM
- Red Hat Enterprise Virtualization (RHEV)
- SUSE Linux Enterprise Server (SLES) KVM

For setup information for VCS for RHEL and SUSE:

See "Installing and configuring Symantec Cluster Server in a kernel-based virtual machine (KVM) environment" on page 52.

---

**Note:** For virtual machine high availability and failover capabilities the virtual machine image must be on the shared storage accessible to all the nodes of the VCS cluster.

---

**Note:** You can also use the cluster functionality of Symantec Storage Foundation HA or Symantec Storage Foundation Cluster File System HA if you need storage management capabilities in addition to virtual machine availability for your KVM host.

# VCS in host monitoring the Virtual Machine as a resource

In this scenario, Symantec Cluster Server (VCS) runs in the host, enabling host-level clustering. Running VCS in the host also enables the monitoring and fail-over of individual guest virtual machines. Each guest virtual machine is simply a process in the KVM architecture and hence can be monitored by VCS running on the host. This capability allows us to monitor the individual virtual machine as an individual resource and restart/fail-over the VM on the same (or another physical) host. To enable support for guest live migration, Symantec recommends that you run Cluster Volume Manager (CVM) in the host.

In this configuration, the physical machines (PMs) hosting VM guests form a cluster. Therefore, VCS does not monitor applications running inside the guest virtual machines. VCS controls and manages the virtual machines with the help of the KVMGuest agent. If a VM guest faults, it fails over to the other host.

**Note:** The VM guests configured as failover service groups in VCS must have same configuration across all hosts. The storage for the VM guests must be accessible to all the hosts in the cluster.

See "Network configuration for VCS cluster across physical machines (PM-PM)" on page 65.

See "Sample configuration" on page 155.

# Validating the virtualization environment for virtual machine availability

The VCS utility `havirtverify` validates the virtualization environment. If the virtualization environment is not valid for VCS to manage virtual machines, it logs an error message indicating that the virtualization environment is invalid and resource state is UNKNOWN. Upon receiving this error message, you must correct the virtualization environment and run the `havirtverify` utility manually to validate the

environment. Upon successful validation, a verification message displays and the VCS resource state clears in the next monitor cycle.

You can also run this utility manually for verifying the environment.

- Run the `havirtverity` utility manually:

```
# /opt/VRTSvcs/bin/KVMGuest/havirtverify resource_name
```

If validation passes, the following message displays:

```
#/opt/VRTSvcs/bin/KVMGuest/havirtverify resource_name
Red Hat Enterprise Virtualization Environment validation successfully
completed for resource resource_name
```

If validation fails, the following message displays:

```
# /opt/VRTSvcs/bin/KVMGuest/havirtverify resource_name
Virtualization environment validation failed for resource resource_name
```

All the log messages of this utility are sent to the engine log file.

See "Sample configuration" on page 155.

See "Sample configurations for a Red Hat Enterprise Virtualization (RHEV) environment" on page 159.

Chapter 13

# Virtual machine availability for live migration

This chapter includes the following topics:



## About live migration

You can enable live migration of guest virtual machines using shared storage through Cluster Volume Manger (CVM) and Cluster File System (CFS), components of Symantec Storage Foundation Cluster File System HA (SFCFSHA). Using CVM significantly reduces planned downtime for individual virtual machines. Individual virtual machines can now be statefully migrated from host to host, enabling better load-balancing, lower machine downtime and path-management of individual physical servers. Physical servers (hosts) can now join and exit the server pool (physical server cluster) at will while the individual guest virtual machines and their corresponding applications continue to run.

For live migration, by using Fast Failover using CVM/CFS in the guest and host, rather than running a single-node Veritas Volume Manager (VxVM) in the host, you can run the CVM/CFS in the host and cluster multiple physical servers within the same server cluster or server pool. This configuration includes Symantec Cluster Server (VCS) also within the host. The significant advantage of creating a cluster of physical servers is that live migration of KVM guest virtual machines from one physical server to another is fully operational and supported.

**Figure 13-1** Live migration setup for Kernel-based Virtual Machine (KVM)

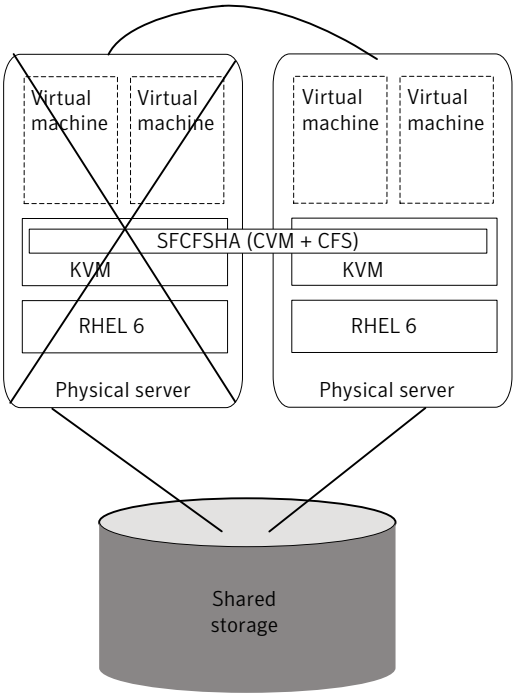


The live migration use case is supported for the following Linux virtualization techologies:

- Red Hat Enterprise Linux (RHEL) KVM
- Red Hat Enterprise Virtualization (RHEV)
- SUSE Linux Enterprise Server (SLES) KVM

**Note:** Storage Foundation is supported only in a KVM environment, it is not supported in a RHEV environment. You can only configure CVM/CFS in a KVM environment.

**Note:** For a RHEV environment, Symantec recommends that you configure Network File System (NFS) or Fibre Channel (FC) type of storage domains.

# Live migration requirements

The following conditions are required for migrating a VM guest from source host to destination host:

- The required guest image must be available on the destination host at the same location.
- The storage and network devices configured in the migrating guest must be identical on source and destination hosts. Any difference may cause the migration process to terminate.
- The KVM hypervisor version on both the hosts should be the same as well as the operating system level.
- For KVM environments, you must set password-less SSH communication between the source and destination host.

VCS-initiated live migration is supported in the following environments:

- RHEL KVM
- SLES KVM
- RHEV

For detailed information about the required and limitation of virtual machine migration, see your Linux virtualization documentation.

# Implementing live migration for virtual machine availability

A virtual machine (VM) can be migrated from one host to another host. This migration can be a live migration or pause migration. You can initiate the migration using:

- The `virsh migrate` command or `virt-manager` console in case of Kernel-based Virtual Machine (KVM) environment
- The RHEV-M console in case of Red Hat Enterprise Virtualization (RHEV) environment
- The Symantec Cluster Server (VCS) `hagrp -migrate` operation (the `hagrp -migrate` command initiates live migration only)

If virtual machine migration is initiated outside VCS (either using the `virsh` commands or the RHEV-M console), VCS monitors the migrated guest and can detect the migration process. VCS changes the resource state according to the virtual machine state, i.e. if the guest is live-migrated from one host to another host, the associated KVMGuest resource is brought online on the host where the guest

is migrated and on the source node the KVMGuest resource state is reported as OFFLINE (Intentional OFFLINE).

For the KVM environment, Symantec recommends the use of CVM and CFS for live migration where a virtual image needs to be simultaneously accessible on a source and destination node.

See "Sample configuration" on page 155.

Symantec Cluster Server (VCS) has introduced a new migrate operation for initiating service group migration. The KVMGuest agent has implemented a "migrate" entry point to initiate virtual machine migration in KVM and RHEV environment. You can initiate a virtual machine live migration using the `hagrp -migrate` command.

The syntax for the command is:

```
#hagrp -migrate service_group_name -to destination_node_name
```

**To verify the password-less SSH requirement for live migration**

- Validate password-less SSH by executing following command on source system:

  ```
  # virsh "connect qemu+ssh://destination_node/system; list"
  ```

  If this command asks for a password, then password-less SSH is not set between source and destination node.

  If proper output is returned, then password-less SSH is set properly.

**To configure VCS to initiate virtual machine migration**

1. To prepare for initiating a virtual machine live migration using `hagrp -migrate` command, you must configure the `PhysicalServer` attribute (system level) of VCS using following command:

   ```
   # hasys -modify  sys_name PhysicalServer physical_server_name
   ```

   For example:

   ```
   # haconf -makerw
   # hasys -modify sys_name PhysicalServer "'hostname'"
   ```

   The `PhysicalServer` name is used while initiating the migration.

2. If `PhysicalServer` attribute is not configured, then the target node name passed to the migrate entry point is used for initiating the migration.

   The KVMGuest Agent `migrate` entry point:

   - For the KVM environment: Agent uses the `virsh migrate` command to initiate virtual machine migration.

- For the RHEV environment: Agent uses REST APIs to initiate virtual machine migration. It also checks whether the virtual machine migration is allowed or not.

See "About the KVMGuest agent" on page 75.

---

**Note:** When a virtual machine is configured for disaster recovery, the virtual machine cannot be migrated across sites.

---

See "Sample configurations for a Red Hat Enterprise Virtualization (RHEV) environment" on page 159.

Chapter 14

# Virtual to virtual clustering in a Red Hat Enterprise Virtualization environment

This chapter includes the following topics:



## Installing and configuring Symantec Cluster Server for Red Hat Enterprise Virtualization (RHEV) virtual-to-virtual clustering

Red Hat Enterprise Virtualization (RHEV) is a server virtualization solution that uses a KVM hypervisor. As KVM forms a core part of the Linux kernel, this virtualization is highly efficient in Linux environments. Platform management infrastructure and application-specific agents, and other tools are the other components of a RHEV setup.

To enable VCS support for in-guest clustering, before you install VCS on the guest virtual machines, you must set up a private network between them. This involves the following steps:

- Add the two NICs to the virtual machine for private communication

**Note:** Symantec recommends that you add one more interface/NIC to the virtual machine for public communication. Also, if the virtual machines for which you are configuring the network run on separate physical hosts, ensure that you set up an LLT communication channel between the physical hosts.

- Attach a switch to each of the two additional NICs

**To create a network on the physical host**

1. From RHEV Manager, create two new logical networks for private LLT heartbeat communication.
2. Assign appropriate physical interfaces to the newly-created logical networks.

**To configure a logical network for virtual machines**

1. Create two network interfaces, of IntelPro 'e1000' type, and associate them with the newly-created logical networks.
2. Repeat step 1 for each virtual machine where you want to monitor application availability with VCS.

To set up a cluster of virtual (guest) machines with Symantec Cluster Server (VCS), perform the following procedures:

- Consult the requirements in:
  *Symantec Cluster Server Release Notes*
- Install VCS on the guest virtual machine:
  *Symantec Cluster Server Installation Guide*
- Configure VCS in the guest virtual machine
  *Symantec Cluster Server Installation Guide*

**Note:** The installation and configuration of VCS inside a virtual machine is similar to that of the physical system. No additional VCS configuration is required to make it work inside the virtual machine.

For more details, see the *Symantec Cluster Server Administrator's Guide*.

## Fencing support for VCS in-guest clusters

VCS supports non-SCSI3, CP server-based fencing in virtual machines to prevent corruption of data disks. SCSI3 fencing is not supported.

For information on configuring fencing, see the *Symantec Cluster Server Installation Guide*.

### Live migration support

VCS in-guest clustering continues to provide high availability of applications on virtual machines, in live migration scenarios initiated by the virtualization technology.

Symantec has tested for live migration support in the RHEV environment under the following conditions:

- Virtual machine image resides on NFS, iSCSI, or FC storage domain

# Storage configuration for VCS in a RHEV environment

To fail over an application from one virtual machine to another, it is mandatory to store the application data on storage shared between the two virtual machines. In an RHEV environment, Symantec has tested application failovers with the application data residing on:

- iSCSI LUNs directly attached to the virtual machine
- NFS exported directory mounted inside virtual machine
- Fibre Channel-based LUNs

---

**Note:** Symantec recommends using a dedicated virtual network for iSCSI storage.

---

Chapter 15

# Virtual to virtual clustering in a Microsoft Hyper-V environment

This chapter includes the following topics:



## Installing and configuring Symantec Cluster Server with Microsoft Hyper-V virtual-to-virtual clustering

The Microsoft Hyper-V role in Windows Server 2008 and Windows Server 2008 R2 is a hypervisor based server virtualization technology for the x86_64 architecture. It provides you with the software infrastructure and management tools that you can use to create and manage a virtualized server computing environment.

To enable VCS support for in-guest clustering, before you install VCS on the guest virtual machines, you must set up a private network between them. This involves the following steps:

- Add two NICs to the virtual machine for private communication

  **Note:** Symantec recommends that you add one more interface/NIC to the virtual machine for public communication. Also, if the virtual machines for which you are configuring the network run on separate physical hosts, ensure that you set up an LLT communication channel between the physical hosts.

- Attach a switch to each of the two additional NICs

**To create a virtual network on the physical host**

1. From the Hyper-V manager, create two virtual networks for private LLT heartbeat communication.
2. Assign appropriate physical interfaces to the newly-created virtual networks.

**To configure the network for the virtual machines**

1. Create two network interfaces of 'Legacy Network Adaptor' type, and associate them with the newly-created virtual networks.
2. Repeat step 1 for each virtual machine where you want to monitor application availability with VCS.

To set up a cluster of virtual (guest) machines with Symantec Cluster Server (VCS), perform the following procedures:

- Consult the requirements in:
  *Symantec Cluster Server Release Notes*
- Install VCS on the guest virtual machine:
  *Symantec Cluster Server Installation Guide*
- Configure VCS in the guest virtual machine
  *Symantec Cluster Server Installation Guide*

---

**Note:** The installation and configuration of VCS inside a virtual machine is similar to that of the physical system. No additional VCS configuration is required to make it work inside the virtual machine.

For more details, see the *Symantec Cluster Server Administrator's Guide*.

---

## Fencing support for VCS in-guest clusters

VCS supports non-SCSI3, CP server-based fencing in virtual machines to prevent corruption of data disks. SCSI3 fencing is not supported.

For information on configuring fencing, see the *Symantec Cluster Server Installation Guide*.

## Live migration support

VCS in-guest clustering continues to provide high availability of applications on virtual machines, in live migration scenarios initiated by the virtualization technology.

Symantec has tested for live migration support in the Hyper-V environment under the following conditions:

- Microsoft Failover Clustering is enabled
- Virtual machine image resides on Microsoft Clustered Shared Volumes

Chapter 16

# Virtual to virtual clustering in a Oracle Virtual Machine (OVM) environment

This chapter includes the following topics:



## Installing and configuring Symantec Cluster Server for Oracle Virtual Machine (OVM) virtual-to-virtual clustering

Oracle VM is an enterprise-grade server virtualization solution that supports guest (virtual machines) that supports various operating systems, including Linux. Based on the Xen hypervisor technology, OVM also provides you with an integrated, Web-based management console.

Before you install VCS on the guest virtual machines, you must set up a private network between them. This involves the following steps:

**To enable VCS support of virtual-to-virtual clustering**

- Set up a private network between the guest virtual machines.
  - Apart from the public NIC on each physical host, create two additional NICs.

**Note:** Symantec recommends that you add one more interface/NIC to the virtual machine for public communication. Also, if the virtual machines for which you are configuring the network run on separate physical hosts, ensure that you set up an LLT communication channel between the physical hosts.

If the virtual machines for which you configure the network run on separate physical hosts, ensure that you create a LLT communication channel between the physical hosts.

- Attach a switch to each of the two additional NICs

**To create a private network on the physical host**

1. From the Oracle VM Manager, create two virtual networks for private LLT heartbeat communication.
2. Assign appropriate physical interfaces to the newly-created virtual networks.

**To configure the network for virtual machines**

1. Create two interfaces (in a network that is created with the option **Create a hybrid network with bonds/ports and VLANS**) and associate the interfaces with the newly-created virtual networks.
2. Repeat step 1 for each virtual machine where you want to monitor availability with VCS.

To set up a cluster of virtual (guest) machines with Symantec Cluster Server (VCS), perform the following procedures:

- Consult the requirements in:
  *Symantec Cluster Server Release Notes*
- Install VCS on the guest virtual machine:
  *Symantec Cluster Server Installation Guide*
- Configure VCS in the guest virtual machine
  *Symantec Cluster Server Installation Guide*

**Note:** The installation and configuration of VCS inside a virtual machine is similar to that of the physical system. No additional VCS configuration is required to make it work inside the virtual machine.

For more details, see the *Symantec Cluster Server Administrator's Guide*.

### Live migration support

Symantec has supported live migration in the OVM environment under the following conditions:

- Virtual machine image resides on NFS data domains

### Fencing support for VCS in-guest clusters

VCS supports non-SCSI3, CP server-based fencing in virtual machines to prevent corruption of data disks.

For information on configuring fencing, see the *Symantec Cluster Server Installation Guide*.

# Storage configuration for VCS support in Oracle Virtual Machine (OVM)

To fail over an application from one virtual machine to another, it is mandatory to store the application data on storage shared between the two virtual machines. In an OVM environment, Symantec has tested application failovers with the application data residing on:

- Local disks
- Shared Network Attached Storage (NFS)
- Shared iSCSI SANs: abstracted LUNs or raw disks accessible over existing network infrastructure
- Fibre Channel SANs connected to one or more host bus adapters (HBAs)

---

**Note:** For more information, see *Oracle* documentation.

---

VCS in-guest clustering continues to provide high availability of applications on virtual machines, in live migration scenarios initiated by the virtualization technology.

Chapter 17

# Disaster recovery for virtual machines in the Red Hat Enterprise Virtualization environment

This chapter includes the following topics:



## About disaster recovery for Red Hat Enterprise Virtualization virtual machines

Red Hat Enterprise Virtualization (RHEV) virtual machines can be configured for disaster recovery (DR) by replicating their boot disks using replication methods such as Hitachi TrueCopy or EMC SRDF. The network configuration for the virtual machines in the primary site may not be effective in the secondary site if the two sites are in different IP subnets. Hence you must make some additional configuration changes to the KVMGuest resource managing the virtual machine.

See Table 4-3 on page 72.

Supported technologies for replicating virtual machines include:

- EMC SRDF
- Hitachi TrueCopy

**Note:** Live migration of virtual machines across replicated sites is not supported.

**Figure 17-1** Schematic of the RHEV DR setup

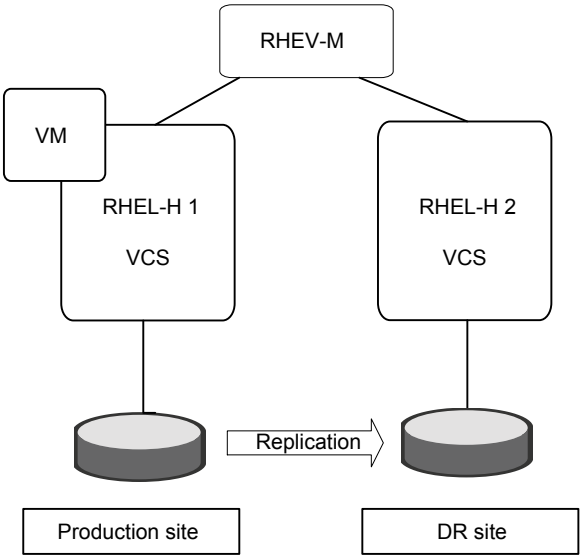


Disaster recovery use cases for virtual machines work in the following way:

- The replication agent takes care of the replication direction. After a disaster event at the primary site, VCS tries to online the replication service group at the secondary site (according to the ClusterFailoverPolicy). The replication resource reverses the replication direction. Reversing the replication direction makes sure that the old secondary LUNs become the new primary LUNs and also are Read-Write enabled on the RHEL-H hosts at the secondary site. This helps RHEV-M activate the Fibre Channel (FC) Storage Domain on the secondary site RHEL-H hosts.
- Before the virtual machine (VM) service group can be brought online, the Storage Pool Manager (SPM) in the datacenter needs to failover to the secondary site. This is achieved by the pre-online trigger script configured on the VM service group. This trigger script checks whether the SPM is still active in the primary RHEV cluster. If so, it deactivates all the RHEL-H hosts in the primary RHEV cluster. Additionally, if the SPM host in the primary RHEV cluster is in the NON_RESPONSIVE state, the trigger fences out the host to enable SPM failover. The trigger script then waits for the SPM to failover to the secondary RHEV cluster. When the SPM successfully fails over to the secondary RHEV cluster, the pre-online trigger script reactivates all the RHEL-H hosts in the primary

RHEV cluster, which were deactivated earlier and proceeds to online the VM service group in the secondary site

**Figure 17-2** VCS Resource dependency diagram

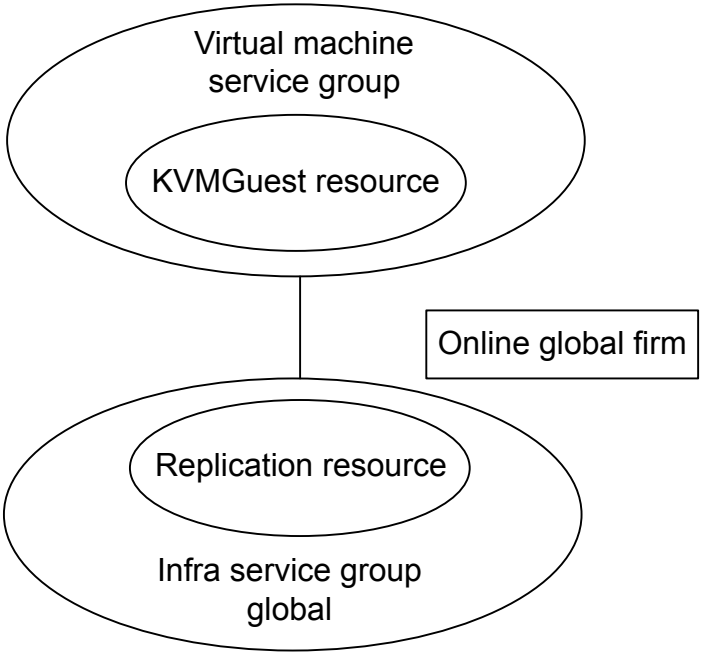


# Configuring Red Hat Enterprise Virtualization (RHEV) virtual machines for disaster recovery using Symantec Cluster Server (VCS)

You can configure new or existing RHEV-based virtual machines for disaster recovery (DR) by setting them up and configuring VCS for DR.

**To set up RHEV-based virtual machines for DR**

1. Configure VCS on both sites in the RHEL-H hosts, with the GCO option.

   For more information about configuring a global cluster: see the *Symantec Storage Foundation and High Availability Solutions Disaster Recovery Implementation Guide*.

2. Configure replication setup using a replication technology such as Hitachi TrueCopy or EMC SRDF.

3 Map the primary LUNs to all the RHEL-H hosts in the primary site.

4 Issue OS level SCSI rescan commands and verify that the LUNs are visible in the output of the `multipath -l` command.

5 Map the secondary LUNs to all the RHEL hosts in the secondary site and verify that they are visible in the output of the `multipath -l` command on all the hosts in the secondary site.

6 Add the RHEL-H hosts to the RHEV-M console.

- Create two RHEV clusters in the same datacenter, representing the two sites.
- Add all the RHEL-H hosts from the primary site to one of the RHEV clusters.
- Similarly, add all the RHEL-H hosts from the secondary site to the second RHEV cluster.

7 Log in to the RHEV-M console and create a Fibre Channel-type Storage Domain on one of the primary site hosts using the primary LUNs.

8 In the RHEV-M console, create a virtual machine and assign a virtual disk carved out of the Fibre Channel Storage Domain created in 7.

- Configure any additional parameters such as NICs and virtual disk for the virtual machine.
- Verify that the virtual machine turns on correctly.
- Install appropriate RHEL operating system inside the guest.
- Configure the network interface with appropriate parameters such as IP address, Netmask, and gateway.
- Make sure that the NIC is not under network manager control. You can disable this setting by editing the `/etc/sysconfig/network-scripts/ifcfg-eth0` file inside the virtual machine and setting NM_CONTROLLED to "no".
- Make sure that the virtual machine does not have a CDROM attached to it. This is necessary since VCS sends the DR payload in the form of a CDROM to the virtual machine.

9 Copy the package VRTSvcsnr from the VCS installation media to the guest and install it. This package installs a lightweight service which starts when the guest boots. The service reconfigures the IP address and Gateway of the guest as specified in the KVMGuest resource.

**To configure VCS for managing RHEV-based virtual machines for DR**

1 Install VCS in the RHEL-H hosts at both the primary and the secondary sites.

- Configure all the VCS nodes in the primary site in a single primary VCS cluster.
- Configure all the VCS nodes in the secondary site in the same secondary VCS cluster.
- Make sure that the RHEV cluster at each site corresponds to the VCS cluster at that site.

See Figure 17-2 on page 138.

2 Create a service group in the primary VCS cluster and add a KVMGuest resource for managing the virtual machine. Repeat this step in the secondary VCS cluster.

3 Configure site-specific parameters for the KVMGuest resource in each VCS cluster.

- The `DROpts` attribute enables you to specify site-specific networking parameters for the virtual machine such as IP Address, Netmask, Gateway, DNSServers, DNSSearchPath and Device. The Device is set to the name of the NIC as seen by the guest, for example eth0.
- Verify that the ConfigureNetwork key in the `DROpts` attribute is set to 1.
- The `DROpts` attribute must be set on the KVMGuest resource in both the clusters.

4 Configure the preonline trigger on the virtual machine service group. The preonline trigger script is located at `/opt/VRTSvcs/bin/sample_triggers/VRTSvcs/preonline_rhev`.

- Create a folder in the `/opt/VRTSvcs` directory on each RHEL-H host to host the trigger script. Copy the trigger script in this folder with the name "preonline". Enable the preonline trigger on the virtual machine service group by setting the PreOnline service group attribute. Also, specify the path (relative to `/opt/VRTSvcs`) in the TriggerPath attribute.

For example:

```
group RHEV_VM_SG1 (
    SystemList = { vcslx317 = 0, vcslx373 = 1 }
    ClusterList = { test_rhevdr_pri = 0, test_rhevdr_sec = 1 }
    AutoStartList = { vcslx317 }
    TriggerPath = "bin/triggers/RHEVDR"
    PreOnline = 1
    )
```

For more information on setting triggers, see the *Symantec Cluster Server Administrator's Guide*.

5 Create a separate service group for managing the replication direction. This task must be performed for each cluster.

- Add the appropriate replication resource (such as Hitachi TrueCopy or EMC SRDF). For details on the appropriate replication agent, see the *Replication Agent Installation and Configuration Guide* for that agent.
- Add an Online Global Firm dependency from the virtual machine (VM) service group to the replication service group.
- Configure the replication service group as global.

6 Configure the postonline trigger on the replication service group. The postonline trigger script is located at `/opt/VRTSvcs/bin/sample_triggers/VRTSvcs/postonline_rhev`.

- Copy the postonline trigger to the same location as the preonline trigger script, with the name "postonline". Enable the postonline trigger on the replication service group by adding the POSTONLINE key to the TriggersEnabled attribute. Also, specify the path (relative to `/opt/VRTSvcs`) in the TriggerPath attribute.
  For example:

```
group SRDF_SG1 (
    SystemList = { vcslx317 = 0, vcslx373 = 1 }
    ClusterList = { test_rhevdr_pri = 0, test_rhevdr_sec = 1 }
    AutoStartList = { vcslx317 }
    TriggerPath = "bin/triggers/RHEVDR"
    TriggersEnabled = { POSTONLINE }
    )
```

  For more information on setting triggers, see the *Symantec Cluster Server Administrator's Guide.*

If you have multiple replicated Storage Domains, the replication direction for all the domains in a datacenter must be the same.

**To align replication for multiple replicated Storage Domains in a datacenter**

1 Add all the replication resources in the same Replication Service Group.

2 If you require different Storage Domains to be replicated in different directions at the same time, configure them in a separate datacenter.

  This is because the Storage Pool Manager (SPM) host requires read-write access to all the Storage Domains in a datacenter.

After completing all the above steps, you can easily switch the virtual machine service group from one site to the other. When you online the replication service group in a site, the replication resource makes sure that the replication direction is

from that site to the remote site. This ensures that all the replicated devices are read-write enabled in the current site.

See "About disaster recovery for Red Hat Enterprise Virtualization virtual machines" on page 136.

**Disaster recovery workflow**

1. Online the replication service group in a site followed by the virtual machine service group.
2. Check the failover by logging into the RHEV-M console. Select the **Hosts** tab of the appropriate datacenter to verify that the SPM is marked on one of the hosts in the site in which the replication service group is online.
3. When you bring the Replication Service Group online, the postonline trigger probes the KVMGuest resources in the parent service group. This is to ensure that the virtual machine service group can go online.
4. When you bring the virtual machine service group online, the preonline trigger performs the following tasks:
   - The trigger checks whether the SPM is in the local cluster. If the SPM is in the local cluster, the trigger checks whether the SPM host is in the UP state. If the SPM host is in the NON_RESPONSIVE state, the trigger fences out the host. This enables RHEV-M to select some other host in the current cluster.
   - If the SPM is in the remote cluster, the trigger deactivates all the hosts in the remote cluster. Additionally, if the remote SPM host is in the NON_RESPONSIVE state, the trigger script fences out the host. This enables RHEV-M to select some other host in the current cluster.
   - The trigger script then waits for 10 minutes for the SPM to failover to the local cluster.
   - When the SPM successfully fails over to the local cluster, the script then reactivates all the remote hosts that were previously deactivated.
   - Then the trigger script proceeds to online the virtual machine service group.
5. When the KVMGuest resource goes online, the KVMGuest agent sets a virtual machine payload on the virtual machine before starting it. This payload contains the site-specific networking parameters that you set in the DROpts attribute for that resource.
6. When the virtual machine starts, the `vcs-net-reconfig` service is loaded and reads the DR parameters from the CDROM and then applies them to the guest. This way, the networking personality of the virtual machine is modified when the virtual machine crosses site boundaries.

### Troubleshooting a disaster recovery configuration

- You can troubleshoot your disaster recovery in the following scenarios:
  - When the service groups are switched to the secondary site, the hosts in the primary site may go into the NON_OPERATIONAL state. To resolve this issue, deactivate the hosts by putting them in maintenance mode, and reactivate them. If the issue is not resolved, log onto the RHEL-H host and restart the `vdsmd` service using the `service vdsmd restart`command. If the issue still persists, please contact RedHat Technical Support.
  - After a DR failover, the DNS configuration of the virtual machine may not change. To resolve this issue, check if the network adapter inside the virtual machine is under Network Manager control. If so, unconfigure the network adapter by editing the `/etc/sysconfig/network-scripts/ifcfg-eth0` file inside the virtual machine and setting NM_CONTROLLED to "no".
  - After a failover to the secondary site, the virtual machine service group does not go online. To resolve this issue, check the state of the SPM in the data center. Make sure that the SPM is active on some host in the secondary RHEV cluster. Additionally, check the VCS engine logs for more information.

Chapter 18

# Multi-tier business service support

This chapter includes the following topics:



## About Virtual Business Services

The Virtual Business Services feature provides visualization, orchestration, and reduced frequency and duration of service disruptions for multi-tier business applications running on heterogeneous operating systems and virtualization technologies. A virtual business service represents the multi-tier application as a consolidated entity that helps you manage operations for a business service. It builds on the high availability and disaster recovery provided for the individual tiers by Symantec products such as Symantec Cluster Server and Symantec ApplicationHA.

Application components that are managed by Symantec Cluster Server, Symantec ApplicationHA, or Microsoft Failover Clustering can be actively managed through a virtual business service.

You can use the Veritas Operations Manager Management Server console to create, configure, and manage virtual business services.

## Sample virtual business service configuration

This section provides a sample virtual business service configuration comprising a multi-tier application. Figure 18-1 shows a Finance application that is dependent

on components that run on three different operating systems and on three different clusters.

- Databases such as Oracle running on Solaris operating systems form the database tier.
- Middleware applications such as WebSphere running on AIX operating systems form the middle tier.
- Web applications such as Apache and IIS running on Windows and Linux virtual machines form the Web tier. This tier is composed of ApplicationHA nodes.

  Each tier can have its own high availability mechanism. For example, you can use Symantec Cluster Server for the databases and middleware applications, and Symantec ApplicationHA for the Web servers.

**Figure 18-1** Sample virtual business service configuration

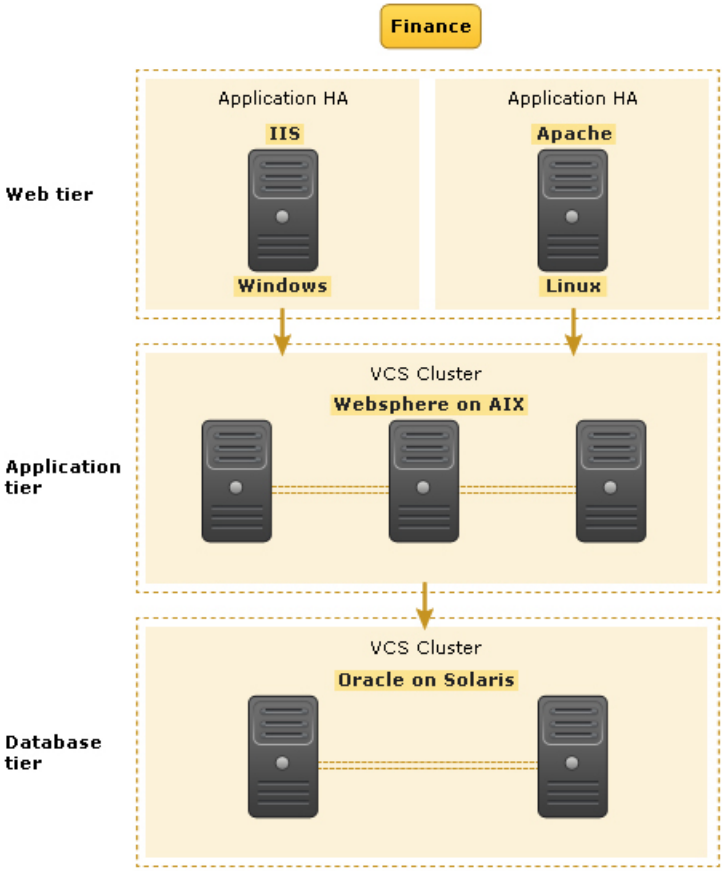


Each time you start the Finance business application, typically you need to bring the components online in the following order – Oracle database, WebSphere, Apache and IIS. In addition, you must bring the virtual machines online before you start the Web tier. To stop the Finance application, you must take the components offline in the reverse order. From the business perspective, the Finance service is unavailable if any of the tiers becomes unavailable.

When you configure the Finance application as a virtual business service, you can specify that the Oracle database must start first, followed by WebSphere and the Web servers. The reverse order automatically applies when you stop the virtual business service. When you start or stop the virtual business service, the components of the service are started or stopped in the defined order.

For more information about Virtual Business Services, refer to the *Virtual Business Service–Availability User's Guide*.

Section 5

# Reference

Appendix A

# Troubleshooting

This appendix includes the following topics:



## Troubleshooting virtual machine live migration

A VCS cluster is formed between virtual machines (VMs) and one of the virtual machines is migrated from one host to another host. During a virtual machine migration, if the VM takes more than 16 seconds to migrate to the target node, one of the VMs panics. In this case, 16 seconds is the default value of the LLT `peerinact`

parameter. You can increase the `peerinact` value to allow sufficient time for the VM to migrate. You can adjust this time based on the environment in which you initiate the VM migration.

To avoid false failovers for virtual machine migration, you can change the `peerinact` value using the following methods:

- Set the `peerinact` value dynamically using lltconfig command:

  ```
  # lltconfig -T peerinact:value
  ```

- Set the `peerinact` value in the `/etc/llttab` file to make the value persistent across reboots.

**To set the peerinact value dynamically using lltconfig command**

1. Determine how long the migrating node is unresponsive in your environment.
2. If that time is less than the default LLT peer inactive timeout of 16 seconds, VCS operates normally.

   If not, increase the peer inactive timeout to an appropriate value on all the nodes in the cluster before beginning the migration.

   For example, to set the LLT `peerinact` timeout to 20 seconds, use the following command:

   ```
   # lltconfig -T peerinact:2000
   ```

   The value of the `peerinact` command is in .01 seconds.
3. Verify that `peerinact` has been set to 20 seconds:

   ```
   # lltconfig -T query

   Current LLT timer values (.01 sec units):
    heartbeat   = 50
    heartbeatlo = 100
    peertrouble = 200
    peerinact   = 2000
    oos         = 10
    retrans     = 10
    service     = 100
    arp         = 30000
    arpreq         = 3000
    Current LLT flow control values (in packets):
    lowwater  = 40
   ```

4 Repeat steps 2 to 3 on other cluster nodes.

5 Reset the value back to the default `peerinact` value using the `lltconfig` command after the migration is complete.

**To make the LLT peerinact value persistent across reboots:**

- Append the following line at the end of `/etc/lltab` file to set the LT `peerinact` value to 20 seconds:

  ```
  set-timer peerinact:2000
  ```

  After appending the above line, `/etc/lltab` file should appear similar to the following:

  ```
  # cat /etc/lltab
  set-node sys1
  set-cluster 1234
  link eth2 eth-00:15:17:48:b5:80 - ether - -
  link eth3 eth-00:15:17:48:b5:81 - ether - -
  set-timer peerinact:2000
  ```

For more information on VCS commands, see the *Symantec Cluster Server Administrator's Guide*.

For attributes related to migration, see the *Symantec Cluster Server Bundled Agents Reference Guide*.

# Live migration storage connectivity in a Red Hat Enterprise Virtualization (RHEV) environment

In a RHEV environment, if a virtual machine (VM) is migrating from one host to another and source host loses storage connectivity then the VM remains in the paused state. This issue is RHEV environment specific.

There is no workaround.

# Troubleshooting Red Hat Enterprise Virtualization (RHEV) virtual machine disaster recovery (DR)

When you fail over the replication service group from one site to another, the hosts in the old site may go into the NON_RESPONSIVE state in the RHEV-M console.

**To resolve the hosts in the NON_RESPONSIVE state in the RHEV-M console**

1. Move the host into MAINTENENCE mode.
2. Try to ACTIVATE the host using the RHEV-M console.
3. If the issue still persists, contact Redhat Support to get it resolved.

# The KVMGuest resource may remain in the online state even if storage connectivity to the host is lost

When a virtual machine is running on a physical host and loses storage connectivity, the virtual machine goes into the PAUSED state. However, the virtual machine process is still running. The KVMGuest resource monitoring the virtual machine reports the state as ONLINE as the virtual machine process is still running and no failover is initiated. The KVMGuest resource is not aware of the storage situation, and therefore does not take any action.

If this issue occurs, either offline the service group or manually switch the service group. This shuts down the virtual machine and starts the virtual machine on another node.

# VCS initiates a virtual machine failover if a host on which a virtual machine is running loses network connectivity

When a virtual machine is running on a physical host and loses network connectivity, such as a public or private communication channel, VCS on each node is not able to communicate. This is a classic split brain situation. VCS running on a node thinks that the other node has crashed and initiates a virtual machine failover.However, the virtual machine is still running on one node while VCS attempts to start same virtual machine on another node.

If this issue occurs, configure disk based fencing to prevent a split brain situation due to a network partition.

# Virtual machine start fails due to having the wrong boot order

When creating a virtual machine, you can specify the boot order. If a virtual machine has the following boot order, the virtual machine start fails as it is not able to find the CD-ROM:

- CD-ROM
- Hard disk

If VCS initiated the virtual machine start, any associated KVMGuest resources also fail. This issue is due to RHEV behavior.

If this issue occurs, manually edit the boot order and remove the CD-ROM from the boot sequence. Then re-initiate the virtual machine start using VCS or the RHEV-M console.

# Virtual machine hangs in the wait_for_launch state and fails to start

When a virtual machine start is initiated through the RHEV-M console, the virtual machine may hang in the wait_for_launch state, and then fails to start. This issue occurs when the `libvirtd` service is unable to process the virtual machine start operation.

There is no workaround.

# VCS fails to start a virtual machine on a host in another RHEV cluster if the DROpts attribute is not set

In the RHEV environment, every host is part of a RHEV cluster. In a local high availability scenario, hosts forming a VCS cluster should be part of a single RHEV cluster. However, in disaster recovery scenarios, you can configure all hosts on the primary site in one RHEV cluster and all hosts on the secondary site in a different RHEV cluster, though they are all part of the same datacenter. During a site failover, when the `DROpts` attribute is set, VCS changes the virtual machine host as per the new RHEV cluster.

If the `DROpts` attribute is not set, VCS does not allow a host from a different RHEV cluster to start the virtual machine. This issue occurs because virtual machine migration does not work across RHEV clusters. Therefore, VCS fails to start the virtual machine on a host that is part of a different cluster.

Symantec recommends configuring hosts in different clusters only in a disaster recovery configuration, and setting the `DROpts` attribute of the KVMGuest agent. For a local high availability scenario, you do not need to set the `DROpts` attribute, and all the hosts forming a VCS cluster should be part of the same RHEV cluster.

# Virtual machine fails to detect attached network cards

A virtual machine may fail to detect an attached network interface. This issue is due to RHEV behavior.

There is no workaround.

# The KVMGuest agent behavior is undefined if any key of the RHEVMInfo attribute is updated using the -add or -delete options of the hares -modify command

If you modify any key of the `RHEVMInfo` attribute using the `-add` or `-delete` options of the `hares -modify` command, the `RHEVMInfo` attribute information sequence changes and can cause the KVMGuest resource behavior to be undefined. The `-add` option adds a new key to any attribute, and the `-delete` option deletes a key from any attribute. These two options should not be used to configure the `RHEVMInfo` attribute.

Use the `-update` option of the `hares -modify` command to modify attribute keys:

```
# hares -modify resource_name RHEVMInfo -update key_name value
```

For example:

```
# hares -modify vmres RHEVMInfo -update User "admin@internal"
```

Appendix B

# Sample configurations

This appendix includes the following topics:



## Sample configuration

Host configuration:

```
include "types.cf"

cluster kvmclus (
)
system sys1 (
)
system sys2 (
)
group rsg (
SystemList = { sys1 = 0, sys2 = 1 }
)

RemoteGroup rvg1 (
IpAddress = "192.203.47.61"
Username = vcsuser
Password = CQIoFQf
GroupName = appsg
VCSSysName = sys1
ControlMode = OnOff
)
```

```
requires group vmgrp online local hard

group vmgrp (
SystemList = { sys1 = 0, sys2 = 1 }
)
Application app1(
                StartProgram = "/usr/bin/virsh start PM1vxfsVM1"
                StopProgram = "/usr/bin/virsh shutdown PM1vxfsVM1"
                PidFiles = { "/var/run/libvirt/qemu/PM1vxfsVM1.pid" }
 )
```

VM guest configuration:

```
include "types.cf"

cluster appclus (
)

system sys1 (
)
group appsg (
SystemList = { sys1 = 0 }
)

Process proc (
PathName = /test
)
```

## Sample configuration 1: Native LVM volumes are used to store the guest image

```
group kvmtest1 (
SystemList = { sys1 = 0, sys2 = 1 }
)
KVMGuest res1 (
GuestName = kvmguest1
GuestConfigFilePath = "/kvmguest/kvmguest1.xml"
DelayAfterGuestOnline = 10
DelayAfterGuestOffline = 35
)
Mount mnt1 (
```

```
BlockDevice = "/dev/mapper/kvmvg-kvmvol"
MountPoint = "/kvmguest"
FSType = ext3
FsckOpt = "-y"
MountOpt = "rw"
)
LVMLogicalVolume lv1 (
VolumeGroup = kvmvg
LogicalVolume = kvmvol
)
LVMVolumeGroup vg1 (
VolumeGroup = kvmvg
)
res1 requires mnt1
mnt1 requires lv1
lv1 requires vg1
```

## Sample configuration 2: VxVM volumes are used to store the guest image

```
group kvmtest2 (
SystemList = { sys1 = 0, sys2 = 1 }
)
KVMGuest res1 (
GuestName = kvmguest1
GuestConfigFilePath = "/kvmguest/kvmguest1.xml"
DelayAfterGuestOnline = 10
DelayAfterGuestOffline = 35
)
Mount mnt1 (
BlockDevice = "/dev/vx/dsk/kvmvg/kvmvol"
MountPoint = "/kvmguest"
FSType = vxfs
FsckOpt = "-y"
MountOpt = "rw"
)
Volume vol1 (
Volume = kvm_vol
DiskGroup = kvm_dg
)
DiskGroup dg1 (
DiskGroup = kvm_dg
)
```

```
res1 requires mnt1
mnt1 requires vol1
vol1 requires dg1
```

## Sample configuration 3: CVM-CFS is used to store the guest image

```
group kvmgrp (
SystemList = { kvmpm1 = 0, kvmpm2 = 1 }
)
KVMGuest kvmres (
GuestName = kvmguest1
GuestConfigFilePath = "/cfsmount/kvmguest1.xml"
DelayAfterGuestOnline = 10
DelayAfterGuestOffline = 35
)

kvmgrp requires group cvm online local firm

group cvm (
SystemList = { kvmpm1 = 0, kvmpm2 = 1 }
AutoFailOver = 0
Parallel = 1
AutoStartList = { kvmpm1, kvmpm2 }
)
CFSMount cfsmount (
MountPoint = "/cfsmount"
BlockDevice = "/dev/vx/dsk/cfsdg/cfsvol"
)
CFSfsckd vxfsckd (
)
CVMCluster cvm_clus (
CVMClustName = kvmcfs
CVMNodeId = { kvmpm1 = 0, kvmpm2 = 1 }
CVMTransport = gab
CVMTimeout = 200
)
CVMVolDg cfsdg (
CVMDiskGroup = cfsdg
CVMVolume = { cfsvol }
CVMActivation = sw
)
CVMVxconfigd cvm_vxconfigd (
Critical = 0
```

```
CVMVxconfigdArgs = { syslog }
)

cfsmount requires cfsdg
cfsmount requires cvm_clus
cvm_clus requires cvm_vxconfigd
vxfsckd requires cvm_clus
```

# Sample configurations for a Red Hat Enterprise Virtualization (RHEV) environment

Sample configuration for a RHEV-based service group:

```
group rhev_grp1 (
        SystemList = { sys1 = 0, sys2 = 1 }
        )

        KVMGuest kvmres1 (
                RHEVMInfo = { Enabled = 1,
                        URL = "https://rhevm-server.example.com:443",
                        User = "admin@internal"
                        Password = bncNfnOnkNphChdHe,
                        Cluster = dc2_cluster1,
                        UseManualRHEVMFencing=1 }
                GuestName = rhevvm1
                DelayAfterGuestOnline = 20
                DelayAfterGuestOffline = 35
                )
```

Sample configuration for an AD-based domain:

```
include "types.cf"

cluster kvmtest (
        UserNames = { admin = bQRjQLqNRmRRpZRlQO }
        Administrators = { admin }
        )

system sys1 (
        )

system sys2 (
        )
```

```
group virt_grp (
        SystemList = { sys1 = 0, sys2 = 1 }
        )

        KVMGuest virt_res (
                RHEVMInfo = { Enabled = 1,
                        URL = "https://rhevm.example.com:443",
                        User = rhevmadmin@example.com",
                        Password = codOgoPolOqiDieIf,
                        Cluster = cluster_NFS,
                        UseManualRHEVMFencing=0 }
                GuestName = VM1
                )
```

Sample configuration for a RHEV-based disaster recovery service group:

```
group VM_SG (
        SystemList = { rhelh_a1 = 0, rhelh_a2 = 1 }
        TriggerPath = "bin/triggers/RHEVDR"
        PreOnline = 1
        OnlineRetryLimit = 2
        )

        KVMGuest kvm_res (
                RHEVMInfo = { Enabled = 1, URL = "https://192.168.72.11:443",
                        User = "admin@internal",
                        Password = CQIoFQf,
                        Cluster = RHEV-PRIM-CLUS,
                        UseManualRHEVMFencing = 1 }
                GuestName = swvm02
                DROpts = { ConfigureNetwork = 1,
                        IPAddress = "192.168.74.21",
                        Netmask = "255.255.252.0",
                        Gateway = "192.168.74.1",
                        DNSServers = "143.127.176.14",
                        DNSSearchPath = "rhevdc.com",
                        Device = eth0 }
                )

        requires group STORAGE online global soft


        // resource dependency tree
```

```
        //
        //      group VM_SG
        //      {
        //      KVMGuest kvm_res
        //      }


group STORAGE (
        SystemList = { rhelh_a1 = 0, rhelh_a2 = 1 }
        ClusterList = { RHEV_SEC = 0, RHEV_PRIM = 1 }
        TriggerPath = "bin/triggers/RHEVDR"
        TriggersEnabled = { POSTONLINE }
        )

        SRDF srdf_res1 (
                GrpName = rhevdr
                )

        SRDF srdf_res2 (
                GrpName = rhevdr2
                )



        // resource dependency tree
        //
        //      group STORAGE
        //      {
        //      SRDF srdf_res1
        //      SRDF srdf_res2
        //      }
```

Sample configuration for a multi-resource configuration in a RHEV environment:

```
system sys1 (
        )

system sys2 (
        )

group rhevgrp1 (
        SystemList = { sys1 = 0, sys2 = 1 }
        )
```

```
        KVMGuest vmres1 (
                RHEVMInfo = { Enabled = 1,
                        URL = "https://rhevm.example.com:443",
                        User = "admin@internal",
                        Password = FRGrJRsROrTLgLHlI,
                        Cluster = vcs_clus,
                        UseManualRHEVMFencing = 0 }
                GuestName = vcsvm1
                DelayAfterGuestOnline = 10
                DelayAfterGuestOffline = 35
                )


group rhevgrp2 (
        SystemList = { sys1 = 0, sys2 = 1 }
        )

        KVMGuest vmres2 (
                RHEVMInfo = { Enabled = 1,
                        URL = "https://rhevm.example.com:443",
                        User = "admin@internal",
                        Password = FRGrJRsROrTLgLHlI,
                        Cluster = vcs_clus,
                        UseManualRHEVMFencing = 0 }
                GuestName = vcsvm2
                DelayAfterGuestOnline = 7
                DelayAfterGuestOffline = 30
                )
```

Sample configuration for RHEV virtual machine migration:

```
group rhevgrp (
        SystemList = { sys1 = 0, sys2 = 1 }
        )

        KVMGuest rhevres (
                RHEVMInfo = { Enabled = 1,
                        URL = "https://rhevm.example.com:443",
                        User = "admin@internal",
                        Password = AMBmEMnMJmOGbGCgD,
                        Cluster = rhev_cluster,
                        UseManualRHEVMFencing=1 }
                GuestName = rhevvm1
```

```
DelayAfterGuestOnline = 15
DelayAfterGuestOffline = 45
)
```

Appendix C

# Where to find more information

This appendix includes the following topics:



## Symantec Storage Foundation and High Availability Solutions product documentation

Symantec Storage Foundation and High Availability Solutions product documentation is available in the Adobe Portable Document Format (PDF) on the product discs or with the downloaded software.

See the release notes for information on documentation changes in this release.

The documentation is available in the `/docs` directory.

Make sure that you are using the current version of documentation. The document version appears on page 2 of each guide. The publication date appears on the title page of each document. Symantec updates the product documents periodically for any errors or corrections. The latest product documentation is available on the Symantec website.

https://sort.symantec.com/documents

# Linux virtualization documentation

For Red Hat documentation:

- Red Hat Enterprise Linux (RHEL):
  https://access.redhat.com/site/documentation/Red_Hat_Enterprise_Linux/
- Red Hat Enterprise Virtualization (RHEV):
  https://access.redhat.com/site/documentation/Red_Hat_Enterprise_Virtualization/
- KVM Whitepaper:
  http://www.redhat.com/resourcelibrary/whitepapers/doc-kvm
- KVM Open source Project Site:
  http://www.linux-kvm.org/page/Main_Page

For SUSE:

- SUSE Linux Enterprise Server (SLES):
  http://www.suse.com/documentation/sles11/book_kvm/?page=/documentation/sles11/book_kvm/data/book_kvm.html
- For SLES11SP2 installation information:
  http://www.suse.com/documentation/sles11
  For a full set of features and capabilities, see the SUSE documentation.

# Service and support

To access the self-service knowledge base, go to the following URL:

http://entsupport.symantec.com

# About Symantec Operations Readiness Tools

Symantec Operations Readiness Tools (SORT) is a website that automates and simplifies some of the most time-consuming administrative tasks. SORT helps you manage your datacenter more efficiently and get the most out of your Symantec products.

SORT can help you do the following:

| | |
|---|---|
| Prepare for your next installation or upgrade | ■ List product installation and upgrade requirements, including operating system versions, memory, disk space, and architecture.<br>■ Analyze systems to determine if they are ready to install or upgrade Symantec products and generate an Installation and Upgrade custom report.<br>■ List patches by product or platform, and in the order they need to be installed. Display and download the most recent patches or historical patches.<br>■ Display Array Support Library (ASL) details by vendor, platform, or Storage Foundation and High Availability (SFHA) version. ASLs make it easier to manage arrays that are connected to SFHA-based servers.<br>■ List VCS and ApplicationHA agents, documentation, and downloads based on the agent type, application, and platform. |
| Identify risks and get server-specific recommendations | ■ Analyze your servers for potential environmental risks. Generate a Risk Assessment custom report with specific recommendations about system availability, storage use, performance, and best practices.<br>■ Display descriptions and solutions for thousands of Symantec error codes. |
| Improve efficiency | ■ Get automatic email notifications about changes to patches, array-specific modules (ASLs/APMs/DDIs/DDLs), documentation, product releases, Hardware Compatibility Lists (HCLs), and VCS/ApplicationHA agents.<br>■ Quickly gather installed Symantec product and license key information from across your production environment. Generate a License/Deployment custom report that includes product names, versions, and platforms, server tiers, Symantec Performance Value Units (SPVUs), and End of Service Life dates.<br>■ List and download Symantec product documentation including product guides, manual pages, compatibility lists, and support articles.<br>■ Access links to important resources on a single page, including Symantec product support, SymConnect forums, customer care, Symantec training and education, Symantec FileConnect, the licensing portal, and my.symantec.com. The page also includes links to key vendor support sites.<br>■ Use a subset of SORT features from your iOS device. Download the application at: https://sort.symantec.com/mobile |

**Note:** Certain features of SORT are not available for all products. Access to SORT is available at no extra cost.

To access SORT, go to:

https://sort.symantec.com