



Red Hat Virtualization 4.1

Disaster Recovery Guide

Configure Red Hat Virtualization 4.1 for Disaster Recovery

Red Hat Virtualization 4.1 Disaster Recovery Guide

Configure Red Hat Virtualization 4.1 for Disaster Recovery

Red Hat Virtualization Documentation Team
Red Hat Customer Content Services
rhev-docs@redhat.com

Legal Notice

Copyright © 2018 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux ® is the registered trademark of Linus Torvalds in the United States and other countries.

Java ® is a registered trademark of Oracle and/or its affiliates.

XFS ® is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL ® is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js ® is an official trademark of Joyent. Red Hat Software Collections is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack ® Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

Red Hat Virtualization can be configured to ensure that the environment remains operational even in the event of a catastrophe. This document provides information and instructions to configure Red Hat Virtualization environments for Disaster Recovery.

Table of Contents

CHAPTER 1. ACTIVE-ACTIVE DISASTER RECOVERY	3
1.1. OVERVIEW	3
1.1.1. Network Considerations	4
1.1.2. Storage Considerations	4
1.2. CONFIGURE A SELF-HOSTED ENGINE STRETCH CLUSTER ENVIRONMENT	5
1.3. CONFIGURE A STANDALONE MANAGER STRETCH CLUSTER ENVIRONMENT	6

CHAPTER 1. ACTIVE-ACTIVE DISASTER RECOVERY

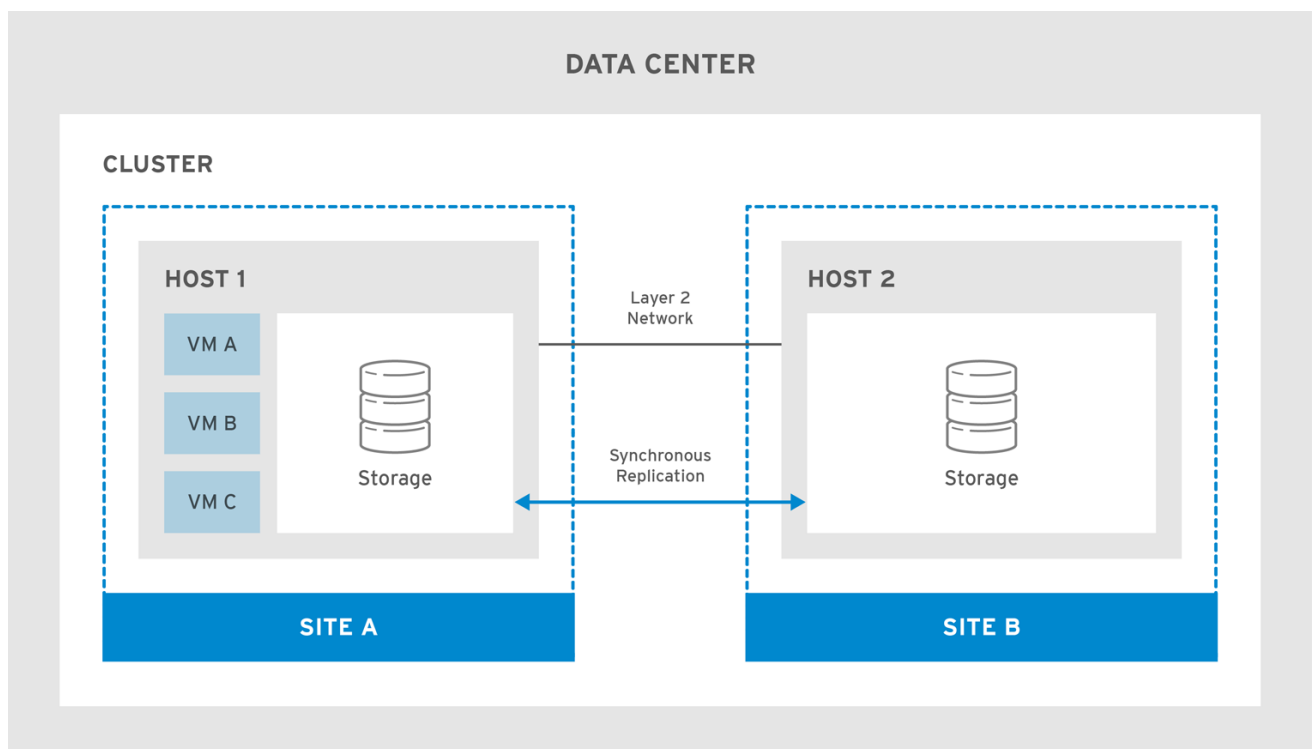
This chapter provides instructions to configure Red Hat Virtualization for disaster recovery using an active-active setup.

1.1. OVERVIEW

Red Hat Virtualization supports an active-active disaster recovery failover configuration that can span two sites. Both sites are active, and if the primary site becomes unavailable, the Red Hat Virtualization environment will continue to operate in the backup site to ensure business continuity.


The active-active failover is achieved by configuring a stretch cluster where hosts capable of running the virtual machines are located in the primary and secondary site. All the hosts belong to the same Red Hat Virtualization cluster.

You require replicated storage that is writeable on both sites to allow virtual machines to migrate between sites and continue running on the site's storage.



RHV_460251_1017

Virtual machines will migrate to the secondary site if the primary site becomes unavailable. The virtual machines will automatically failback to the primary site when the site becomes available and the storage is replicated in both sites.



RHV_460251_1017

IMPORTANT

To ensure virtual machine failover and failback works:

- Virtual machines must be configured to be highly available, and each virtual machine must have a lease on a target storage domain to ensure the virtual machine can start even without power management.
- Soft enforced virtual machine to host affinity must be configured to ensure the virtual machines only start on the selected hosts.

For more information see [Improving Uptime with Virtual Machine High Availability](#) and [Affinity Groups](#) in the *Virtual Machine Management Guide*.

The stretched cluster configuration can be implemented using a self-hosted engine environment, or a standalone Manager environment. For more information about the different types of deployments see [Architecture](#) in the *Product Guide*.

1.1.1. Network Considerations

All hosts in the cluster must be on the same broadcast domain over an L2 network. This means that connectivity between the two sites needs to be L2.

The maximum latency requirements between the sites across the L2 network differs for the two setups. The standalone Manager environment requires a maximum latency of 100ms, while the self-hosted engine environment requires a maximum latency of 7ms.

1.1.2. Storage Considerations

The storage domain for Red Hat Virtualization can be made of either block devices (SAN - iSCSI or FCP) or a file system (NAS - NFS, GlusterFS, or other POSIX compliant file systems). For more information about Red Hat Virtualization storage see [Storage](#) in the *Administration Guide*.

The sites require synchronously replicated storage that is writeable on both sites with shared layer 2 (L2) network connectivity. The replicated storage is required to allow virtual machines to migrate between sites and continue running on the site's storage. All storage replication options supported by Red Hat Enterprise Linux 7 and later can be used in the stretch cluster.



IMPORTANT

If you have a custom multipath configuration that is recommended by the storage vendor, copy the `.conf` file to the `/etc/multipath/conf.d/` directory. The custom settings will override settings in the VDSMs `multipath.conf` file. Do not modify the VDSM file directly.

1.2. CONFIGURE A SELF-HOSTED ENGINE STRETCH CLUSTER ENVIRONMENT

This procedure provides instructions to configure a stretch cluster using a self-hosted engine deployment.

Prerequisites:

- A writable storage server in both sites with L2 network connectivity.
- Real-time storage replication service to duplicate the storage.

Limitations:

- Maximum 7ms latency between sites.

Configuring the Self-hosted Engine Stretch Cluster

1. Deploy the self-hosted engine. See [Deploying Self-Hosted Engine](#) in the *Self-hosted Engine Guide*.
2. Install additional self-hosted engine nodes in each site and add them to your cluster. See [Installing Additional Hosts to a Self-Hosted Environment](#) in the *Self-hosted Engine Guide*.
3. Optional. Install additional standard hosts. See [Adding a Host to the Red hat Virtualization Manager](#) in the *Installation Guide*.
4. Configure the SPM priority to be higher on all hosts in the primary site to ensure SPM failover to the secondary site occurs only when all hosts in the primary site are unavailable. See [SPM Priority](#) in the *Administration Guide*.
5. Configure all virtual machines that need to failover as highly available, and ensure that the virtual machine has a lease on the target storage domain. See [Configuring a Highly Available Virtual Machine](#) in the *Virtual Machine Management Guide*.
6. Configure virtual machine to host soft affinity and define the behavior you expect from the affinity group. See [Affinity Groups](#) in the *Virtual Machine Management Guide* and [Scheduling Policies](#) in the *Administration Guide*.

The active-active failover can be manually performed by placing the main site's hosts into maintenance mode.

1.3. CONFIGURE A STANDALONE MANAGER STRETCH CLUSTER ENVIRONMENT

This procedure provides instructions to configure a stretch cluster using a standalone Manager deployment.

Prerequisites:

- A writable storage server in both sites with L2 network connectivity.
- Real-time storage replication service to duplicate the storage.

Limitations:

- Maximum 100ms latency between sites.



IMPORTANT

The Manager must be highly available for virtual machines to failover and failback between sites. If the Manager goes down with the site, the virtual machines will not failover.

The standalone Manager is only highly available when managed externally. For example:

- Using Red Hat's High Availability Add-On.
- As a highly available virtual machine in a separate virtualization environment.
- Using Red Hat Enterprise Linux Cluster Suite.
- In a public cloud.

Configuring the Standalone Manager Stretch Cluster

1. Install and configure the Red Hat Virtualization Manager. See [Installing the Red Hat Virtualization Manager](#) in the *Installation Guide*.
2. Install the hosts in each site and add them to the cluster. See [Installing Hosts](#) and [Adding a Host to the Red hat Virtualization Manager](#) in the *Installation Guide*.
3. Configure the SPM priority to be higher on all hosts in the primary site to ensure SPM failover to the secondary site occurs only when all hosts in the primary site are unavailable. See [SPM Priority](#) in the *Administration Guide*.
4. Configure all virtual machines that need to failover as highly available, and ensure that the virtual machine has a lease on the target storage domain. See [Configuring a Highly Available Virtual Machine](#) in the *Virtual Machine Management Guide*.
5. Configure virtual machine to host soft affinity and define the behavior you expect from the affinity group. See [Affinity Groups](#) in the *Virtual Machine Management Guide* and [Scheduling Policies](#) in the *Administration Guide*.

The active-active failover can be manually performed by placing the main site's hosts into maintenance mode.

