

vRealize Suite 6.0 Disaster Recovery by Using Site Recovery Manager 6.1

vRealize Suite 6.0

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see <http://www.vmware.com/support/pubs>.

EN-002062-00

vmware[®]

You can find the most up-to-date technical documentation on the VMware Web site at:

<http://www.vmware.com/support/>

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com

Copyright © 2016 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

	vRealize Suite 6.0 Disaster Recovery by Using Site Recovery Manager 6.1	5
1	Disaster Recovery Introduction	7
	Overview of VMware Site Recovery Manager	8
	Site Recovery Manager Workflow	11
2	Types of Replication Technologies	15
	Array-Based Replication Versus vSphere Replication	15
	Using Array-Based Replication with Site Recovery Manager	16
	Using vSphere Replication with Site Recovery Manager	17
	Using Array-Based Replication and vSphere Replication with Site Recovery Manager	18
3	Configuring Site Recovery Manager	21
	Configure Virtual Machines for vSphere Replication	21
	Configure Virtual Machines for Array Based Replication	22
	Create Protection Groups	22
	Create a Recovery Plan	24
	Edit a Recovery Plan	24
4	Configuring vRealize Suite Components for Disaster Recovery	27
	vRealize Automation Disaster Recovery	27
	vRealize Orchestrator Disaster Recovery	39
	vRealize Operations Manager Disaster Recovery	39
	vRealize Log Insight Disaster Recovery	45
5	Testing and Executing a Recovery Plan	53
	Test a Recovery Plan	53
	Clean Up After Testing a Recovery Plan	54
	Run a Recovery Plan	54
	Cancel a Test or Recovery	56
6	Perform a Failback	57
	Configuring vRealize Suite Components Post Failback	58
	Index	59

vRealize Suite 6.0 Disaster Recovery by Using Site Recovery Manager 6.1

The *vRealize Suite 6.0 Disaster Recovery by Using Site Recovery Manager 6.1* provides information about how to protect your vRealize Suite components by using Site Recovery Manager, which is disaster recovery automation software that provides policy-based management, non-disruptive testing, and automated orchestration.

To protect your vRealize Suite 6.0 components, Site Recovery Manager 6.1 automates every aspect of executing a disaster recovery plan to accelerate recovery.

Intended Audience

This information is intended for anyone who wants to implement Site Recovery Manager to protect the vRealize Suite components from a disaster. This information is written for experienced Windows or Linux system administrators who are familiar with virtual machine technology and datacenter operations.

VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to <http://www.vmware.com/support/pubs>.

Disaster Recovery Introduction

vRealize Suite 6.0 Disaster Recovery by Using Site Recovery Manager 6.1 describes how to implement and use Site Recovery Manager with vSphere Replication or array-based replication to protect the vRealize Suite components.

The following vRealize Suite 6.0 components were tested for disaster recovery by using Site Recovery Manager 6.1:

- vRealize Automation 6.2.3
 - CloudClient 3.4.1
- vRealize Orchestrator 6.0.2
- vRealize Operations Manager 6.1.0

NOTE vRealize Operations Manager disaster recovery was tested completely with vSphere Replication. However, vRealize Operations Manager disaster recovery was tested with the array-based replication in a limited manner, only by using EMC RecoverPoint.

- vRealize Log Insight 2.5 and 3.0

The following products were used for implementing disaster recovery for vRealize Suite 6.0:

- VMware vSphere Replication 6.0
- Array-based replication
 - EMC RecoverPoint Storage Replication Adapter
 - vCenter 6.0

For information about installing, upgrading, and configuring VMware Site Recovery Manager, see the [VMware Site Recovery Manager 6.1 Documentation](#).

To protect the vRealize Suite components, you must configure the Site Recovery Manager tool and the vRealize Suite components in the following order:

- 1 [Chapter 3, “Configuring Site Recovery Manager,”](#) on page 21. You must configure Site Recovery Manager by following the guidelines to protect your vRealize Suite components.
- 2 [Chapter 4, “Configuring vRealize Suite Components for Disaster Recovery,”](#) on page 27. You must configure the vRealize Suite components that you want to protect by using Site Recovery Manager.
- 3 [Chapter 5, “Testing and Executing a Recovery Plan,”](#) on page 53. You must test the recovery plans that you have created to verify that the recovery is successfully completed without any data loss, before a disaster recovery situation occurs.

This chapter includes the following topics:

- [“Overview of VMware Site Recovery Manager,”](#) on page 8
- [“Site Recovery Manager Workflow,”](#) on page 11

Overview of VMware Site Recovery Manager

VMware Site Recovery Manager is a business continuity and disaster recovery solution that helps you to plan, test, and run the recovery of virtual machines between a protected vCenter Server site and a recovery vCenter Server site.

You can configure Site Recovery Manager to protect virtual machines in different ways.

Datastore groups	Protect the virtual machines in datastore groups by using third-party disk replication mechanisms to configure array-based replication. Array-based replication surfaces replicated datastores to recover virtual machine workloads.
Individual virtual machines	Protect the individual virtual machines on a host by using Site Recovery Manager in combination with VMware vSphere Replication.
Storage policies	Protect virtual machines based on their association with specific storage policies. Protecting virtual machines by using storage policies requires array-based replication.

You can use Site Recovery Manager to implement different types of recovery from the protected site to the recovery site.

Planned migration	The orderly evacuation of virtual machines from the protected site to the recovery site. Planned migration prevents data loss when migrating workloads in an orderly fashion. For planned migration to succeed, both sites must be running and fully functioning.
Disaster recovery	Similar to planned migration except that disaster recovery does not require that both sites be up and running, for example if the protected site goes offline unexpectedly. During a disaster recovery operation, failure of operations on the protected site is reported but is otherwise ignored.

Site Recovery Manager orchestrates the recovery process with the replication mechanisms, to minimize data loss and system down time.

- At the protected site, Site Recovery Manager shuts down virtual machines cleanly and synchronizes storage, if the protected site is still running.
- Site Recovery Manager powers on the replicated virtual machines at the recovery site according to a recovery plan.

A recovery plan specifies the order in which virtual machines start up on the recovery site. A recovery plan specifies network parameters, such as IP addresses, and can contain user-specified scripts that Site Recovery Manager can run to perform custom recovery actions on virtual machines.

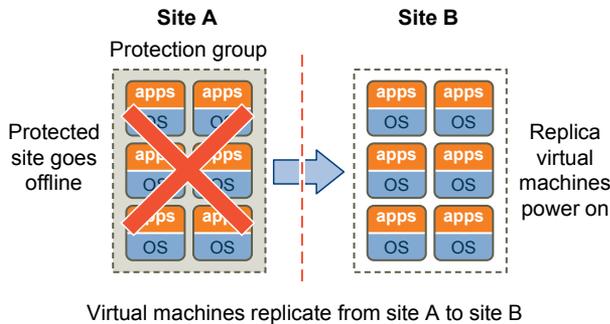
Site Recovery Manager lets you test recovery plans. You conduct tests by using a temporary copy of the replicated data in a way that does not disrupt ongoing operations at either site.

About Protected Sites and Recovery Sites

In a typical Site Recovery Manager installation, the protected site provides business-critical datacenter services. The recovery site is an alternative infrastructure to which Site Recovery Manager can migrate these services.

The protected site can be any site where vCenter Server supports a critical business need. The recovery site can be located thousands of miles away from the protected site. Conversely, the recovery site can be in the same room as a way of establishing redundancy. The recovery site is usually located in a facility that is unlikely to be affected by environmental, infrastructure, or other disturbances that affect the protected site.

Figure 1-1. Site Recovery Manager Protected and Recovery Sites



The vSphere configurations at each site must meet requirements for Site Recovery Manager.

- You must run the same version of Site Recovery Manager on both sites.
- You must run the same version of vCenter Server on both sites.
- The version of vCenter Server must be compatible with the version of Site Recovery Manager. For information about compatibility between vCenter Server and Site Recovery Manager versions, see *vCenter Server Requirements* in the *Compatibility Matrixes for Site Recovery Manager 6.1* at <https://www.vmware.com/support/srm/srm-compat-matrix-6-1.html>.
- Each site must have at least one datacenter.
- If you are using array-based replication, the same replication technology must be available at both sites, and the arrays must be paired.
- If you are using vSphere Replication, you require a vSphere Replication appliance on both sites. The vSphere Replication appliances must be connected to each other.
- The vSphere Replication appliances must be of the same version.
- The vSphere Replication version must be compatible with the version of Site Recovery Manager. For information about compatibility between vSphere Replication and Site Recovery Manager versions, see *vSphere Replication Requirements* in the *Compatibility Matrixes for Site Recovery Manager 6.1* at <https://www.vmware.com/support/srm/srm-compat-matrix-6-1.html>.
- The recovery site must have hardware, network, and storage resources that can support the same virtual machines and workloads as the protected site. You can oversubscribe the recovery site by running additional virtual machines there that are not protected. In this case, during a recovery you must suspend noncritical virtual machines on the recovery site.
- The sites must be connected by a reliable IP network. If you are using array-based replication, ensure that your network connectivity meets the arrays' network requirements.
- The recovery site should have access to comparable public and private networks as the protected site, although not necessarily the same range of network addresses.

Heterogeneous Configurations on the Protected and Recovery Sites

Some components in the Site Recovery Manager and vCenter Server installations must be identical on each site. Because the protected and recovery sites are often in different physical locations, some components on the protected site can be of a different type to their counterparts on the recovery site.

Although components can be different on each site, you must use the types and versions of these components that Site Recovery Manager supports. See the *Compatibility Matrixes for Site Recovery Manager 6.1* at <https://www.vmware.com/support/srm/srm-compat-matrix-6-1.html> for information.

Table 1-1. Heterogeneity of Site Recovery Manager Components Between Sites

Component	Heterogeneous or Identical Installations
Site Recovery Manager Server	Must be the same version on both sites.
vCenter Server and Platform Services Controller	Must be the same version on both sites. The Site Recovery Manager version must be compatible with the vCenter Server and Platform Services Controller version.
vSphere Replication	Must be the same version on both sites. The vSphere Replication version must be compatible with the Site Recovery Manager version and the vCenter Server version.
vCenter Server Appliance or standard vCenter Server instance	Can be different on each site. You can run a vCenter Server Appliance on one site and a standard vCenter Server instance on the other site.
Storage arrays for array-based replication	Can be different versions on each site. You can use different versions of the same type of storage array on each site. The Site Recovery Manager Server instance on each site requires the appropriate storage replication adapter (SRA) for each version of storage array for that site. Check SRA compatibility with all versions of your storage arrays to ensure compatibility.
Site Recovery Manager database	Can be different on each site. You can use different versions of the same type of database on each site, or different types of database on each site.
Host operating system of the Site Recovery Manager Server installation	Can be different on each site. You can run different versions of the host operating system and the host operating system can run in different locales on each site.
Host operating system of the vCenter Server installation	Can be different on each site. You can run different versions of the host operating system and the host operating system can run in different locales on each site.

Example: Heterogenous Configurations on the Protected and Recovery Sites

The Site Recovery Manager and vCenter Server installations might be in different countries, with different setups.

- Site A in Japan:
 - Site Recovery Manager Server runs on Windows Server 2008 in the Japanese locale
 - Site Recovery Manager extends a vCenter Server Appliance instance
 - Site Recovery Manager Server uses the embedded Site Recovery Manager database
- Site B in the United States:
 - Site Recovery Manager Server runs on Windows Server 2012 in the English locale

- Site Recovery Manager extends a standard vCenter Server instance that runs on Windows Server 2008 in the English locale
- Site Recovery Manager Server uses an Oracle Server database

Site Recovery Manager Workflow

When you create or modify a recovery plan, test it before you use it for planned migration or for disaster recovery.

Testing a Recovery Plan

By testing a recovery plan, you ensure that the virtual machines that the plan protects recover correctly to the recovery site. If you do not test recovery plans, a disaster recovery situation might not recover all virtual machines, resulting in data loss.

If you use vSphere Replication, when you test a recovery plan, the virtual machine on the protected site can still synchronize with the replica virtual machine disk files at the recovery site. The vSphere Replication server creates redo logs on the virtual machine disk files at the recovery site, so that synchronization can continue normally. When you perform cleanup after running a test, the vSphere Replication server removes the redo logs from the disks at the recovery site and persists the changes accumulated in the logs to VM disks.

If you use array-based replication, when you test a recovery plan, the virtual machines on the protected site are still replicated to the replica virtual machines' disk files at the recovery site. During test recovery, the array creates a snapshot of the volumes hosting the virtual machines' disk files at the recovery site. Array replication continues normally while the test is in progress. When you perform cleanup after running a test, the array removes the snapshots that were created earlier as part of the test recovery workflow.

You can run a recovery plan test as often as necessary. You can cancel a recovery plan test at any time. Before running a failover or another test, you must successfully run a cleanup operation.

Performing a Planned Migration

You can run a recovery plan under planned circumstances to migrate virtual machines from the protected site to the recovery site. You can also run a recovery plan under unplanned circumstances if the protected site suffers an unforeseen event that might result in data loss.

During a planned migration, Site Recovery Manager synchronizes the virtual machine data at the recovery site with the virtual machines on the protected site. Site Recovery Manager attempts to gracefully shut down the protected machines and performs a final synchronization to prevent data loss, and powers on the virtual machines at the recovery site. If errors occur during a planned migration, the plan stops so that you can resolve the errors and rerun the plan. You can reprotect the virtual machines after the recovery.

After Site Recovery Manager completes the final replication, Site Recovery Manager makes changes at both sites that require significant time and effort to reverse. Because of this time and effort, you must assign the privilege to test a recovery plan and the privilege to run a recovery plan separately.

Performing a Disaster Recovery

During disaster recoveries, Site Recovery Manager first attempts a storage synchronization. If it succeeds, Site Recovery Manager uses the synchronized storage state to recover virtual machines at the recovery site to their most recent available state, according to the recovery point objective (RPO) that you set when you configure your replication technology.

When you run a recovery plan to perform a disaster recovery, Site Recovery Manager attempts to shut down the virtual machines on the protected site. If Site Recovery Manager cannot shut down the virtual machines, Site Recovery Manager still starts the copies at the recovery site. In case the protected site comes back online after disaster recovery, the recovery plan goes into an inconsistent state where production virtual machines are running on both sites, known as a split-brain scenario. Site Recovery Manager detects this state and allows you to run the plan once more to power off the virtual machines on the protected site. Then the recovery plan goes back to a consistent state and you can run reprotect.

Restoring the Pre-Recovery Site Configuration By Performing Failback

To restore the original configuration of the protected and recovery sites after a recovery, you can perform a sequence of optional procedures known as failback.

After a planned migration or a disaster recovery, the former recovery site becomes the protected site. Immediately after the recovery, the new protected site has no recovery site to which to recover. If you run reprotect, the new protected site is protected by the original protection site, reversing the original direction of protection. See [Reprotecting Virtual Machines After a Recovery](#) for information about reprotect.

To restore the configuration of the protected and recovery sites to their initial configuration before the recovery, you perform failback.

To perform failback, you run a sequence of reprotect and planned migration operations.

- 1 Perform a reprotect. The recovery site becomes the protected site. The former protected site becomes the recovery site.
- 2 Perform a planned migration to shut down the virtual machines on the protected site and start up the virtual machines on the recovery site. To avoid interruptions in virtual machine availability, you might want to run a test before you start the planned migration. If the test identifies errors, you can resolve them before you perform the planned migration.
- 3 Perform a second reprotect, to revert the protected and recovery sites to their original configuration before the recovery.

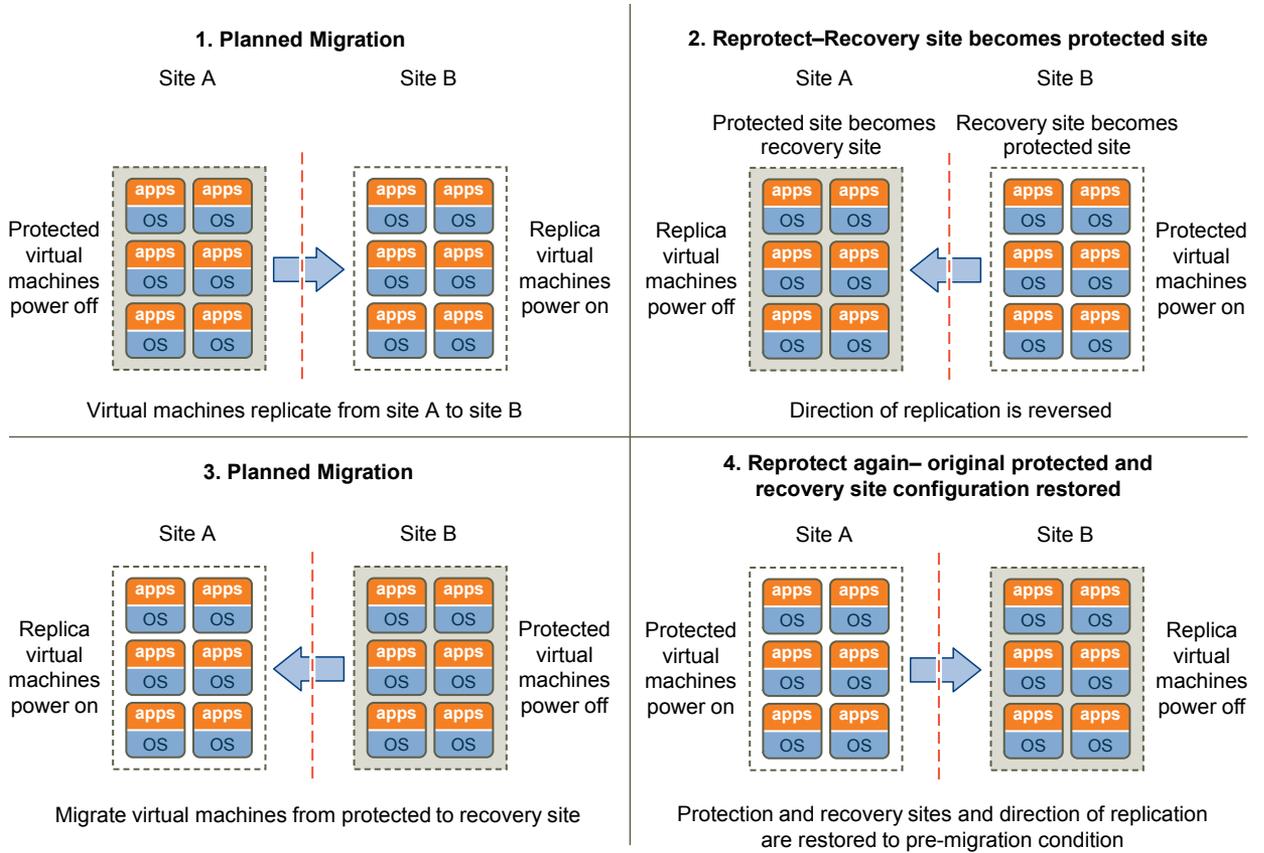
You can configure and run a failback when you are ready to restore services to the original protected site, after you have brought it back online after an incident.

Example: Performing a Failback Operation

Site A is the protected site and B is the recovery site. A recovery occurs, migrating the virtual machines from site A to site B. To restore site A as the protected site, you perform a failback.

- 1 Virtual machines replicate from site A to site B.
- 2 Perform a reprotect. Site B, the former recovery site, becomes the protected site. Site Recovery Manager uses the protection information to establish the protection of site B. Site A becomes the recovery site.
- 3 Perform a planned migration to recover the protected virtual machines on site B to site A.
- 4 Perform a second reprotect. Site A becomes the protected site and site B becomes the recovery site.

Figure 1-2. Site Recovery Manager Failback Process



Types of Replication Technologies

You must configure replication on the virtual machines that you want to protect. Site Recovery Manager supports two types of replication technologies, array-based replication and vSphere Replication.

This chapter includes the following topics:

- [“Array-Based Replication Versus vSphere Replication,”](#) on page 15
- [“Using Array-Based Replication with Site Recovery Manager,”](#) on page 16
- [“Using vSphere Replication with Site Recovery Manager,”](#) on page 17
- [“Using Array-Based Replication and vSphere Replication with Site Recovery Manager,”](#) on page 18

Array-Based Replication Versus vSphere Replication

You can replicate virtual machines by using either array-based replication, vSphere Replication, or a combination of both. You cannot use both the replication technologies to protect the same virtual machine.

The following table lists the differences between the two replication technologies and helps you to decide why you must use one technology rather than its alternative.

Characteristics	Array-Based Replication	vSphere Replication
Type	Replication using the storage layer	Replication using the host or vSphere layer
Recovery point objective min or max	0 up to max supported by vendor	5 minutes to 24 hours
Scale	Scales up to 5,000 VMs protected, 2,000 simultaneously recoverable through vCenter and Site Recovery Manager pairing	Scales up to 2000 VMs (protected and recoverable) through vCenter and Site Recovery Manager pairing
Write-order fidelity	Supports write-order fidelity within and across multiple VMs in the same consistency group	Supports write-order fidelity on the disks or VMDKs that comprise a VM. Consistency cannot be guaranteed across multiple VMs.
Replication level	Replicates at the LUN or VMFS,, or NFS volume level	Replicates at the VM level
Replication configuration	Replication is configured and managed on the storage array.	Replication is configured and managed in the vSphere Web Client.
Array and vendor types	Requires same storage replication solution at both sites, for example, EMC RecoverPoint, NetApp vFiler, IBM SVC	Can support any storage solution at either end including local storage if it is covered by the vSphere HCL
Storage supported	Replication supported on FC, sCSI, or NFS storage only	Supports replicating VMs on local, attached, Virtual SAN, FC, sCSI, or NFS storage.

Characteristics	Array-Based Replication	vSphere Replication
Cost	Replication and snapshot licensing is required.	vSphere Replication is included in vSphere Essentials Plus Kit version 5.1 and later.
Deployment	Deployment is fairly involved and must include storage administration and possibly networking.	Deployment requirements are minimal: OVF at each site and start configuring replications.
Application consistency	Depending on the array, application consistency might be supported with the addition of agents to the VM.	Supports Volume Shadow Copy Service and Linux file system application consistency
Fault Tolerance (FT) VMs	Can replicate UP FT-protected VMs. After a VM is recovered, it is no longer FT enabled. Does not support FT VMs that are configured for SMP.	Cannot replicate FT-protected VMs
Powered off VMs, templates, linked clones, or ISOs	Can replicate powered off VMs, templates, and linked clones, if all nodes in the snapshot tree are also replicated, and ISOs	Can replicate only powered on VMs. Cannot replicate powered off VMs, templates, linked clones, ISOs, or any non-VM files.
Raw device mapping (RDM) support	Physical and virtual mode RDMs can be replicated.	Only virtual mode RDMs can be replicated.
Microsoft Cluster Service (MSCS) support	VMs that are part of a MSCS cluster can be replicated.	Cannot replicate VMs that are part of an MSCS cluster. vSphere Replication cannot replicate disks in multiwriter mode.
vApp support	Replicating vApps is supported.	You cannot replicate vApps. You can replicate VMs that are part of a vApp and create a vApp at the recovery site that they are recovered to.
vSphere versions supported	Hosts running vSphere 3.5 through 6.x are supported.	Hosts must be running vSphere 5.0 or later.
Multiple point-in-time (MPIT) snapshots	MPIT snapshots or rollback is supported by some supported array vendors, for example, EMC RecoverPoint.	Supports up to 24 recovery points
Snapshots	Supports replicating VMs with snapshots and maintaining the snapshot tree	Supports replicating VMs with snapshots. However, the tree is collapsed at the target site.
Response to host failure	Replication is not impacted.	Host failure and the VM restarting on another host trigger a full sync.

Using Array-Based Replication with Site Recovery Manager

When you use array-based replication, one or more storage arrays at the protected site replicate data to peer arrays at the recovery site. With storage replication adapters (SRAs), you can integrate Site Recovery Manager with a wide variety of arrays.

To use array-based replication with Site Recovery Manager, you must configure replication first before you can configure Site Recovery Manager to use it.

If your storage array supports consistency groups, Site Recovery Manager is compatible with vSphere Storage DRS and vSphere Storage vMotion. You can use Storage DRS and Storage vMotion to move virtual machine files within a consistency group that Site Recovery Manager protects. If your storage array does not support consistency groups, you cannot use Storage DRS and Storage vMotion in combination with Site Recovery Manager.

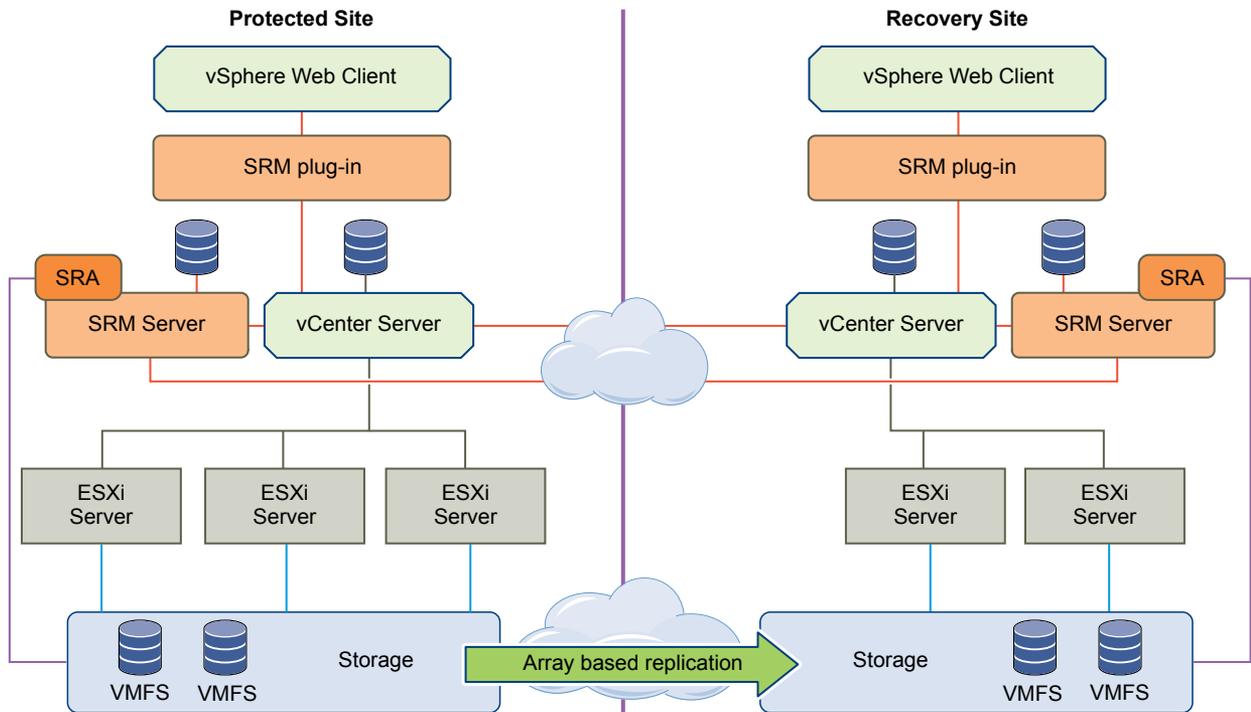
You can protect virtual machines that contain disks that use VMware vSphere Flash Read Cache storage. Since the host to which a virtual machine recovers might not be configured for Flash Read Cache, Site Recovery Manager disables Flash Read Cache on disks when it starts the virtual machines on the recovery site. Site Recovery Manager sets the reservation to zero. Before performing a recovery on a virtual machine that is configured to use vSphere Flash Read Cache, take a note of virtual machine's cache reservation from the vSphere Web Client. After the recovery, you can migrate the virtual machine to a host with Flash Read Cache storage and manually restore the original Flash Read Cache setting on the virtual machine.

If you protect virtual machines by using storage policy protection groups, you must replicate those virtual machines by using array-based replication.

Storage Replication Adapters

Storage replication adapters are not part of a Site Recovery Manager release. Your array vendor develops and supports them. You must install an SRA specific to each array that you use with Site Recovery Manager on the Site Recovery Manager Server host. Site Recovery Manager supports the use of multiple SRAs.

Figure 2-1. Site Recovery Manager Architecture with Array-Based Replication

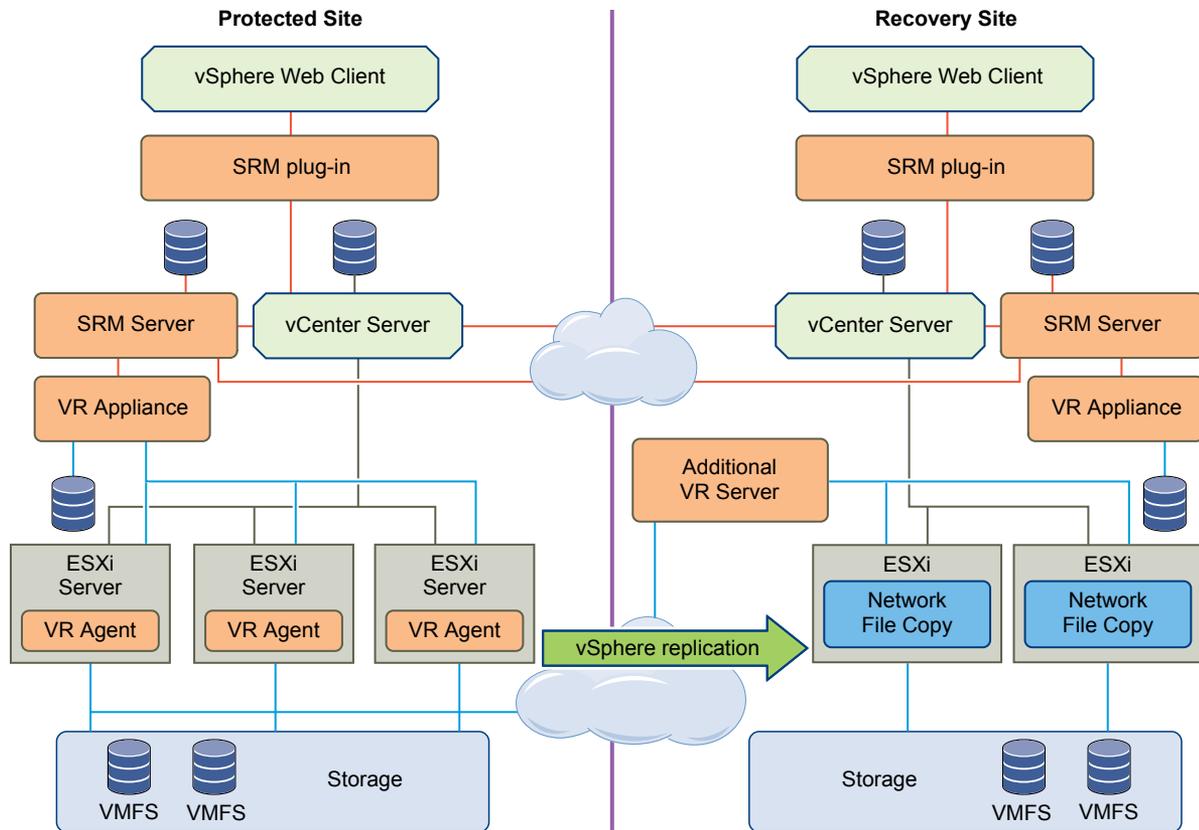


Using vSphere Replication with Site Recovery Manager

Site Recovery Manager can use vSphere Replication to replicate data to servers at the recovery site.

vSphere Replication does not require storage arrays. The vSphere Replication storage replication source and target can be any storage device, including, but not limited to, storage arrays.

You can configure vSphere Replication to regularly create and retain snapshots of protected virtual machines on the recovery site. Taking multiple point-in-time (PIT) snapshots of virtual machines allows you to retain more than one replica of a virtual machine on the recovery site. Each snapshot reflects the state of the virtual machine at a certain point in time. You can select which snapshot to recover when you use vSphere Replication to perform a recovery.

Figure 2-2. Site Recovery Manager Architecture with vSphere Replication

Using vSphere Replication and Site Recovery Manager with vSphere Storage vMotion and vSphere Storage DRS

vSphere Replication is compatible with vSphere Storage vMotion and vSphere Storage DRS on the protected site. You can use Storage vMotion and Storage DRS to move the disk files of a virtual machine that vSphere Replication protects, with no impact on replication.

Using vSphere Replication and VMware Virtual SAN Storage with Site Recovery Manager

You can use VMware Virtual SAN storage with vSphere Replication and Site Recovery Manager.

Using Array-Based Replication and vSphere Replication with Site Recovery Manager

You can use a combination of array-based replication and vSphere Replication in your Site Recovery Manager deployment.

To create a mixed Site Recovery Manager deployment that uses array-based replication and vSphere Replication, you must configure the protected and recovery sites for both types of replication.

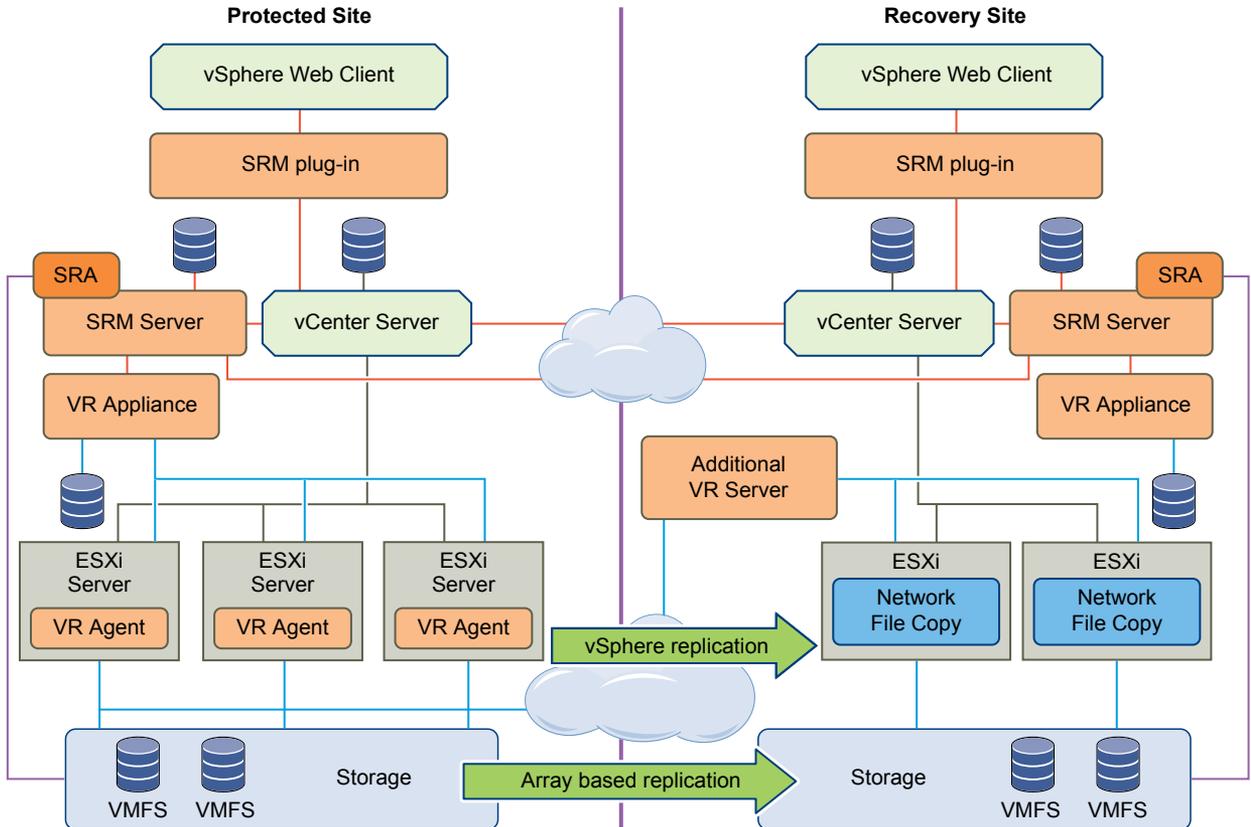
- Set up and connect the storage arrays and install the appropriate storage replication adapters (SRA) on both sites.
- Deploy vSphere Replication appliances on both sites and configure the connection between the appliances.

- Configure virtual machines for replication using either array-based replication or vSphere Replication, as appropriate.

NOTE Do not attempt to configure vSphere Replication on a virtual machine that resides on a datastore that you replicate by using array-based replication.

You create array-based protection groups for virtual machines that you configure with array-based replication, and vSphere Replication protection groups for virtual machines that you configure with vSphere Replication. You cannot mix replication types in a protection group. You can mix array-based protection groups and vSphere Replication protection groups in the same recovery plan.

Figure 2-3. Site Recovery Manager Architecture with Array-Based Replication and vSphere Replication



Configuring Site Recovery Manager

You must configure Site Recovery Manager to protect your vRealize Suite components. Secure this protection by completing the common configuration tasks for Site Recovery Manager.

Prepare the Environment

You must ensure that you meet the following prerequisites before you start configuring Site Recovery Manager.

- Verify that vSphere 6.0 is installed on the protected and recovery sites.
- Verify that you are using Site Recovery Manager 6.1.
- Verify that your vRealize Suite components are in a ready state.

This chapter includes the following topics:

- [“Configure Virtual Machines for vSphere Replication,”](#) on page 21
- [“Configure Virtual Machines for Array Based Replication,”](#) on page 22
- [“Create Protection Groups,”](#) on page 22
- [“Create a Recovery Plan,”](#) on page 24
- [“Edit a Recovery Plan,”](#) on page 24

Configure Virtual Machines for vSphere Replication

You must configure the virtual machines for vSphere Replication or array based replication in order to use Site Recovery Manager.

To enable vSphere Replication on the required virtual machines, perform the following steps.

Procedure

- 1 In the vSphere Web Client, select a virtual machine on which vSphere Replication should be enabled and click **Actions > All vSphere Replication Actions > Configure Replication**.
- 2 In the Replication type window, select **Replicate to a vCenter Server** and click **Next**.
- 3 In the Target site window, select the vCenter for the recovery site and click **Next**.
- 4 In the Replication server window, select a vSphere Replication server and click **Next**.
- 5 In the Target location window, click **Edit** and select the target datastore, where the replicated files will be stored and click **Next**.
- 6 In the Replication options window, keep the default setting and click **Next**.

- 7 In the Recovery settings window, enter time for **Recovery Point Objective (RPO)** and **Point in time instances**, and click **Next**.
- 8 In the Ready to complete window, verify the settings and click **Finish**.
- 9 Repeat these steps for all virtual machines on which vSphere Replication must be enabled.

Configure Virtual Machines for Array Based Replication

You must configure the virtual machines for vSphere Replication or array based replication in order to use Site Recovery Manager.

Ensure that the virtual machines are located on the same datastore drive array configured in the array based replication protection group. To enable array based replication on the required virtual machines, perform the following steps.

Procedure

- 1 In the vSphere Web Client, click **Site Recovery > Protection Groups**.
- 2 Select the **Array Based Replication** protection group.
- 3 On the **Related Objects** tab, click **Virtual Machines**.
- 4 Right-click the virtual machine and select **Configure Protection**, if the status of the virtual machine is **Not Configured** or **Mapping Missing**.
- 5 Selected the required options on the **VM Protection Properties** window and click **OK**.
- 6 To apply these mappings to all protected virtual machines on the site, select the **Save as Inventory Mapping** check box for each resource.
If you do not select the check box, the mapping is only applied to the selected virtual machine.
- 7 Click **OK**.
 - If Site Recovery Manager successfully applied inventory mappings to the virtual machine, the status of the virtual machine is **OK**.
 - If Site Recovery Manager was unable to apply some or all of the inventory mappings, the status of the virtual machine is **Not Configured** or **Mapping Missing**.

Create Protection Groups

You create protection groups to enable Site Recovery Manager to protect virtual machines.

You can organize protection groups in folders. Different views in the vSphere Web Client display the names of the protection groups, but they do not display the folder names. If you have two protection groups with the same name in different folders, it might be difficult to tell them apart in some views in the vSphere Web Client. Consequently, ensure that protection group names are unique across all folders. In environments in which not all users have view privileges for all folders, to be sure of the uniqueness of protection group names, do not place protection groups in folders.

When you create protection groups, wait to ensure that the operations finish as expected. Make sure that Site Recovery Manager creates the protection group and that the protection of the virtual machines in the group is successful.

Prerequisites

Verify that you performed one of the following tasks:

- Included virtual machines in datastores for which you configured array-based replication
- Satisfied the requirements in [Prerequisites for Storage Policy Protection Group](#) and reviewed the [Limitations of Storage Policy Protection Groups](#).

- Configured vSphere Replication on virtual machines
- Performed a combination of some or all of the above

Procedure

- 1 In the vSphere Web Client, click **Site Recovery > Protection Groups**.
- 2 On the **Objects** tab, click the icon to create a protection group.
- 3 On the Name and location page, enter a name and description for the protection group, select a pair of sites or a folder, and click **Next**.
- 4 On the Protection group type page, select the direction of protection and the protection group type, and click **Next**.

Option	Action
Create an array-based replication protection group	Select Datastore groups (array-based replication) and select an array pair.
Create a vSphere Replication protection group	Select Individual VMs (vSphere Replication) .
Create a storage policy protection group	Select Storage Policies (array-based replication) .

- 5 Select datastore groups, virtual machines, or storage policies to add to the protection group.

Option	Action
Array-based replication protection groups	Select datastore groups and click Next . When you select a datastore group, the virtual machines that the group contains appear in the Virtual machines table.
vSphere Replication protection groups	Select virtual machines from the list, and click Next . Only virtual machines that you configured for vSphere Replication and that are not already in a protection group appear in the list.
Storage policy protection groups	Select storage policies from the list, and click Next .

- 6 Review your settings and click **Finish**.

You can monitor the progress of the creation of the protection group on the **Objects** tab.

- For array-based replication and vSphere Replication protection groups, if Site Recovery Manager successfully applied inventory mappings to the protected virtual machines, the protection status of the protection group is OK.
- For storage policy protection groups, if Site Recovery Manager successfully protected all of the virtual machines associated with the storage policy, the protection status of the protection group is OK.
- For array-based replication and vSphere Replication protection groups, if you did not configure inventory mappings, or if Site Recovery Manager was unable to apply them, the protection status of the protection group is Not Configured.
- For storage policy protection groups, if Site Recovery Manager could not protect all of the virtual machines associated with the storage policy, the protection status of the protection group is Not Configured.

What to do next

For array-based replication and vSphere Replication protection groups, if the protection status of the protection group is Not Configured, apply inventory mappings to the virtual machines:

- To apply site-wide inventory mappings, or to check that inventory mappings that you have already set are valid, see [Select Inventory Mappings](#) in *Site Recovery Manager Installation and Configuration*. To apply these mappings to all of the virtual machines, see [Apply Inventory Mappings to All Members of a Protection Group](#).
- To apply inventory mappings to each virtual machine in the protection group individually, see [Configure Inventory Mappings for an Individual Virtual Machine in a Protection Group](#).

For storage policy protection groups, if the protection status of the protection group is Not Configured, verify that you satisfied the prerequisites in [Prerequisites for Storage Policy Protection Group](#), review [Limitations of Storage Policy Protection Groups](#), modify the storage policy implementation accordingly, and attempt to recreate the protection group.

Create a Recovery Plan

You create a recovery plan to establish how Site Recovery Manager recovers virtual machines.

Procedure

- 1 In the vSphere Web Client, select **Site Recovery > Recovery Plans**.
- 2 On the **Objects** tab, click the icon to create a recovery plan.
- 3 Enter a name and description for the plan, select a folder, then click **Next**.
- 4 Select the recovery site and click **Next**.
- 5 Select the group type from the menu.

Option	Description
VM protection groups	Select this option to create a recovery plan that contains array-based replication and vSphere Replication protection groups.
Storage policy protection groups	Select this option to create a recovery plan that contains storage policy protection groups.

The default is **VM protection groups**.

NOTE If using stretched storage, select **Storage policy protection groups** for the group type.

- 6 Select one or more protection groups for the plan to recover, and click **Next**.
- 7 Click the **Test Network** value, select a network to use during test recovery, and click **Next**.
The default option is to create an isolated network automatically.
- 8 Review the summary information and click **Finish** to create the recovery plan.

Edit a Recovery Plan

You can edit a recovery plan to change the properties that you specified when you created it. You can edit recovery plans from the protected site or from the recovery site.

Procedure

- 1 In the vSphere Web Client, select **Site Recovery > Recovery Plans**.

- 2 Right-click a recovery plan, and select **Edit Plan**.

You can also edit a recovery plan by clicking the **Edit recovery plan** icon in the **Recovery Steps** view in the **Monitor** tab.

- 3 (Optional) Change the name or description of the plan in the **Recovery Plan Name** text box, and click **Next**.

- 4 On the Recovery site page, click **Next**.

You cannot change the recovery site.

- 5 (Optional) Select or deselect one or more protection groups to add them to or remove them from the plan, and click **Next**.

- 6 (Optional) Click the test network to select a different test network on the recovery site, and click **Next**.

- 7 Review the summary information and click **Finish** to make the specified changes to the recovery plan.

You can monitor the update of the plan in the Recent Tasks view.

Configuring vRealize Suite Components for Disaster Recovery

4

You can configure vRealize Suite components that you want to protect by using the following information.

This chapter includes the following topics:

- “vRealize Automation Disaster Recovery,” on page 27
- “vRealize Orchestrator Disaster Recovery,” on page 39
- “vRealize Operations Manager Disaster Recovery,” on page 39
- “vRealize Log Insight Disaster Recovery,” on page 45

vRealize Automation Disaster Recovery

You can use the following information specific to disaster recovery for vRealize Automation 6.2.3.

Virtual Machines managed by vRealize Automation are associated with a vRealize Automation reservation. Each reservation resides on a vRealize Automation compute resource that is associated with a vRealize Automation business group. These resource virtual machines can be protected by Site Recovery Manager.

Each virtual machine which is protected using Site Recovery Manager is assigned to a protection group and a recovery plan. When Site Recovery Manager fails over a protected virtual machine from the protected site to a recovery site, the virtual machines change clusters characteristics placed on the mapped datastore. After a failover occurs, a scheduled vRealize Automation data collection is executed against the recovery site, the virtual machine entry in the vRealize Automation database is updated with the new cluster, storage, and vSphere ManagedObjectReference. However, the reservation is not modified. This failure to modify the reservation might result in an error for any ongoing vRealize Automation resource actions, such as power state change, and the task might not be performed. To resolve this issue, each virtual machine must be updated within vRealize Automation to point to the new reservation.

Change the reservation for the virtual machines within vRealize Automation by using either of the following methods:

- Select the **Infrastructure** tab and select **Machines > Managed Machines > Change Reservation** to update the reservation for each machine manually.
- Automate the update of the reservation details dynamically:
 - a Configure vRealize Automation endpoints with a custom property
 - b Configure CloudClient scripts on the Site Recovery Manager protected and recovery servers
 - c Configure Site Recovery Manager to invoke CloudClient scripts during the execution of the recovery plan.

NOTE For CloudClient to automatically update the reservation after recovery, the recovery vCenter endpoint should be configured with reservation that includes the arrays used for the reservation.

- For array-based replication, you should execute at least one sample recovery to be able to create the proper reservation for the recovery vCenter. The replication storage is not available on the recovery site, until you do so. Following are the additional steps required
 - a Execute array-based replication so that the storage is available on the recovery site.
 - b Perform a data collection for the recovery site in vRealize Automation.
 - c Edit the reservation for the recovery site by adding the replication storage.
 - d Re-protect and execute a recovery back to the protected site and re-protect again.
 - e In vCenter, navigate to **Site Recovery > Sites > Advanced Settings > Storage Provider** and set property **storageProvider.fixRecoveredDatastoreNames** to **true**. This prevents the array-based replication to change the name of the storage when a recovery is executed.

NOTE CloudClient does not support virtual machines provisioned by Advanced Service Designer (ASD). For all virtual machines provisioned by ASD, you should manually change the reservation after recovery.

Support Matrix

Disaster recovery for vRealize Automation 6.2.3 successfully works with Site Recovery Manager 5.8 and 6.1, only when you use the following combination of CloudClient and Site Recovery Manager versions.

Support	vRealize Automation	CloudClient	Site Recovery Manager	vCenter
	6.2.3	3.3	5.8	5.5
	6.2.3	3.4.1	6.1	6.0

CloudClient supports automatic provisioning of the managed virtual machines, however if you want to manually import the managed virtual machines, see [“Bulk Import, Update, or Migrate Virtual Machines,”](#) on page 34.

Configuring Custom Properties for vRealize Automation

To automate the reservation update, you must configure custom properties on each vSphere endpoint, where vRealize Automation managed virtual machines are protected by Site Recovery Manager to automate the reservation update. You must also configure and invoke CloudClient scripts on the Site Recovery Manager servers.

Configure a Custom Property for vRealize Automation EndPoints

You must configure a custom property for vRealize Automation endpoints on the protected and recovery sites.

Perform the following steps to edit the vRealize Automation endpoints for the protected and recovery vCenter servers.

Procedure

- 1 Log in to vRealize Automation.
- 2 On the **Infrastructure** tab, select **Endpoints > Endpoints**.
- 3 Select the endpoint for the protected site or the recovery site and click **Edit**.
- 4 Click **New Property**.
- 5 Enter **vcac.srm.vcenterinfo** as the **Name**, enter the FQDN of the protected or recovery vCenter server in the **Value** text box, and click **OK**.

- 6 Repeat this process for all the endpoints.

For example, following is the edit endpoint window for the recovery server.

Edit Endpoint - vSphere (vCenter)

Manage a specific endpoint.

General

*** Name:**

Description:

*** Address:**

*** Credentials:**

Specify manager for network and security platform

Custom properties:

Name	Value
vcac.srm.vcenterinfo	ra-vcenter-res-b1

To avoid conflict with vRealize Automation properties, use a prefix such as a company or feature name for all custom property names.

Configure Reservations and Execute Data Collections

Depending on the replication technology that you use, vSphere Replication, array based replication, or mixed, ensure that vRealize Automation data collection is executed against the recovery and protected sites and the reservations are configured correctly, before using CloudClient as part of the Recovery Plan execution.

NOTE For array based replication, the configured datastore drives are only available at the one of the sites before or after recovery operations. This requires a test recovery plan to be executed in order to discover the datastore drives, during the vRealize Automation data collection for the recovery and the protected sites, so that you can update the reservations accordingly.

You need to complete the following steps to configure reservations and execute data collections.

Procedure

- 1 Perform a data collection for the protected site in vRealize Automation.
- 2 Edit the reservation for the protected site by adding the replication storage.
- 3 Execute array based replication. The configured storage should now be available on the recovery site.
- 4 Perform a data collection for the recovery site in vRealize Automation.
- 5 Edit the reservation for the recovery site by adding the replication storage.
- 6 Re-protect and execute a recovery back to the protected site.

- 7 Re-protect again.

Configure CloudClient Scripts on the Site Recovery Manager Servers

vRealize CloudClient is a command-line utility that provides verb-based access with a unified interface across vRealize Automation APIs.

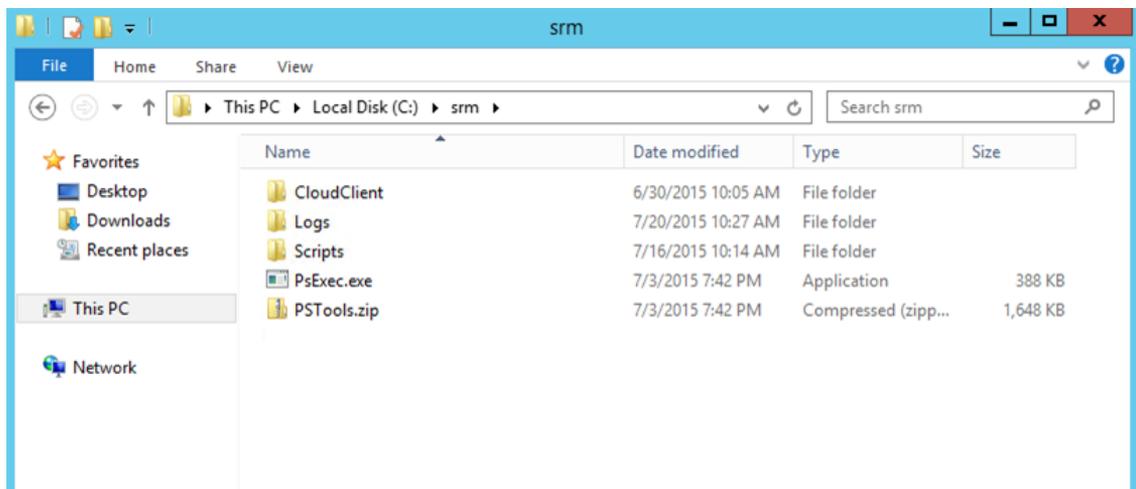
You must install and configure CloudClient batch files for pre-recovery and post-recovery steps, so that machines that vRealize Automation manages are mapped against new datastores at the recovery site. You must install CloudClient on the protected and the recovery Site Recovery Manager servers. For more information about vRealize CloudClient, see <https://developercenter.vmware.com/tool/cloudclient>.

You can manually import virtual machines when CloudClient fails or when the managed virtual machines are not moved to the recovered site by using Site Recovery Manager. For information about how to manually import virtual machines, see “Bulk Import, Update, or Migrate Virtual Machines,” on page 34.

Perform the following steps on the Site Recovery Manager servers for the protected and recovery site.

Procedure

- 1 Create a folder C:\srm.
- 2 Extract and install CloudClient in C:\srm.
- 3 Create a folder named Scripts in the srm directory.



- 4 Create scripts that Site Recovery Manager call out as pre and post fail-over custom steps.
 - a In the c:\srm\Scripts\ folder, create a file named pre_fail-over.bat, and include the following contents: `c:\srm\CloudClient\bin\cloudclient.bat srm protected vra vms stop datacollection --srmsrver <SRM_PROTECTED_HOST> --srmuser <SRM_USER> --srmpassword <SRM_PASSWORD> --recoveryplan %VMware_RecoveryName% --recoverymode %VMware_RecoveryMode%`
 - b In the c:\srm\Scripts\ folder, create a file named post_fail-over.bat, and include the following contents: `c:\srm\CloudClient\bin\cloudclient.bat srm protected vra vms start datacollection --recoveryplan %VMware_RecoveryName% --recoverymode %VMware_RecoveryMode% --vcenterinfo %VMware_VC_Host%`
 - c Verify that the Site Recovery Manager service is running with a user who has administrator permissions on the local machine. You will need to configure the script with the same account that has rights to query the Site Recovery Manager API for testing if an administrator account was used.
- 5 Configure the c:\srm\CloudClient\log4j2.xml debug level file for troubleshooting purpose.

Configure Site Recovery Manager to Invoke CloudClient Scripts

Each recovery plan that contains virtual machines managed by vRealize Automation requires a step at the beginning and end of the plan that invoke CloudClient scripts.

The CloudClient scripts updates the reservations for all vRealize Automation managed machines, except the ASD machines, even if they belong to different vRealize Automation tenants. In the CloudClient.properties file, you should provide the vRealize Automation server, tenant, tenant administrator credentials for only one of the tenants, IaaS Web server, and IaaS Windows Administrator credentials.

NOTE To create the CloudClient.properties file, start Windows Command prompt and navigate to `c:\srm\CloudClient\bin` folder and run `cloudclient.bat login autologinfile`. For more information, see the CloudClient documentation.

Procedure

- 1 Select the **Recovery Plan** in the Site Recovery Manager and navigate to **Monitor > Recovery Steps**.
- 2 In the **View** drop-down menu, right-click the first recovery step and select **Add Step Before** from the context menu.
- 3 Select the first step in the plan and click **Add Step**.
 - a Select **Command on SRM Server** from the **Type** drop-down menu.
 - b In the **Name** text box, enter **CloudClient_Prefail-over_Callout**.
 - c In the **Content** text box, enter
`c:\windows\system32\cmd.exe /c c:\srm\Scripts\pre_fail-over.bat`
 - d In **Timeout**, enter 5 minutes.

Edit Recovery Plan Step - 1. Command: Pre-Failover-Step

Type: Command on SRM Server
 Prompt (requires a user to acknowledge the prompt before the plan continues)

Name: CloudClient_Prefail_Over_Callout

Content: c:\Windows\System32\cmd.exe /c C:\srm\Scripts\pre_failover.bat

Timeout: 5 minutes 0 seconds

OK Cancel

- 4 Select **Prompt** to verify that vRealize Automation is fully operational before you execute post-failover step.

NOTE Before you execute the recovery, verify that the vRealize Automation is completely operational, all services in the management console are registered, and the load balancer pool detects the members as active.

- 5 Right-click the last step in the Recovery Plan and from the context menu, click **Add Step After**.
 - a Select **Command on SRM Server** for **Type**.
 - b In the **Name** text box, enter **CloudClient_Postfail-over_Callout**.
 - c In the **Content** text box, enter **c:\windows\system32\cmd.exe /c c:\srm\Scripts\post_failover.bat**.
 - d In **Timeout**, select **15** minutes.

Edit Recovery Plan Step - 16. Command: Post-Failover-Step

Type: Command on SRM Server
 Prompt (requires a user to acknowledge the prompt before the plan continues)

Name:

Content:

Timeout: minutes seconds

OK Cancel

- 6 Select **Prompt** and enter the **Name** and **Content** in the respective text boxes.
- 7 Repeat steps 1 through 6 for all recovery plans that have protected virtual machines that are managed by vRealize Automation.

NOTE CloudClient does not support XaaS provisioned VMs. For all VMs that are provisioned by XaaS, you must manually change the reservation after the recovery.

Configure Load Balancer

You can configure the load balancer that is used with your vRealize Automation system. The F5 load balancer is used as an example in this document, however you can use similar load balancer of your choice.

You should not change the IP addresses, so that the load balancer responds when the environment moves to the recovery site.

If you change the IP addresses of the virtual machines, you can configure the load balancer to be a combination of all active and inactive IP addresses at the same time. It is important that you set up appropriate health monitors. Even if you are able to ping the protected and recovery sites without setting up health monitors, you might still see some issues. The following example shows how the vRealize Automation virtual appliances were configured during testing.

Current Members						
<input checked="" type="checkbox"/>	Status	Member	Address	Service Port	FQDN	Ephemeral
<input type="checkbox"/>	◆	siteB.vn-vra-2.	443		No
<input type="checkbox"/>	◆	siteB.vn-vra-1.	443		No
<input type="checkbox"/>	●	vn-vra-2.	443		No
<input type="checkbox"/>	●	vn-vra-1.	443		No

For more information about configuring load balancer, see *vRealize Automation Load Balancing* guide on <https://www.vmware.com/support/pubs/vrealize-automation-pubs.html>

Update DNS

It is recommended that you do not change the DNS names of the nodes, so that you do not have to update the DNS records.

If you change the IP addresses, the DNS records should be updated while the recovery plan is in progress. You should replace the old IP addresses with the new IP addresses.

Reset RabbitMQ

You might experience an issue with the RabbitMQ queues while the environment is trying to get operational again.

To resolve the issue, navigate to the vRealize Automation Web interface on the vRealize Automation and reset the messaging for the primary VA node and every VA node.

To reset the messaging from the Web interface, select **vRealize Automation Settings > Messaging** and click **Reset Rabbitmq**.

To reset the messaging from your vRealize Automation appliance (Linux), run `/sbin/service rabbitmq-server stop` command, followed by `/sbin/service rabbitmq-server start` command.

Bulk Import, Update, or Migrate Virtual Machines

You can use the Bulk Import feature to import one or more virtual machines to a vRealize Automation deployment. You can also use this feature to update one or more virtual machines without the need to re-import them or to migrate machines from one environment to another.

The Bulk Import feature imports virtual machines intact with defining data such as reservation, storage path, blueprint, owner, and any custom properties. Bulk Import supports the following administrative tasks:

- Import one or more unmanaged virtual machines so that they can be managed in a vRealize Automation deployment
- Import one or more managed virtual machines from a vRealize Automation deployment into an upgraded deployment
- Make a global change to a virtual machine property, such as a storage path
- Migrate a virtual machine from one environment to another

You can execute the Bulk Import feature commands using either the vRealize Automation console or the CloudUtil command-line interface. For more information about using the CloudUtil command-line interface, see the *Life Cycle Extensibility* documentation.

Prerequisites

Log in to the vRealize Automation console as a **fabric administrator** and as a **business group manager**.

Generate Virtual Machine CSV Data File

You generate a virtual machine CSV data file to import, update, or migrate virtual machines to a vRealize Automation deployment.

Prerequisites

Log in to the vRealize Automation console as a **fabric administrator** and as a **business group manager**.

Procedure

- 1 Select **Infrastructure > Administration > Bulk Imports**.
- 2 Click **Generate CSV File**.
- 3 Select the machine type from the **Machines** drop-down menu.

Option	Description
Managed	Virtual machine is managed in a vRealize Automation deployment and can be viewed in the console. Select this option if you are updating a machine or migrating from one environment to another.
Unmanaged	Virtual machine exists in a hypervisor but is not managed in a vRealize Automation deployment and cannot be viewed in the console. Select this option if you are importing a virtual machine.

- 4 Select the **Business group** default value.
- 5 Select the **Owner** default value.
- 6 Select the **Blueprint** default value.

If you select **Unmanaged** for the machine type and select a value for **Business group** and **Blueprint**, you might see the following results in the CSV data file:

- Host Reservation (Name or ID) = INVALID_RESERVATION
- Host To Storage (Name or ID) = INVALID_HOST_RESERVATION_TO_STORAGE

This happens when you do not have a reservation in the selected business group for the host machine that also hosts the unmanaged machine. If you have a reservation in that business group for the unmanaged machine's host, the Host Reservation and Host To Storage values fill in properly.

- 7 Select the resource type from the **Resource** drop-down menu.

Option	Description
Endpoint	Information required to access a virtualization host.
Compute Resource	Information required to access a group of virtual machines performing a similar function.

- 8 Select the name of the virtual machine resource from the **Name** drop-down menu.
- 9 Click **OK**.

Edit Virtual Machine CSV Data File

Before you import or update one or more virtual machines, you must edit the virtual machine CSV data file so that each machine value matches a value that exists in the target deployment. If you are migrating a virtual machine from one environment to another, editing is optional.

To import, update or migrate virtual machines contained in a CSV data file, each machine must be associated with a reservation, storage location, blueprint, and owner that already exists in the target vRealize Automation deployment. All of the values for each machine must be present in the target vRealize Automation deployment for the operation to succeed. You can change the values for reservation, storage location, blueprint, and owner for any operation on each machine by editing the CSV file.

If you are importing a virtual machine that uses a static IP address, you must append the appropriate command to the CSV file.

Prerequisites

[“Generate Virtual Machine CSV Data File,”](#) on page 35

Procedure

- 1 Open the CSV file and edit the data categories so that they match existing categories in the target vRealize Automation deployment.

Heading	Comment
# Import--Yes or No	Can change to No to prevent a particular machine from being imported.
Virtual Machine Name	Do not change.
Virtual Machine ID	Do not change because it is ignored during the import process.
Host Reservation (Name or ID)	Must match the name of a reservation in the target vRealize Automation instance.
Host To Storage (Name or ID)	Must match the name of a storage location in the target vRealize Automation instance.
Blueprint (Name or ID)	Must match a blueprint in the target vRealize Automation instance.
Owner Name	Must match a domain user in the target vRealize Automation instance.

- 2 If you are importing a virtual machine with a static IP address, append a command in the following form to the CSV file.

```
,VirtualMachine.Network#.Address, w.x.y.z, HOP
```

Configure the command with the appropriate information for your virtual machine.

- Change the # to the number of the network interface being configured with this static IP address. For example, `VirtualMachineNetwork0.Address`
- Change `w.x.y.z` to be the static IP address for the virtual machine.
- The `HOP` string sets the visibility of the property. This default property is removed from the virtual machine after a successful import.

NOTE For a successful import, the IP address must be available in a properly configured address pool. If the address cannot be found or is already in use, the import will succeed without the static IP address definition, and an error will be logged.

- 3 Save the CSV file and close it.

Import, Update, or Migrate One or More Virtual Machines

After you edit the virtual machine CSV data file, you can import, update, or migrate one or more virtual machines into a vRealize Automation deployment.

You can import a managed machine or an unmanaged machine. You can migrate or update only managed machines. A managed machine is a virtual machine that is managed in a vRealize Automation deployment and that you can view in the console. An unmanaged machine is a virtual machine that exists in a hypervisor but is not managed in a vRealize Automation deployment and cannot be viewed in the console.

Prerequisites

[“Edit Virtual Machine CSV Data File,”](#) on page 36

Procedure

- 1 Select **Infrastructure > Administration > Bulk Imports**.
- 2 Click **New Bulk Import**.
- 3 Enter a name for this task in the **Name** text box.
- 4 Enter the CSV file name in the **CSV file** text box by browsing to the CSV file name.
- 5 Import the file using these options.

- Select **Now** to begin the import, update, or migrate process immediately.
- Select a start date and time in the **Start time** drop-down menu.

NOTE The specified start time is the server's local time and not the local time of the user's workstation.

- Select the number of seconds to delay each virtual machine registration in the **Delay (seconds)** drop-down menu.

NOTE To specify no delay, leave the option blank. Selecting this option slows the import process. Select this option when you import a large number of virtual machines.

- Select the total number of machines being registered at a given time in the **Batch size** menu.

NOTE To specify no limit, leave the option blank. Selecting this option slows the import process. Select this option when you import a large number of virtual machines.

- If you are importing virtual machines, select **Ignore managed machines** to omit managed machines during the import process.

NOTE By selecting this option, you can rerun the import without editing the CSV file to exclude machines that are already successfully imported.

- If you are updating virtual machines, do not select **Ignore managed machines**.
- If you are migrating machines, select **Ignore managed machines** in the target environment so that you can reprocess the CSV file.

- Select **Skip user validation** to omit validating users during the import process.

NOTE Selecting this option sets a machine's owner to the value listed in the Owner column of the CSV data file without verifying that the user exists. Selecting this option can decrease the import time.

- Select **Test import** to run the import process without importing machines.

NOTE Testing the import process allows you to test the CSV file for errors before you actually import the machines.

- 6 Click **OK**.

The progress of the operation appears on the Bulk Import Details page.

Start Up vRealize Automation

When you start vRealize Automation from the beginning, such as after a power outage or a controlled shutdown, you must start its components in a specified order.

Prerequisites

Verify that the load balancers that your deployment uses are running.

Procedure

- 1 Start the MSSQL database machine. If you are using a legacy PostgreSQL standalone database, start that machine as well.
- 2 Start your Identity Appliance or vSphere SSO appliance and wait for the startup to finish.
- 3 Start the primary vRealize Automation appliance instance and wait for the startup to finish.

The primary vRealize Automation appliance instance contains the writeable Appliance Database, if applicable, and is the last appliance that you shut down in an ordered shutdown procedure.

- 4 (Optional) If you are running a distributed deployment, start the secondary virtual appliances and wait for the startup to finish.

You must wait for one appliance to boot before you start up another appliance. Make sure that all services, besides IaaS and vRealize Orchestrator, are running on the appliance before you start another appliance.

- 5 Start the primary Web node and wait for the startup to finish.
- 6 (Optional) If you are running a distributed deployment, start all secondary Web nodes and wait 5 minutes.
- 7 Start the primary Manager Service node and wait for 2 to 5 minutes, depending on your site configuration.
- 8 Start the Distributed Execution Manager Orchestrator and Workers and all vRealize Automation agents.

You can start these components in any order and you do not need to wait for one startup to finish before you start another.

- 9 Verify that the startup succeeded.
 - a Navigate to the vRealize Automation appliance management console by using its fully qualified domain name, `https://vra-va-hostname.domain.name:5480/`.
 - b Click the **Services** tab.
 - c Click the **Refresh** tab to monitor the progress of service startup.

When all services are listed as registered, the system is ready to use.

vRealize Orchestrator Disaster Recovery

You can use the following information specific to disaster recovery for vRealize Orchestrator 6.0.2.

Do not change the IP addresses and DNS names for the virtual machines at the protected and recovery sites. The disaster recovery information for vRealize Automation applies to vRealize Orchestrator.

vRealize Operations Manager Disaster Recovery

You can use the following information specific to disaster recovery for vRealize Operations Manager 6.1.

NOTE vRealize Operations Manager disaster recovery was tested completely with vSphere Replication. However, vRealize Operations Manager was tested with the array-based replication only with EMC RecoverPoint, in a limited manner.

Guidelines for vRealize Operations Manager Migration and Recovery

You can use the following guidelines for recovering vRealize Operations Manager by using Site Recovery Manager.

- Migrate or recover vRealize Operations Manager virtual machines to an identical network configuration.

If the recovery site is configured to have the same network configuration as the protected site and a mapping is created between the identical networks, configure all replicated vRealize Operations Manager virtual machines to be started with the same IPs, because these virtual machines are the protected nodes. The recovered system will become operational after the planned migration or disaster recovery has finished successfully.
- Migrate or recover vRealize Operations Manager virtual machines to a different network configuration (DHCP). If the recovery site is configured to have a different network configuration than the protected site and a DHCP server exists, all protected vRealize Operations Manager instances must be reconfigured to use the new network addresses.
- Migrate or recover vRealize Operations Manager virtual machines to a different network configuration (static IPs). The recovery site might be configured to have a different network configuration than the protected site and the Site Recovery Manager might explicitly configure the static IPs for the vRealize Operations Manager instances. As a result, you must reconfigure all of the nodes to use the new network addresses.
- Testing a recovery plan. After you execute a test recovery plan, the virtual machines on the protected site remain powered on. The powered on state can cause some communication issues between the vRealize Operations Manager nodes on both sites. Always clean up the executed recovery test.

NOTE The load balancing configuration for vRealize Operations Manager disaster recovery is beyond the scope of this document.

Change IP Address for Single-Node vRealize Operations Manager

If your disaster recovery plan requires, change the IP address for the vRealize Operations Manager node to update the new IP address for the virtual machine.

Procedure

- 1 Open a console on a virtual machine or start an SSH session on the virtual machine, if SSH is enabled.

- 2 Take the slice offline. From `/usr/lib/vmware-vcopssuite/utilities/sliceConfiguration/bin`, run the `$VMWARE_PYTHON_BIN ./vcopsConfigureRoles.py --action bringSliceOffline --offlineReason <Put Comment Here> command`.
- 3 Shut down the virtual machine gracefully and change the IP address.
- 4 Power on the virtual machine. If the slice was not offline, the reboot might take 15 minutes. If it takes more than 15 minutes, take the slice offline.
- 5 Take the slice offline if the virtual machine is not powered on in 15 minutes. From `/usr/lib/vmware-vcopssuite/utilities/sliceConfiguration/bin`, run the `$VMWARE_PYTHON_BIN ./vcopsConfigureRoles.py --action bringSliceOffline --offlineReason <Put Comment Here> command`.
- 6 Verify that the slice is offline. Run the service `vmware-casa stop` command to stop the CaSA service.
- 7 Edit the `/storage/db/casa/webapp/hsqldb/casa.db.script` file to update the old IP address with the new IP address for the master node for all the instances.
- 8 Update the `$VCOPS_BASE/user/conf/cassandra/cassandra.yaml` file with the new IP address for `address`, `broadcast_rpc_address`, and `seeds` values.
- 9 Run the service `vmware-casa start` command to start the CaSA service.
- 10 Update the master node IP. From `/usr/lib/vmware-vcopssuite/utilities/sliceConfiguration/bin`, run the `$VMWARE_PYTHON_BIN ./vcopsConfigureRoles.py --adminCS <New IP Here>`.

This command updates the IP address in the following files. You can also change the IP address in these files manually.

- `/usr/lib/vmware-vcopssuite/utilities/sliceConfiguration/data/roleState.properties`
 - `/usr/lib/vmware-vcops/user/conf/gemfire.properties`
 - `/usr/lib/vmware-vcops/user/conf/persistence/persistence.properties`
- 11 Bring the slice online. From `/usr/lib/vmware-vcopssuite/utilities/sliceConfiguration/bin`, run the `$VMWARE_PYTHON_BIN ./vcopsConfigureRoles.py --action bringSliceOnline` command.
 - 12 Verify that the slice is online. Log in to the administration UI of the master node and verify that the node is up and collecting data.

Change IP Address for Multinode vRealize Operations Manager Configuration

If your disaster recovery plan requires, change the IP address for the vRealize Operations Manager multinode clustered configuration to update the new IP addresses on all the nodes, with or without high availability enabled.

Procedure

- 1 Open a console on the virtual machine or start an SSH session on the virtual machine if SSH is enabled.
- 2 Take slice offline on all the nodes. From `/usr/lib/vmware-vcopssuite/utilities/sliceConfiguration/bin`, run the `$VMWARE_PYTHON_BIN ./vcopsConfigureRoles.py --action bringSliceOffline --offlineReason <Put Comment Here> command`.
- 3 Shut down the virtual machines on which the IP addresses must change and change the required IP addresses, gracefully.
- 4 Power on the virtual machines. If the nodes were not offline, the reboot might take 15 minutes. If it takes more than 15 minutes, take the slice offline on all the nodes.

- 5 If the virtual machines are not powered on in 15 minutes, take the slice offline on all the nodes. From `/usr/lib/vmware-vcopssuite/utilities/sliceConfiguration/bin`, run `$VMWARE_PYTHON_BIN ./vcopsConfigureRoles.py --action bringSliceOffline --offlineReason <Put Comment Here>`.
- 6 Verify that all nodes are offline. On all the nodes, starting with the master node, run the service `vmware-casa stop` command to stop the CaSA service.
- 7 Edit the `/storage/db/casa/webapp/hsqldb/casa.db.script` file to update the old IP address with the new IP address for every node for all the instances.
- 8 Update the `$VCOPS_BASE/user/conf/cassandra/cassandra.yaml` file with the new IP addresses for the `listen_address` and `broadcast_rpc_address` values. Change the `seeds` value to the new master node IP address and add the replica node IP address if high availability is enabled.
- 9 On all the nodes starting with the master node, run the service `vmware-casa start` command to start the CaSA service.
- 10 Update the IP addresses of the master node and replica node. From `/usr/lib/vmware-vcopssuite/utilities/sliceConfiguration/bin`, run the `$VMWARE_PYTHON_BIN ./vcopsConfigureRoles.py --adminCS <Master IP Here>, <replica IP here>` command.

This command updates the IP address in the following files. You can also change the IP addresses in these files manually.

- `/usr/lib/vmware-vcopssuite/utilities/sliceConfiguration/data/roleState.properties`
 - `/usr/lib/vmware-vcops/user/conf/gemfire.properties`
 - `/usr/lib/vmware-vcops/user/conf/persistence/persistence.properties`
- 11 Verify that the `bind-address` property in these files points to the new IP address. You can edit the files manually, if required.
 - `/usr/lib/vmware-vcops/user/conf/gemfire.locator.properties`
 - `/usr/lib/vmware-vcops/user/conf/gemfire.native.properties`
 - 12 On all the nodes starting with the master node, bring the slice online. From `/usr/lib/vmware-vcopssuite/utilities/sliceConfiguration/bin`, run the `$VMWARE_PYTHON_BIN ./vcopsConfigureRoles.py --action bringSliceOnline` command.

Do not wait for the script to finish before bringing online the next node. All nodes must be brought online in parallel, because components must rediscover their remote partners during startup.
 - 13 Verify that all the nodes are online. Log in to the administration UI of the master node and verify that all nodes in the cluster are up and collecting data.
 - 14 Enable high availability from the Admin UI, if required.

Change the IP Address After Restoring Clusters on a Remote Host

After you have restored a vRealize Operations Manager cluster to a remote host, change the IP address of the master nodes and data nodes to point to the new host.

Prerequisites

- Verify that the restore job has completed successfully.
- Verify that the datastore on the new host has sufficient capacity for the new cluster.

Procedure

- 1 Shut down the vRealize Operations Manager cluster at the original location.

- 2 In the Virtual Appliance Management Interface (VAMI), access the machine from the vCenter console and run the `'/opt/vmware/share/vami/vami_set_network eth0 STATICV4 <new IP> <netmask> <gateway>` command to change the IP address for each node in the cluster.

For example, ``/opt/vmware/share/vami/vami_set_network eth0 STATICV4 10.145.152.170 255.255.252.0 10.145.155.253``
- 3 When the command has successfully run, restart the network, and reboot each node.
- 4 If you are using a remote collector, power on the remote collector node.
- 5 Access the master, data, and remote collector node using SSH, and run the ``$VMWARE_PYTHON_BIN /usr/lib/vmware-vcopssuite/utilities/sliceConfiguration/bin/vcopsConfigureRoles.py --action=bringSliceOffline --offlineReason="restore cluster"`` command to take the cluster offline.
- 6 Update the CaSA database on the master nodes, and then update the CaSA database on the data nodes.
 - a Run the `vmware-casa stop` command to stop the CaSA service.
 - b Edit the `/storage/db/casa/webapp/hsqldb/casa.db.script` file, replacing all instances of the old IP address and with the new IP address.
 - c Run the `vmware-casa start` command to start the CaSA service.
- 7 Edit the following configuration files, and replace all instances of the old IP address with the new IP address:
 - `/usr/lib/vmware-vcopssuite/utilities/sliceConfiguration/data/roleState.properties`
 - `/usr/lib/vmware-vcops/user/conf/gemfire.properties`
 - `/usr/lib/vmware-vcops/user/conf/gemfire.locator.properties`. This configuration file only runs on the master node. Edit the `locator` parameter accordingly.
 - `/usr/lib/vmware-vcops/user/conf/gemfire.native.properties`
 - `/usr/lib/vmware-vcops/user/conf/persistence/persistence.properties`
- 8 Edit the Cassandra configuration file `/usr/lib/vmware-vcops/user/conf/cassandra/cassandra.yaml` so that the `seeds` parameter points to the new IP address of the master node. The other IP address in this file points to the IP addresses of the nodes.
- 9 Log in to the vRealize Operations Manager administration interface, and restart the vRealize Operations Manager cluster.

Reconnect End Point Operations Management Agents

When vRealize Operations Manager is restored to an environment that uses a different IP configuration and different networks, you must reconnect the End Point Operations (EPOps) Management agents.

When the IP addresses of the vRealize Operations Manager cluster change, you must change the IP addresses of the EPOps Management agents to communicate properly with the newly configured IP address of the cluster. The agents only need to be reconfigured if they are configured to point directly at an IP or if the DNS name changes. If the agent points at a DNS name and only the IP is changing, you must only update the DNS record to reflect the new IP address and there is no need to reconfigure the agent.

If a dedicated load balancer is in use for the EPOps Management agents, you must also configure the load balancer.

Procedure

- 1 Log in to the machine where the EPOps agent is installed.
- 2 Navigate to the folder where the installation files are stored.

- 3 Delete the data folder.
- 4 Run the script for the `epops-agent` installation again and provide the new IP address of the vRealize Operations Manager instance.

Reconnect Remote Collector Node

After recovery, you must reconfigure the remote collector node to connect to the changed IP cluster.

NOTE It is optional to protect remote collector nodes by using Site Recovery Manager. If you decide not to protect a remote collector node by using Site Recovery Manager, you need to redeploy a new remote collector node after recovery.

Use the following steps to configure the remote collector node for non high availability operation.

Procedure

- 1 Shut down the cluster at the protected site.
- 2 Change the IP addresses for each node in the cluster. Access the virtual machines from the vCenter console and run the `/opt/vmware/share/vami/vami_set_network eth0 STATICV4 <new IP> <netmask> <gateway>` command on the master node and the data nodes.
- 3 After the command is run successfully, restart the network. When the correct network is assigned, reboot each node.
- 4 Power on the remote collector node.

NOTE Do not change the IP address for the remote collector node, if it was not protected as part of the protection group.

- 5 Take the cluster offline by accessing master, data, and remote collector nodes by using ssh and run `$VMWARE_PYTHON_BIN /usr/lib/vmware-vcopssuite/utilities/sliceConfiguration/bin/vcopsConfigureRoles.py --action=bringSliceOffline --offlineReason="restore cluster"` command.
- 6 Update the Cluster and Slice Administration (CaSA) database. Complete the following steps by starting with the master node, followed by the data and remote collector nodes respectively.
 - a Run `service vmware-casa stop` command to shut down CaSA.
 - b Edit `/storage/db/casa/webapp/hsqldb/casa.db.script` to replace the old IP addresses with the new IP addresses.
 - c Run `service vmware-casa start` command to start CaSA.
- 7 Edit the following configuration files on all the nodes starting with the master node, followed by the data and remote collector node respectively.
 - `/usr/lib/vmware-vcopssuite/utilities/sliceConfiguration/data/roleState.properties`
 - `/usr/lib/vmware-vcops/user/conf/gemfire.properties`
 - `/usr/lib/vmware-vcops/user/conf/gemfire.locator.properties`
 - `/usr/lib/vmware-vcops/user/conf/gemfire.native.properties`
 - `/usr/lib/vmware-vcops/user/conf/persistence/persistence.properties`

NOTE The Gemfire Locator runs on the master node only. Edit the locator locations accordingly.

- 8 Edit the Cassandra configuration `/usr/lib/vmware-vcops/user/conf/cassandra/cassandra.yaml`. The `seeds` variable needs to point to the master node. The 2 other variables are the IP address of the node.

- 9 Start the cluster. Access the Admin UI from a browser and bring the cluster online.

Starting vRealize Operations Manager

When you start vRealize Operations Manager from the beginning in a clustered configuration, you must start its components in a specific order.

You must power on the nodes in the following order for a successful recovery:

- 1 Master node
- 2 Replica node
- 3 Data nodes
- 4 Remote Collector nodes

Ensure that each node is online before attempting to start the next component.

Planning the Capacity and Time for vRealize Operations Manager Cluster

You must plan the capacity and time required for a cluster size of the vRealize Operations Manager system.

Calculate Bandwidth for vSphere Replication

To determine the bandwidth that vSphere Replication requires to replicate virtual machines efficiently, you should calculate the maximum data change rate within an RPO period divided by the link speed.

If you have groups of virtual machines that have different RPO periods, consider that the replication of all machines in all groups will trigger at the same time while calculating the bandwidth. This way the biggest amount of data will be considered.

For example, you might have four groups with RPO of 15 minutes, 1 hour, 4 hours, and 24 hours. Factor in all the different RPOs in the environment, the subset of virtual machines in your environment that is replicated, the change rate of the data within that subset, the amount of data changes within each configured RPO, and the link speeds in your network.

Examine how data change rate, traffic rates, and the link speed meet the RPO, and then look at the aggregate of each group.

- 1 Identify the average data change rate within the RPO by calculating the average change rate over a longer period, and then dividing it by the RPO.
- 2 Calculate how much traffic this data change rate generates in each RPO period.
- 3 Measure the traffic against your link speed.

For example, a data change rate of 100 GB requires approximately 200 hours to replicate on a T1 network, 30 hours to replicate on a 10 Mbps network, 3 hours on a 100Mbps network.

For information about calculating bandwidth requirements for vSphere Replication, see the Knowledge Base article at <http://kb.vmware.com/kb/2037268>.

For estimating average RPO, network bandwidth requirements, or the number of virtual machines that can be replicated, use the vSphere Replication calculator available on <http://www.vmware.com/vrcalculator>.

Sample Test Measurements

The following table displays the results of the internal measurements for a controlled lab environment, where vRealize Operations Manager 6.1 instance was replicated to a recovery site. This information can be used as a guideline while calculating necessary replication bandwidth. Results may vary in your environment.

20,000 VMs, 8 nodes, HA is disabled				
RPO time configured on the replication settings	1 hour	4 hours	6 hours	8 hours
Max amount of data accumulated for replication on the vRealize Operations Manager cluster during that period (all nodes)	< 16GB	< 60GB	< 130GB	> 130GB
Bandwidth required to satisfy RPO time	~ 23 MB/sec	~ 26 MB/sec	~ 31 MB/sec	> 31 MB/sec

Interval of time for each replication was measured by measuring the start and the stop time.

The test lab consisted of the following infrastructure and characteristics :

- 2 vCenter instances with vSphere Replication and Site Recovery Manager 6.0.
- Replication interval was initially set to 4h, later changed to 6h and finally to 8h. Independent measurements were performed for each time interval.
- vRealize Operations Manager cluster configuration consisted of 8 nodes:
 - vROps-1. Master node
 - vROps-2. Replica node
 - vROps-3. Data node 1
 - vROps-4. Data node 2
 - vROps-5. Data node 3
 - vROps-6. Data node 4
 - vROps-7. Remote Collector node 1
 - vROps-8. Remote Collector node 2
- vRealize Operations Manager was monitoring 20,000 virtual machines at the vCenter
- End Point Operations (EPOps) Management Agents were not configured
- High Availability was disabled while testing.

vRealize Log Insight Disaster Recovery

You can use the following information specific to disaster recovery for vRealize Log Insight 2.5 and 3.0 by using Site Recovery Manager.

Guidelines for Protecting vRealize Log Insight

To guard against expensive data center downtime, use the following guidelines for disaster recovery operations.

- Allocate enough resources at the protected and recovery sites. Verify that enough CPU resources and storage are allocated to protected and recovery sites, because some of the operations of disaster recovery setup are resource intensive.
- vRealize Log Insight does not support quiesced snapshots. If you are using a disaster recovery tool that supports quiesced snapshots, make sure to disable quiescing.
- The choice of replication type is critical when you are configuring any virtual machine for disaster recovery. Consider Recovery Point Objective (RPO), Recovery Time Objective (RTO), Cost and Scalability when you are planning the replication type to use.
- Use static IP addresses for all nodes in a vRealize Log Insight cluster.
 - Using static IP addresses eliminates the need to update the IP addresses of vRealize Log Insight cluster nodes each time the IP address of a vRealize Log Insight node changes.

- vRealize Log Insight includes all node IP addresses in each cluster node configuration file at `/storage/core/loginsight/config/loginsight-config.xml#<n>` where `<n>` is the largest number.
- Some products that integrate with vRealize Log Insight to feed their logs, use a fully qualified domain name (FQDN) or IP address as the syslog target. For example, vSphere ESXi, vSphere, and vRealize Operations Manager use the nodes of the cluster master's or the load balancer's (if configured) FQDN or IP address as the syslog target.
- Use an FQDN for all nodes in the vRealize Log Insight cluster.
 - For the master node, when you use a load balancer, a fully resolvable FQDN is required. Otherwise, the ESXi hosts fail to feed the syslog messages to vRealize Log Insight or to any remote target.
 - Using an FQDN saves time on post-restore and recovery configuration changes, assuming that the same FQDN can be resolved on the recovery site.
 - For system alerts, vRealize Log Insight uses FQDN host names if available instead of IP addresses.
 - Assuming that only underlying IP addresses change post-backup and recovery or disaster recovery operations, using FQDN eliminates the need to change the syslog target address (master node FQDN or internal load balancer FQDN) on all the external devices feeding logs to the vRealize Log Insight cluster.
 - With vRealize Log Insight 2.5, you must update the configuration file, located at `/storage/core/loginsight/config/loginsight-config.xml#<n>` where `<n>` is the largest number. This configuration file replaces the worker node IP address with the new IP address used for the restored nodes because the FQDN is not used for worker node addresses in the configuration file. You need to make this change only on the master node to synchronize the changes with all the worker nodes.
- Join requests from a vRealize Log Insight worker node should use the FQDN of the vRealize Log Insight master node.
 - Beginning in vRealize Log Insight 2.5, the master node host value in the configuration file on each of the nodes, located at `/storage/core/loginsight/config/loginsight-config.xml#<n>`, is based on the value used by the first worker node sending a join request. Using the FQDN of the master node for the join request prevents making any manual changes to the master node host value post-disaster recovery. Otherwise, the worker nodes cannot rejoin the master node until the master node host name is updated in the configuration files on all restored cluster nodes.
- Provide static IP addresses as well as optional virtual IP addresses for the load balancer.
 - When configuring an integrated load balancer, provide the optional FQDN for the virtual IP address. This optional FQDN enables vRealize Log Insight to revert to the FQDN when an IP address is not reachable for any reason.

Post-Recovery Configuration Change Guidelines

Depending on the recovery target and IP customizations applied during the backup configuration, manual configuration changes are required to one or more vRealize Log Insight nodes before the restored site can become fully functional.

Recovering to the Same Host

You can restore vRealize Log Insight cluster to the same host by using any back up tool.

- All network, IP, and FQDN settings that are used for the production environment should be preserved in the restored site.
- The original copy of the cluster is overwritten with the restored version unless a new name is provided to the virtual machine, during the restore process.

- If the same IP addresses and FQDNs are used for the restored cluster nodes as per the default settings, power down the existing cluster before beginning the restore.
- After a successful restoration and validation, delete the old copy to conserve resources and to prevent potential issues.

Recovering to a Different Host

You must perform manual configuration on vRealize Log Insight, if you are restoring to a different host cluster. For information about changes that are specific to vRealize Log Insight 2.5 and 3.0 versions, see [“Restoring to a Different Host,”](#) on page 47. It is assumed that the restored vRealize Log Insight nodes have been assigned different IP addresses and FQDNs than their source counterparts from which a backup was taken.

Recovering vRealize Log Insight Forwarders

The manual instructions for recovering vRealize Log Insight forwarders are the same as that of the vRealize Log Insight server as described above.

Recovering vRealize Log Insight Agents

If the complete agent OS is backed up, follow the tool-specific workflow to recover the agent OS.

- If agent configuration is made on the client side, that is on agent OS, replace the agent.ini using the backup copy.
- If configuration changes are made on the server side, that is vRealize Log Insight master node, no backup and recovery is required for the agent virtual machines.

Confirming the Restoration

You must confirm that all restored vRealize Log Insight clusters are fully functional.

- Verify that you can access the vRealize Log Insight user interface using the Internal Load Balancer (ILB) IP address or FQDN (if configured) as well as access all individual cluster nodes using respective IP addresses or FQDNs.
- From the vRealize Log Insight Administration page:
 - Verify the status of cluster nodes from the cluster page and make sure the ILB, if configured, is also in an active state.
 - Verify the vSphere integration. If required, reconfigure the integration. This occurs when the ILB and/or the master node IP address or FQDN is changed post-recovery.
 - Verify the vRealize Operations Manager integration and reconfigure again if needed.
 - Verify that all content packs and UI features are functioning correctly.
 - Verify that vRealize Log Insight forwarders and agents, if configured, are functioning correctly.
- Verify that other key features of vRealize Log Insight are functioning as expected.

Restoring to a Different Host

When you restore your system to a different host, you should make some configuration changes on the vRealize Log Insight cluster.

The configuration changes listed are specific to vRealize Log Insight 2.5 and 3.0. It is assumed that the restored vRealize Log Insight nodes are assigned different IP addresses and FQDNs than their source counterparts from which the backup was taken.

Procedure

- 1 List all new IP addresses and FQDNs that were assigned to each vRealize Log Insight node.
- 2 Perform the following configuration changes on the master node:
 - a Power on the master node, if it is not ON.

NOTE Steps b through e are applicable for vRealize Log Insight 2.5. You can not make changes to the configuration files directly from the appliance console for vRealize Log Insight 3.0 and higher. To make changes to the internal configuration options by using the web UI interface for vRealize Log Insight 3.0 and higher, refer to the Knowledge Base article [KB 2123058](#).

- b Use SSH to connect as a root user to the node's virtual appliance.
- c If the vRealize Log Insight service is running, stop the service first by running this command `service loginsight stop`.
- d Run `cd /storage/core/loginsight/config`
- e Run `cp loginsight-config.xml#<n> backup-loginsight-config.xml` where <n> represents the largest number that is automatically suffixed to `loginsight-config.xml` during configuration changes.
- f Open the copied version of the configuration file in your favorite editor or in the vRealize Log Insight 3.0 web UI and look for lines that resemble the following lines. This configuration change is applicable to both vRealize Log Insight 2.5 and 3.0.

```
<distributed overwrite-children="true">
  <daemon host="prod-es-vrli1.domain.com" port="16520" token="c4c4c6a7-f85c-4f28-
a48f-43aeea27cd0e">
    <service-group name="standalone" />
  </daemon>
  <daemon host="192.168.1.73" port="16520" token="a5c65b52-aff5-43ea-8a6d-38807ebc6167">
    <service-group name="workernode" />
  </daemon>
  <daemon host="192.168.1.74" port="16520" token="a2b57cb5-a6ac-48ee-8e10-17134e1e462e">
    <service-group name="workernode" />
  </daemon>
</distributed>
```

In this code snippet, there are three nodes. The first one is the master node which shows `<service-group name=standalone>` and the remaining two nodes are worker nodes and show `<service-group name="workernode">`.

- g For the master node, in the newly recovered environment, verify if the DNS entry that was used in the pre-recovery environment can be reused.
 - If the DNS entry can be reused, you only need to update the DNS entry to point to the new IP address of the master node.
 - If the DNS entry cannot be reused, replace the master node entry with the new DNS name, pointing to the new IP address.
 - If the DNS name cannot be assigned, as a last option, update the configuration entry with the new IP address.
- h Update the worker node IP addresses to reflect the new IP addresses.

- i In the same configuration file, look for entries that represent NTP, SMTP and database, and appenders sections.

This applies to vRealize Log Insight 2.5 and 3.0.

NOTE The `<logging><appenders>...</appenders></logging>` section is applicable only to the vRealize Log Insight 2.5 and is not available for vRealize Log Insight 3.0.

```
<ntp>
  <ntp-servers value="ntp1.domain.com, ntp2.domain.com" />
</ntp>

<smtp>
  <server value="smtp.domain.com" />
  <default-sender value="source.domain.com@domain.com" />
</smtp>

<database>
  <password value="xserttt" />
  <host value="vrli-node1.domain.com" />
  <port value="12543" />
</database>

<logging>
  <appenders>
    <appender name="REMOTE"
class="com.vmware.loginsight.common.logging.ThriftSocketAppender">
      <param name="RemoteHost" value="vqli-node1.domain.com" />
    </appender>
  </appenders>
</logging>
```

- If the configured NTP server values are not valid in the new environment, update these in the `<ntp>...</ntp>` section.
- If the configured SMTP server values are not valid in the new environment, update these in the `<smtp>...</smtp>` section.
- Optionally, change the `default-sender` value in the SMTP section. The value can be any value, but as a good practice, you should represent the source from where the email was sent.
- In the `<database>...</database>` section, change the `host` value to point to the master node FQDN or IP address.
- In the `<logging><appenders>...</appenders></logging>` section, change the parameter value for `RemoteHost` to reflect the new master node FQDN or IP address.

- j In the same configuration file, update the vRealize Log Insight ILB configuration section

For a vRealize Log Insight 3.0 appliance,

```
<load-balancer>
<leadership-lease-renewal-secs value="5" />
<high-availability-enabled value="true" />
<high-availability-ip value="10.158.128.165" />
<high-availability-fqdn value="LB-FQDN.eng.vmware.com" />
<layer4-enabled value="true" />
<ui-balancing-enabled value="true" />
</load-balancer>
```

For a vRealize Log Insight 2.5 appliance,

```
<load-balancer>
  <leadership-lease-renewal-secs value="5" />
  <high-availability-enabled value="true" />
  <high-availability-ip value="192.168.1.75" />
  <layer4-enabled value="true" />
</load-balancer>
```

- k Under the `<load-balancer>...</load-balancer>` section, update the `high-availability-ip` value if it is different from the current setting.
- l In the vRealize Log Insight 3.0, make sure to also update the FQDN of the load balancer.
- m Rename the updated configuration file to finish the changes.

NOTE This step is applicable for vRealize Log Insight 2.5 only. In vRealize Log Insight 3.0 the changes are made through web UI.

Run : `mv backup-loginsight-config.xml loginsight-config.xml#<n+1>` where n represents the current maximum number suffixed to the `loginsight-config.xml` files.

- n For vRealize Log Insight 2.5, restart the vRealize Log Insight service and run : `service loginsight start`.

NOTE For vRealize Log Insight 3.0, this can be achieved from the web UI by going to the Cluster tab on the Administration page. For each node listed, select its hostname or IP address to open the details panel and click **Restart Log Insight**. The configuration changes are automatically applied to all cluster nodes.

- o Wait for two minutes after the vRealize Log Insight service starts in order to give enough time for Cassandra services to come up before bringing other worker nodes online.

NOTE You can skip steps 3 to 9 for vRealize Log Insight 3.0. These steps are only applicable for vRealize Log Insight 2.5.

- 3 SSH onto the first worker node using root credentials.
- 4 Stop the vRealize Log Insight service and run : `service loginsight stop`.
- 5 Copy the latest `loginsight-config.xml` file from the master node to the worker node.
- 6 On the worker node, run : `scp root@[master-node-ip]:/storage/core/loginsight/config/loginsight-config.xml#<n> /storage/core/loginsight/config/`
- 7 Run : `service loginsight start`.

- 8 Wait for 2 minutes after the vRealize Log Insight service starts in order to give enough time for Cassandra service to start completely.
- 9 Repeat the steps for each worker node.

Testing and Executing a Recovery Plan

5

By testing a recovery plan, you ensure that the virtual machines that the plan protects recover correctly to the recovery site. If you do not test recovery plans, a disaster recovery situation might not recover all virtual machines, resulting in data loss.

This chapter includes the following topics:

- [“Test a Recovery Plan,”](#) on page 53
- [“Clean Up After Testing a Recovery Plan,”](#) on page 54
- [“Run a Recovery Plan,”](#) on page 54
- [“Cancel a Test or Recovery,”](#) on page 56

Test a Recovery Plan

When you test a recovery plan, Site Recovery Manager runs the virtual machines of the recovery plan on a test network and on a temporary snapshot of replicated data at the recovery site. Site Recovery Manager does not disrupt operations at the protected site.

Testing a recovery plan runs all the steps in the plan, except for powering down virtual machines at the protected site and forcing devices at the recovery site to assume mastership of replicated data. If the plan requires the suspension of local virtual machines at the recovery site, Site Recovery Manager suspends those virtual machines during the test. Running a test of a recovery plan makes no other changes to the production environment at either site.

Testing a recovery plan creates a snapshot on the recovery site of all of the disk files of the virtual machines in the recovery plan. The creation of the snapshots adds to the I/O latency on the storage. If you notice slower response times when you test recovery plans and you are using VMware Virtual SAN storage, monitor the I/O latency by using the monitoring tool in the Virtual SAN interface.

Procedure

- 1 In the vSphere Web Client, select **Site Recovery > Recovery Plans**, and select a recovery plan.
- 2 Right-click the plan and select **Test**.

You can also run a test by clicking the **Test recovery plan** icon in the **Recovery Steps** view in the **Monitor** tab.

- 3 (Optional) Select **Replicate recent changes to recovery site**.

Selecting this option ensures that the recovery site has the latest copy of protected virtual machines, but means that the synchronization might take more time.

- 4 Click **Next**.
- 5 Review the test information and click **Finish**.

- 6 Click the **Recovery Steps** in the **Monitor** tab to monitor the progress of the test and respond to messages.

The **Recovery Steps** tab displays the progress of individual steps. The Test task in Recent Tasks tracks overall progress.

NOTE Site Recovery Manager runs recovery steps in the prescribed order, except that it does not wait for the Prepare Storage step to finish for all protection groups before continuing to the next steps.

What to do next

Run a cleanup operation after the recovery plan test finishes to restore the recovery plan to its original state from before the test.

Clean Up After Testing a Recovery Plan

After you test a recovery plan, you can return the recovery plan to the Ready state by running a cleanup operation. You must complete the cleanup operation before you can run a failover or another test.

Site Recovery Manager performs several cleanup operations after a test.

- Powers off the recovered virtual machines.
- Replaces recovered virtual machines with placeholders, preserving their identity and configuration information.
- Cleans up replicated storage snapshots that the recovered virtual machines used during the test.

Prerequisites

Verify that you tested a recovery plan.

Procedure

- 1 In the vSphere Web Client, select **Site Recovery > Recovery Plans**, and select a recovery plan.

- 2 Right-click the recovery plan and select **Cleanup**.

You can also run cleanup by clicking the cleanup icon in the **Recovery Steps** view in the **Monitor** tab.

- 3 Review the cleanup information and click **Next**.

- 4 Click **Finish**.

- 5 After the cleanup finishes, if it reports errors, run the cleanup again, selecting the **Force Cleanup** option.

The **Force Cleanup** option forces the removal of virtual machines, ignoring any errors, and returns the plan to the Ready state. If necessary, run cleanup several times with the **Force Cleanup** option, until the cleanup succeeds.

Run a Recovery Plan

When you run a recovery plan, Site Recovery Manager migrates all virtual machines in the recovery plan to the recovery site. Site Recovery Manager attempts to shut down the corresponding virtual machines on the protected site.



CAUTION A recovery plan makes significant alterations in the configurations of the protected and recovery sites and it stops replication. Do not run any recovery plan that you have not tested. Reversing these changes might cost significant time and effort and can result in prolonged service downtime.

Prerequisites

- To use forced recovery, you must first enable this function. You enable forced recovery by enabling the **recovery.forceRecovery** setting as described in [Change Recovery Settings](#).
- Ensure that you have configured full inventory mappings. If you have only configured temporary placeholder inventory mappings and you run a planned migration with the **Enable vMotion of eligible VMs** option, planned migration fails, even though both sites are running.
- To use the **Enable vMotion of eligible VMs** option with planned migration, enable vMotion on the virtual machines. For instructions about enabling vMotion on virtual machines, see [Enable vSphere vMotion for Planned Migration](#).

Procedure

- 1 In the vSphere Web Client, select **Site Recovery > Recovery Plans**, and select a recovery plan.
- 2 Right-click the recovery plan and select **Run**.
- 3 Review the information in the confirmation prompt, and select **I understand that this process will permanently alter the virtual machines and infrastructure of both the protected and recovery datacenters**.
- 4 Select the type of recovery to run.

Option	Description
Planned Migration	Recovers virtual machines to the recovery site when both sites are running. If errors occur on the protected site during a planned migration, the planned migration operation fails. If your array supports stretched storage, select the Enable vMotion of eligible VMs check box.
Disaster Recovery	Recovers virtual machines to the recovery site if the protected site experiences a problem. If errors occur on the protected site during a disaster recovery, the disaster recovery continues and does not fail.

- 5 (Optional) Select the **Forced Recovery - recovery site operations only** check box.
This option is available if you enabled the forced recovery function and you selected **Disaster Recovery**.
- 6 Click **Next**.
- 7 Review the recovery information and click **Finish**.

When you run planned migration of a recovery plan that contains a storage policy protection group, Site Recovery Manager checks that the protection groups are synchronized on both of the protected and recovery sites before it runs the recovery plan. This check happens when you click **Finish**. If the protection group is synchronized on both sites, the planned migration begins. If the protection group is not synchronized at both sites, you see the error *The peer site has not finished synchronizing changes to protection group*.

- If you see this error, dismiss the error message and click **Finish** again. If the synchronization has completed, the planned migration begins.
 - If the error persists, cancel the planned migration, wait a short time, and attempt to run the planned migration again.
- 8 Click the **Monitor** tab and click **Recovery Steps**.

The **Recovery Steps** tab displays the progress of individual steps. The Recent Tasks area reports the progress of the overall plan.

Cancel a Test or Recovery

You can cancel a recovery plan test whenever the status is test in progress or failover in progress.

When you cancel a test or recovery, Site Recovery Manager does not start processes, and uses certain rules to stop processes that are in progress. Canceling a failover requires you to re-run the failover.

- Processes that cannot be stopped, such as powering on or waiting for a heartbeat, run to completion before the cancellation finishes.
- Processes that add or remove storage devices are undone by cleanup operations if you cancel.

The time it takes to cancel a test or recovery depends on the type and number of processes that are currently in progress.

Procedure

- 1 In the vSphere Web Client, select **Site Recovery > Recovery Plans**, and select a recovery plan.
- 2 Right-click the recovery plan and select **Cancel**. You can also cancel the plan from the Recovery Steps tab.

What to do next

Run a cleanup after canceling a test.

Perform a Failback

After Site Recovery Manager performs a recovery, you can perform a failback to restore the original configuration of the protected and recovery sites.

To aid comprehension, the original protected site from before a recovery is site A. The original recovery site is site B. After a recovery from site A to site B, the recovered virtual machines are running on site B without protection.

Prerequisites

Verify that the following conditions are in place.

- You have performed a recovery, either as part of a planned migration or as part of a disaster recovery.
- The original protected site, site A, is running.
- If you performed a disaster recovery, you must perform a planned migration recovery when the hosts and datastores on the original protected site, site A, are running again.
- You did not run reprotect since the recovery.

Procedure

- 1 In the vSphere Web Client, select **Site Recovery > Recovery Plans**.
- 2 Right-click a recovery plan and select **Reprotect**.
- 3 Select the check box to confirm that you understand that the reprotect operation is irreversible and click **Next**.
- 4 Determine whether to enable **Force Cleanup** and click **Next**.

This option is only available after you have run reprotect once and errors occurred. Enabling this option forces the removal of virtual machines, ignoring errors, and returns the recovery plan to the ready state.
- 5 Review the reprotect information and click **Finish**.
- 6 In the **Monitor** tab, click **Recovery Steps** to monitor the reprotect operation until it finishes.
- 7 (Optional) If necessary, rerun reprotect until it finishes without errors.

At the end of the reprotect operation, Site Recovery Manager has reversed replication, so that the original recovery site, site B, is now the protected site.
- 8 (Optional) After the test completes, right-click the recovery plan and select **Cleanup** to clean up the recovery plan.
- 9 Right-click the recovery plan and select **Recovery** to run the recovery plan as a planned migration.

- 10 In the **Monitor** tab, click **Recovery Steps** to monitor the planned migration until it finishes.

The planned migration shuts down the virtual machines on the new protected site, site B, and starts up the virtual machines on the new recovery site, site A. If necessary, rerun the planned migration until it finishes without errors.

When the planned migration completes, the virtual machines are running on the original protected site, site A, but the virtual machines are not protected. The virtual machines on the original recovery site, site B, are powered off.

- 11 Right-click the recovery plan and select **Reprotect** and follow the instructions of the wizard to perform a second reprotect operation.

Running reprotect again reestablishes protection in the original direction from before the recovery.

You restored the protected and recovery sites to their original configuration before the recovery. The protected site is site A, and the recovery site is site B.

Configuring vRealize Suite Components Post Failback

After you perform failback, you might have to reconfigure the vRealize Suite components to restore the original configuration of the protected and recovery sites.

- For vRealize Automation, you do not have to reconfigure any components after failback. If for any reason, the failback fails, see related information under [“Configure Load Balancer,”](#) on page 33.
- For vRealize Operations Manager, if your disaster recovery plan requires you to change IP addresses after restoring the clusters on a remote host, see [“Change the IP Address After Restoring Clusters on a Remote Host,”](#) on page 41. You would have to reconnect the remote collector node, if you have protected it by using Site Recovery Manager. For more information on how to reconnect the remote collector node, see [“Reconnect Remote Collector Node,”](#) on page 43
- For vRealize Orchestrator and vRealize Log Insight, you do not have to reconfigure any components after the failback.

Index

A

- all paths down (APD) **54**
- array based recovery plan, create **24**
- array-based replication, and vSphere Replication **18**
- array-based replication versus vSphere replication **15**

C

- capacity planning **44**
- change IP address after a restore job **41**
- change IP address for the node **39**
- change IP address for the clustered configuration **40**
- cleanup, recovery plan **54**
- CloudClient **33**
- configure reservations **29**
- configure CloudClient scripts **30**
- configure custom properties **28**
- configure remote collector node **43**
- configure virtual machines for array based replication **22**
- configure virtual machines for vSphere replication **21**
- configuring vRealize Suite **27**
- CSV data file
 - edit **36**
 - generate **35**
 - virtual machine **35**

D

- datastore, protected **16**
- datastore group **16**
- disaster recovery **8**

E

- edit endpoint **33**
- end point management agents **42**
- execute data collections **29**

F

- failback
 - diagram **12**
 - perform **57**
- failover, effects of **54**
- Flash Read Cache **16**

- forced recovery **54**

H

- host-based replication **17**

I

- introduction **7**
- invoke CloudClient scripts **31**

M

- MPIT **17**

O

- overview **8**

P

- planned migration **8**
- point-in-time recovery **17**
- post failback **58**
- post-recovery configuration change guidelines **46**
- prepare the environment **21**
- protected and recovery sites
 - different configurations **10**
 - heterogeneous **10**
- protected site **9**
- protection groups
 - array-based replication **22**
 - create **22**
 - storage policy **22**
 - vSphere Replication **22**

R

- recovering to a different host **47**
- recovery, of datastores in APD state **54**
- recovery plan
 - cleanup **53, 54**
 - force cleanup **53**
 - forced recovery **54**
 - running **54**
 - testing **53**
 - to change properties of **24**
- recovery plan test **11**
- recovery site **9**
- recovery test, to cancel **56**
- replication, array-based **16**

- replication types **15**
- reset RabbitMQ **34**
- restore a system to a remote location **41**

S

- SRM architecture diagram
 - array-based replication, storage policy protection **16**
 - array-based replication and vSphere Replication **18**
 - vSphere Replication **17**
- start up vRealize Operations Manager **44**
- Storage DRS, with array-based replication **16**
- storage policies, and array-based replication **16**
- storage replication adapter **16**
- Storage vMotion, with array-based replication **16**

T

- test **11**
- test recovery plan **53**
- types of replication **15**

V

- virtual machine
 - CSV data file **35**
 - edit CSV data file **36**
 - import **37**
 - managed **35, 37**
 - unmanaged **35, 37**
- Virtual SAN **17, 53**
- virtual machine{CSV data file **34**
- vRealize Automation, starting up **38**
- vRealize Orchestrator **27**
- vRealize Operations Manager single-node **39**
- vRealize Operations Manager disaster recovery **39**
- vRealize Orchestrator disaster recovery **39**
- vRealize Suite components **7**
- vRelaize Log Insight disaster recovery **45**
- VSAN **17, 53**
- vSphere Replication
 - and array-based replication **18**
 - introduction **17**
- vSphere Replication server, role **17**
- vSphere Replication management server, role **17**