Veritas™ Resiliency Platform 2.1 Solutions for Microsoft Hyper-V

VERITAS™

Veritas Resiliency Platform: Solutions for Microsoft Hyper-V

Last updated: 2017-03-15

Document version: Document version: 2.1 Rev 1

Legal Notice

500 E Middlefield Road
Mountain View, CA 94043

http://www.veritas.com

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

https://www.veritas.com/support

You can manage your Veritas account information at the following URL:

https://my.veritas.com

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)	CustomerCare@veritas.com
Japan	CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The document version appears on page 2 of each guide. The latest documentation is available on the Veritas website:

https://sort.veritas.com/documents

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

doc.feedback@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

http://www.veritas.com/community/

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Section 4   Managing disaster recovery

Chapter 7   Rehearsing DR operations to ensure DR readiness

Chapter 8   Performing disaster recovery operations

Chapter 9   Managing resiliency plans

Chapter 10   Monitoring risks, reports, and activites

Section

1

Overview of Resiliency Platform

Overview of Resiliency Platform

This chapter includes the following topics:

- About Veritas Resiliency Platform

- About disaster recovery using Resiliency Platform

- About Resiliency Platform features and components

- About Resiliency Platform capabilities

- About managing Hyper-V virtual machines using Resiliency Platform

- About permissions for operations in the console

About Veritas Resiliency Platform

Veritas Resiliency Platform offers a unified solution that helps you proactively maintain business uptime across private, public, and hybrid clouds. Resiliency Platform gives you complete automation for all resiliency operations involving the virtual machines, applications, and multi-tier business-services in your data center. It safeguards the current technology investments by plugging into your existing environments and infrastructure.

For data replication, you can use the Veritas Data Mover or any third-party solution that is supported by Veritas Resiliency Platform. For a list of supported vendors and products, see *Veritas Resiliency Platform Hardware and Software Compatibility Guide*.

Resiliency Platform has the following core capabilities:

Security and Compliance	Veritas Resiliency Platform provides enhanced data encryption (for data-in-flight and data-at-rest) as well as choice of data residency.
Predictability	Customers can predictably meet critical business Recovery Time Objectives (RTOs) and Recovery Point Objectives (RPOs).
Compliance	Customers can prove compliance to internal and external business continuity mandates with audit reporting and non-disruptive, real-time disaster recovery testing.
Automation	Customers get complete automation for all resiliency operations including recovery run books, and start and stop recovery orchestration for multi-tier applications. This reduces risk of downtime from human error.
Flexibility	Customers get the flexibility to keep their existing infrastructures and can innovate on their terms, with the flexibility that Resiliency Platform provides, to enable workload migration across sites and even to the cloud.

About disaster recovery using Resiliency Platform

A comprehensive disaster recovery strategy ensures that your mission-critical IT functions can continue during and after a disaster and any unforeseen risk can be mitigated to the extent possible.

Veritas Resiliency Platform lets you perform disaster recovery operations on your critical IT services. This section introduces you to the key features of Resiliency Platform:

- Monitoring of data center assets - storage, virtual machines, and applications.

- Ability to group your virtual machines or applications in resiliency groups based on your production environment and business needs.

- Making business services more resilient by providing the ability to perform disaster recovery operations on virtual machines and applications. For example, migrate and take over.

- Ability to replicate data from virtual machines on source data centers to target data centers using Resiliency Platform Data Mover integrated with VMware API I/O filtering framework or array-based replication technologies provided by array vendors.

- Resiliency plan (a sequential execution of predefined steps) to automate site-level recovery operations on your IT infrastructure in the event of downtime.

- Auto-discovery and real-time tracking for recovery objectives.

- Ability to perform non-disruptive testing (rehearsal) on your virtual machines and applications to ensure that your infrastructure is adequately prepared for protection in the event of disaster.

- Reporting capabilities providing details about resiliency health of applications and virtual machines.

See "Understanding the role of resiliency groups in disaster recovery operations" on page 60.

About Resiliency Platform features and components

The following is a brief introduction to Veritas Resiliency Platform key components and their relationships. Administrators responsible for deploying and configuring the product need to understand these in more detail.

resiliency domain	The logical scope of a Resiliency Platform deployment. It can extend across multiple data centers.
Resiliency Manager	The component that provides resiliency capabilities within a resiliency domain. It is composed of loosely coupled services, a distributed data repository, and a management console. The Resiliency Manager is deployed as a virtual appliance.
Infrastructure Management Server (IMS)	The component that discovers, monitors, and manages the asset infrastructure within a data center. The IMS transmits information about the asset infrastructure to the Resiliency Manager. The IMS is deployed as a virtual appliance. To achieve scale, multiple IMSs can be deployed in the same data center.
Veritas InfoScale Operations Manager Management Server	The component that allows discovery of InfoScale applications that are already configured in Veritas InfoScale Operations Manager. Also referred to as Veritas InfoScale Operations Manager server. You can manage the InfoScale applications that are already configured in Veritas InfoScale Operations Manager on Linux, Solaris, AIX as well as Windows platform.

Replication Gateway	The component of Veritas Resiliency Platform Data Mover that is deployed as a virtual appliance on both data centers and used to perform replication between the data centers.
data center	For a disaster recovery use case, the resiliency domain must contain at least two data centers in different locations, a production data center and recovery data center. Each data center has a Resiliency Manager and one or more IMSs.
asset infrastructure	The data center assets that you add to Resiliency Platform for discovery and monitoring by the IMS.
	The asset infrastructure can include hosts (Windows or Linux servers), virtualization servers for Hyper-V and VMware, and enclosures (storage arrays). Once the asset infrastructure is discovered by the IMS, the discovered virtual machines or applications are listed in the console as assets to manage or protect.
resiliency group	The unit of management and control in Resiliency Platform. You organize related assets into a resiliency group and manage and monitor them as a single entity.
service objective	A template to define the type of operations and technologies that are supported for a group of assets. You apply a service objective to each resiliency group.
	A template which identifies the characteristics of a service. These could be availability related characteristics such as local redundancy, and number of nodes in a cluster or DR characteristics such as remote recovery, Recovery Point Objective (RPO) SLAs, rehearsal support etc. Service objective is applied when a group of assets are being added to a resiliency group.
	Resiliency Platform monitors the resiliency groups based on the service objective definition and raises the risks as applicable.
Virtual Business Service (VBS)	A multi-tier business service where each VBS tier hosts one or more resiliency groups. A VBS lets you group multiple services as a single unit for visualization, automation, and controlled start and stop in the desired order. VBS uses the vertical grouping mechanism to group the multiple services.You can also perform operations such as migrate, takeover, resync, rehearsal on the entire VBS.

For more information on the above components, refer to the Deployment Guide.

About Resiliency Platform capabilities

Resiliency Platform helps you monitor and manage recovery across multiple data centers. It provides the following capabilities.

Table 1-1   Resiliency Platform capabilities

Capability	More information
Configuring virtual machines and applications for remote recovery operations or basic monitoring	See "Managing virtual machines for basic monitoring" on page 52.
Starting and stopping resiliency groups for maintenance	See "Starting a resiliency group" on page 53. See "Stopping a resiliency group" on page 53.
Rehearsing disaster recovery	See "Performing the rehearsal operation" on page 72. See "Performing cleanup rehearsal " on page 73.
Migrating a resiliency group	See "Migrating a resiliency group of virtual machines" on page 75.
Taking over resiliency groups	See "Taking over a resiliency group of virtual machines" on page 76.
Performing the resync operation	See "Performing the resync operation" on page 77.
Managing activities and resiliency plans	See "Managing activities" on page 105.
Displaying an overview of your resiliency domain including the number and health of your resiliency groups	See "About the Resiliency Platform Dashboard" on page 91. See "Displaying resiliency group information and status" on page 54.
Monitoring risks for protected assets	See "About risk insight" on page 94.
Viewing reports	See "Viewing reports" on page 103.

About managing Hyper-V virtual machines using Resiliency Platform

You can use Veritas Resiliency Platform to manage and protect your Hyper-V virtual machines configured in the resiliency domain.

Note: Make sure you enable all integration services for Hyper-V virtual machines.

The unit of management and control in Veritas Resiliency Platform is the resiliency group. Related virtual machines are organized into a resiliency group and managed and protected as a single entity.

Resiliency Platform supports workload management (start and stop) operations and recovery operations on resiliency groups.

Workload management lets you perform the tasks required for routine maintenance activities, for example, stop a resiliency group, update the required software components, and then restart the resiliency group.

If you configure a resiliency group for disaster recovery, you can perform tasks such as migrate your resiliency group to another data center, or perform the rehearse operation on the resiliency group.

You can set up replication using Hyper-V Replica or storage array (for example, EMC Symmetrix).

The detailed information about resiliency group management, virtual machine disaster recovery operations, and supported replication technologies is provided in the subsequent chapters of this guide.

See "Protecting Hyper-V virtual machines using array-based replication - an overview" on page 20.

About permissions for operations in the console

Users that are configured for Resiliency Platform have permission by default to view the web console but not to perform any operations. Permissions for operations must be assigned separately by a Resiliency Platform administrator, who assigns the appropriate personas to users or groups. A persona is a role with access to a set of operations. The administrator can further limit the scope of some operations by selecting the objects, such as resiliency groups, to which the user has access.

For example, an administrator can assign one user the permission to perform operations on resiliency group RG1 and assign another user the permission to perform operations on RG2. If more resiliency groups are added later, the

administrator needs to update permissions to assign access to the new resiliency groups.

Some objects, such as resiliency plans or virtual business services, can include multiple resiliency groups. To perform an operation on such an object, a user must have access to all its resiliency groups. Otherwise, the operation fails.

For more information on setting up user access to operations, refer to the *Deployment Guide*.

Overview of Amazon Web Services

This chapter includes the following topics:

- About AWS support in Resiliency Platform
- Using Amazon Web Services for remote recovery- an overview

About AWS support in Resiliency Platform

Veritas resiliency Platform 2.1 introduces the support for recovery of your assets to AWS environment.

Using Veritas Resiliency Platform 2.1, you can configure and protect your VMware and Hyper-V virtual machines for recovery to AWS using the Resiliency Platform Data Mover. You will need one license for recovery and one license for Resiliency Platform Data Mover.

Using Amazon Web Services for remote recovery- an overview

The following is a summary of the steps that are required to configure and protect your assets for recovery in Amazon Web Services (AWS) and where to go for more information on each step.

Table 2-1   Process overview

Step	More information
Download and deploy the appropriate Resiliency Platform virtual appliances for the following components: ■ In cloud: Resiliency Manager, IMS, and Replication Gateway ■ On-premises: IMS and Replication Gateway	For more information refer to the Deployment guide.
Configure the virtual appliances as Resiliency Platform components	For more information refer to the Deployment guide.
Set up the resiliency domain and add cloud configuration using the Getting Started wizard in the web console	For more information refer to the Deployment guide.
Configure the settings for the resiliency domain	For more information refer to the Deployment guide.
Add the asset infrastructure: ■ Add virtualization servers (vCenter server, Hyper-V server) ■ Prepare host for replication ■ Create gateway pair ■ Network customization (optional)	For more information refer to the Deployment guide.
Create resiliency groups for the virtual machines to be managed	See "Managing virtual machines for basic monitoring" on page 52. See "Managing virtual machines for remote recovery (DR) in Amazon Web Services" on page 67.
(Optional) Implement custom resiliency plans	See "Creating a new resiliency plan" on page 85.
(Optional) Configure virtual business services	For more information refer to the Solutions Guide for Virtual Business Services.

Table 2-1   Process overview *(continued)*

Step	More information
Perform disaster recovery operations.	See "Performing the rehearsal operation" on page 72.
	See "Performing cleanup rehearsal " on page 73.
	See "Migrating a resiliency group of virtual machines" on page 75.
	See "Taking over a resiliency group of virtual machines" on page 76.
	See "Performing the resync operation" on page 77.

Section 2

Preparing your environment

Using array-based replication

This chapter includes the following topics:

Protecting Hyper-V virtual machines using array-based replication - an overview

This section lists the key steps required to configure and perform the disaster recovery of Hyper-V virtual machines using array-based replication.

Table 3-1   Configure and perform disaster recovery using array-based replication

Action	Description	Refer to
Set up your replication environment	Complete the configuration steps required to set up Hyper-V Replica or array-based replication environment	See "Configuring Hyper-V virtual machines for disaster recovery using Hyper-V Replica" on page 22. See "Configuring Hyper-V virtual machines for disaster recovery using EMC SRDF" on page 26. See "Configuring Hyper-V virtual machines for disaster recovery using EMC RecoverPoint" on page 30. See "Configuring Hyper-V virtual machines for disaster recovery using Hitachi TrueCopy Universal Replicator" on page 32. See "Configuring Hyper-V virtual machines for disaster recovery using HPE 3PAR Remote Copy" on page 36.
Add the asset infrastructure	Add the virtualization servers to Resiliency Platform For array-based replication, add the storage arrays or the discovery hosts as required to Resiliency Platform	Refer to the Deployment Guide.
Prepare the virtual machines	Ensure that prerequisites are configured on virtual machines	See "Prerequisites for configuring Hyper-V virtual machines for disaster recovery" on page 62. See "Limitations for virtual machine disaster recovery" on page 63.
Configure network settings	Configure network settings for mapping between data centers	See "Setting up network mapping between production and recovery data centers" on page 46.
Configure your assets for disaster recovery	Group the virtual machines in a resiliency group and apply the appropriate service definition	See "Managing virtual machines for remote recovery (DR) using 3rd party replication technology" on page 64.

Table 3-1   Configure and perform disaster recovery using array-based
replication *(continued)*

Action	Description	Refer to
Rehearse DR operations	Test your disaster recovery environment to ensure DR readiness	See "Performing the rehearsal operation" on page 72. See "Performing cleanup rehearsal " on page 73.
Disaster recovery operations	Perform the disaster recovery operations	See "Migrating a resiliency group of virtual machines" on page 75. See "Taking over a resiliency group of virtual machines" on page 76. See "Performing the resync operation" on page 77.

Configuring Hyper-V virtual machines for disaster recovery using Hyper-V Replica

Prerequisites

To replicate data using Hyper-V Replica, you need to first configure Hyper-V Replica in your environment.

Refer to Microsoft documentation for configuring Hyper-V replica with and without Microsoft Failover clustering.

Hyper-V Replica with Microsoft Failover Clustering

This section lists the pre-requisites to enable data replication using Hyper-V Replica when the hosts are part of a Microsoft failover cluster.

■ Enable the Hyper-V and Failover Cluster roles on the Windows Server 2012 R2 hosts at the production and recovery (DR) data centers.

■ Ensure that the Microsoft failover cluster is already created using the required nodes at the production and recovery data centers.

■ In case of certificate-based authentication, ensure that you have specified broker-level certificate in the **Hyper-V Replica broker replication settings**.

■ Ensure that the cluster shared volumes (CSVs) are present at either the production or the recovery data centers' Microsoft failover clusters.

■ Ensure that all virtual machines are part of the Microsoft failover cluster. The data and configuration of the virtual machines will be stored on the cluster shared volume.

- Ensure that Hyper-V Replica Broker is configured on a node of the Microsoft failover clusters at the production and recovery data centers.

- On the Replica broker replication settings page, ensure to add the remote site broker details.

- Ensure that replication is already enabled for the virtual machines at the primary site.

- On Hyper-V host, verify whether the replication state for a given virtual machine is **Replicating**.

Note: Disaster recovery for cloned virtual machines in Hyper-V replica is currently not supported.

Once you have performed the necessary configurations, proceed with Resiliency Platform specific tasks.

Resiliency Platform configurations:

- Add Hyper-V virtual machines under respective data center using the **+ Hyper-V Server** option in **Settings** > **Infrastructure** > **Virtualization**. Provide domain user's credentials if Hyper-V Replica is managed by a domain user. Note that the primary and remote hosts must not be the part of the same data center.

Hyper-V Replica without Microsoft Failover Clustering

This section lists the pre-requisites to enable data replication using Hyper-V Replica when Microsoft Failover Clustering is not used.

- Enable the Hyper-V role on the Windows Server 2012 R2 hosts at the production and recovery (DR) data centers.

- Ensure to create and configure the required virtual machines on the host at the production data center.

- Ensure to enable the replication for the required virtual machine using the Hyper-V Manager at the production data center. It replicates the virtual machine boot disk (.VHDX) to the recovery data center.

Once you have performed the necessary configurations, proceed with Resiliency Platform specific tasks.

Resiliency Platform configurations:

- Add Hyper-V virtual machines under respective data center using the **+ Hyper-V Server** option in **Settings** > **Infrastructure** > **Virtualization**. Provide domain user's credential if Hyper-V Replica is managed by a domain user.

See "Protecting Hyper-V virtual machines using array-based replication - an overview" on page 20.

Configuring Hyper-V virtual machines for disaster recovery using EMC SRDF

EMC SRDF with Microsoft Failover Clustering

This section lists the pre-requisites to enable data replication using EMC SRDF when the hosts are part of a Microsoft failover cluster. For EMC SRDF-based replication, all virtual machines consuming storage from a consistency group must belong to the same resiliency group.

- Ensure that EMC Solutions Enabler (version v7.4, or later) is installed on a host and the SRDF device groups are already set up for replication between the primary and remote arrays.

- Ensure that EMC Symmetrix Gatekeeper device is presented to the array control host. You can designate any host, including the Hyper-V Server, as the array control host.

Note: The SRDF R1 and R2 LUNs must be on different hosts from different data centers.

- Ensure to enable the Hyper-V and Failover Cluster roles on the Windows Server 2012 R2 hosts at the production and recovery (DR) data centers.

- Ensure that you have created Microsoft failover cluster using the required nodes at the production and recovery data centers.

- Create virtual machines on the primary data center's Microsoft failover cluster with their data disks (`.vhdx`) on the replicated CSVs. In order to share the virtual machine configurations between cluster nodes, make sure to have another CSV (non-replicated). The user must set default virtual machine location to point to the non replicated CSV.
  Ensure that you have all the integration services enabled for these virtual machines.

- Ensure to create virtual machines in the Microsoft failover cluster at the production data center.

- Ensure that the Hyper-V virtual network switch name that is used by the replicated virtual machines, is same across all the Hyper-V hosts.

Once you have performed the necessary configurations, proceed with Resiliency Platform specific tasks.

Resiliency Platform configurations:

■ Add Microsoft Windows 2012 R2 host using the **+ Hyper-V Server** option in **Settings** > **Infrastructure** > **Virtualization**.

■ Discovery hosts can be added using the **+ Discovery Host** option in **Settings** > **Infrastructure** > **Storage** > **EMC** tab.
   Add the array control host where the SRDF device groups are configured, to the each IMS using the **Add Hosts** operation.

■ Add EMC Symmetrix enclosure using the **+ EMC Enclosure** option. Provide the discovery host name and the SYMCLI location on the discovery host. This operation returns the list of all Symmetrix arrays (local and remote) that are accessible to the host. To configure disaster recovery for the virtual machines, select only local arrays. This step is optional.

Default SymCLI location on Linux host   /opt/emc/SYMCLI/bin/

Default SymCLI location on Windows host   C:\Program Files\EMC\SYMCLI\bin

- Perform add discovery host and add enclosure (optional) operations at the disaster recovery data center as well.

Limitations:

- EMC SRDF LUN-based replication (without device group) and replication in the adaptive copy mode are not supported.

- Raw Device Mapping (RDM) on replicated disks to virtual machine is not supported.

- The rehearsal operation for resiliency groups using EMC Symmetrix Timefinder SNAP is not supported in Microsoft Failover Cluster environment.
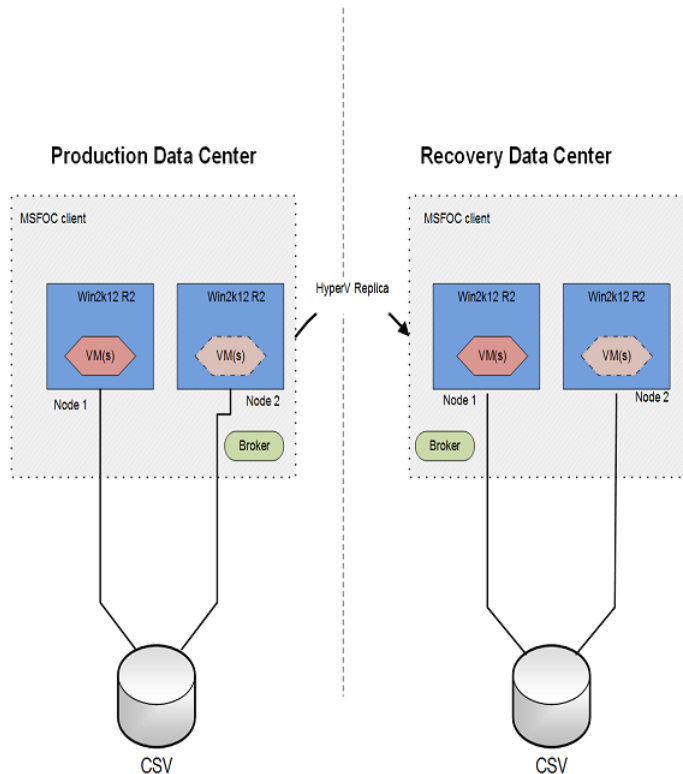
EMC SRDF without Microsoft Failover Clustering

This section lists the pre-requisites to enable data replication using EMC SRDF when the hosts are not part of a Microsoft failover cluster.

- Ensure that EMC Symmetrix Solutions Enabler (version v7.4, or later) is installed on the hosts and SRDF device groups are already set up for the replication between the primary and remote arrays.

- Ensure that EMC Symmetrix Gatekeeper device is present on the array control host. You can designate any host, including the Hyper-V Server, as the array control host.

   Note: The replicated and primary LUNs must be on different hosts from different data centers.

- Ensure that you have created the volumes on the primary Hyper-V host where the LUNs are read/write enabled.
   Veritas Resiliency Platform supports only one volume per replicated disk. We do not support Windows Storage Space Storage Pool.

- Ensure that you have created virtual machines at the primary data center under the Hyper-V Manager and kept their data files (`.vhdx`) on the replicated volumes. Do not keep their configuration files (.xml) on any replicated drive. Also, make sure that the default virtual machine configuration location under **Hyper-V Manager Settings** is not on a replicated drive.

- Ensure the respective remote disks (Read only - R2 remote disk and snapshot) are in the offline state on the Hyper-V server at the DR data center. And also verify that no drive letter is assigned to the volumes present on these offline disks.

> **Note:** For the support of rehearsal operation, you must add the snapshot devices to the DR data center's SRDF device group, and thereafter map them to the DR data center's Hyper-V hosts.

Once you have performed the necessary configurations, proceed with Resiliency Platform specific tasks.
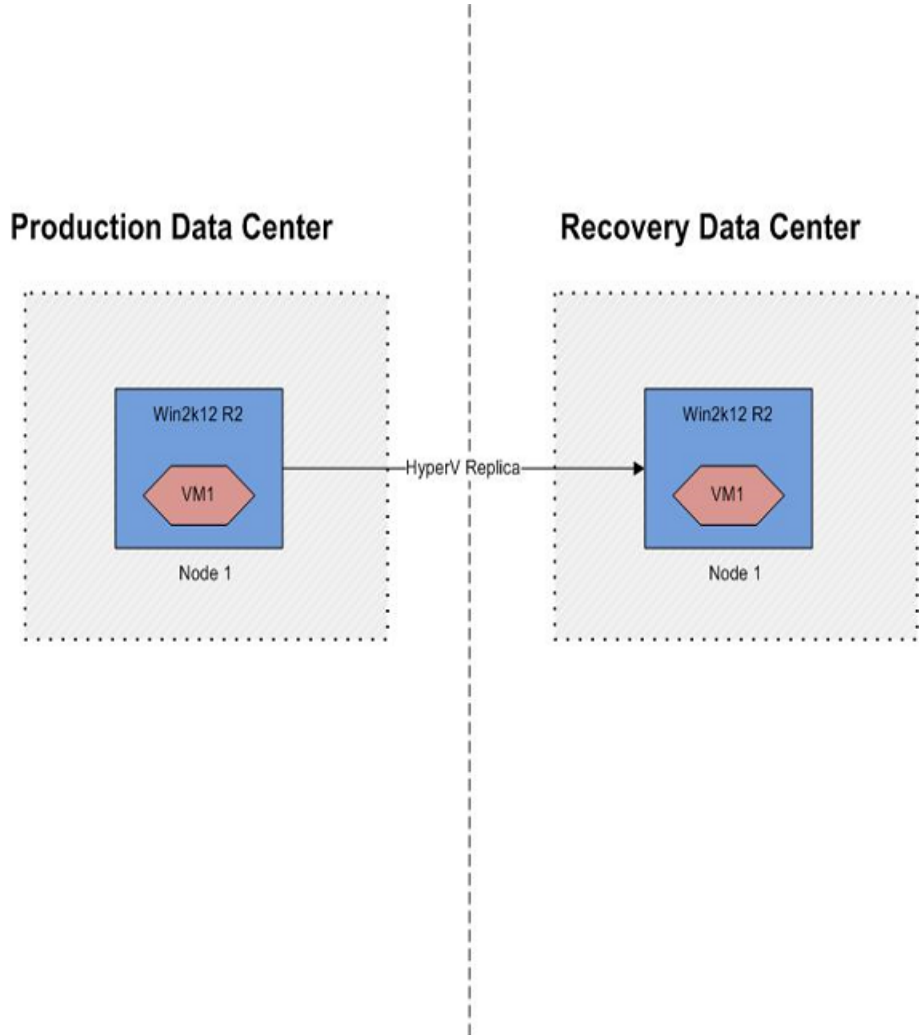
Veritas Resiliency Platform configurations:

- Add the host where the SRDF device groups are configured, using the **+ Discovery Host** option in **Settings** > **Infrastructure** > **Storage** > **EMC** tab.

- Add Symmetrix enclosure using the **+ EMC Enclosure** option. Provide the discovery host name and the SYMCLI location on the discovery host. This operation returns the list of all Symmetrix arrays (local and remote) that are accessible to the host. To configure disaster recovery for the virtual machines, select one or more local arrays.

| Default SymCLI location on Linux host | /opt/emc/SYMCLI/bin/ |
| Default SymCLI location on Windows host | C:\Program Files\EMC\SYMCLI\bin |

- Perform add discovery host and add enclosure operations at the disaster recovery data center as well.

Limitations:

- Logical grouping of disks (Windows Server Storage space storage pool) is not supported.
- Raw Device Mapping (RDM) on replicated disks to virtual machines is not supported.

Configuring Hyper-V virtual machines for disaster recovery using EMC RecoverPoint

EMC RecoverPoint with Microsoft Failover Clustering

This section lists the pre-requisites to enable data replication using EMC RecoverPoint when the hosts are part of a Microsoft failover cluster. For EMC RecoverPoint-based replication, all virtual machines consuming storage from a consistency group must belong to the same resiliency group.
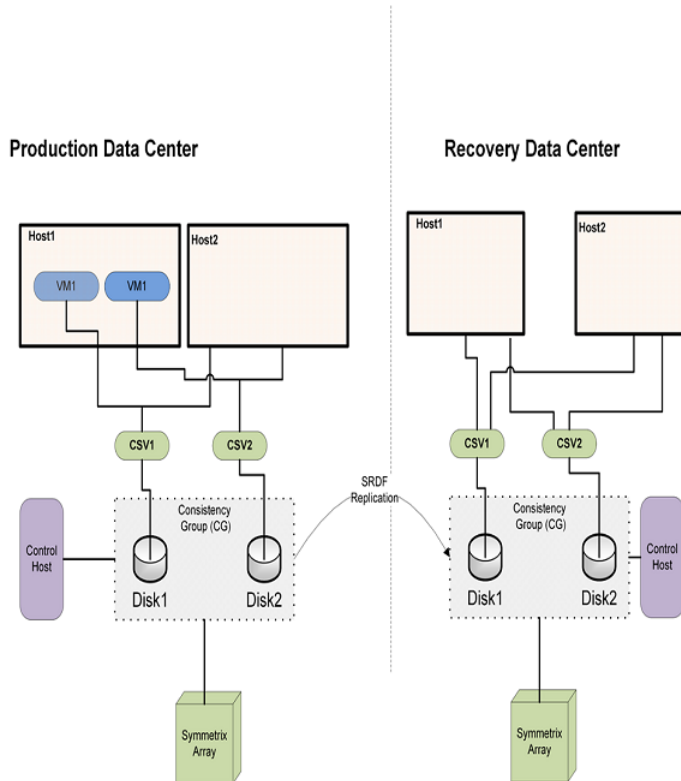
Prerequisites:

- Ensure that you are using EMC RecoverPoint 4.1.
- Discovery host should be able to communicate with Recoverpoint using SSH.

Once you have performed the necessary configurations, proceed with Resiliency Platform specific tasks.

Resiliency Platform configurations:

- Add Microsoft Windows 2012 R2 host using **Add Hosts** option under Infrastructure Management Server (IMS).
- Add the array control host where the RecoverPoint device groups are configured, to the each IMS using the **Add Hosts** operation.
- Add EMC Symmetrix enclosure using the **Add Enclosure** option. Provide the discovery host name and the SYMCLI location on the discovery host. This operation returns the list of all Symmetrix arrays (local and remote) that are accessible to the host. To configure disaster recovery for the virtual machines, select only local arrays.

Default SymCLI location on Linux host   /opt/emc/SYMCLI/bin/

Default SymCLI location on Windows host   C:\Program Files\EMC\SYMCLI\bin

- Perform add host and add enclosure operations for the IMS at the disaster recovery data center as well.

Limitations:

- Continuous data protection (CDP) and concurrent local and remote (CLR) data protection are not supported.

- Consistency group configured for more than 2 copies is not supported.

EMC RecoverPoint without Microsoft Failover Clustering

This section lists the prerequisites to enable data replication using EMC RecoverPoint when the hosts are not part of a Microsoft failover cluster.

Prerequisites:

- Ensure that you are using EMC RecoverPoint 4.1.

- Discovery host should be able to communicate with Recoverpoint using SSH.

Once you have performed the necessary configurations, proceed with Resiliency Platform specific tasks.

Veritas Resiliency Platform configurations:

- Add the host where the RecoverPoint device groups are configured, to the Infrastructure Management Server (IMS) using the **Add Hosts** operation.

- Add Symmetrix enclosure using the **Add Enclosure** option. Provide the discovery host name and the SYMCLI location on the discovery host. This operation returns the list of all Symmetrix arrays (local and remote) that are accessible to the host. To configure disaster recovery for the virtual machines, select one or more local arrays.

Default SymCLI location on Linux host   /opt/emc/SYMCLI/bin/

Default SymCLI location on Windows host   C:\Program Files\EMC\SYMCLI\bin

- Perform add host and add enclosure operations for IMS at the disaster recovery data center as well.

Limitations:

- Continuous data protection (CDP) and concurrent local and remote (CLR) data protection are not supported.

■ Consistency group configured for more than 2 copies is not supported.

Configuring Hyper-V virtual machines for disaster recovery using Hitachi TrueCopy Universal Replicator

Hitachi TrueCopy/Universal Replicator with Microsoft Failover Clustering

This section lists the pre-requisites to enable data replication using Hitachi TrueCopy/Universal Replicator when the hosts are part of a Microsoft failover cluster. For Hitachi TrueCopy/Universal Replicator based replication, all virtual machines consuming storage from a device group must belong to the same resiliency group.

Prerequisites:

■ Ensure that HORCM is installed on a host and the Hitachi TrueCopy/Universal Replicator device groups are already set up for replication between the primary and remote arrays.

■ Ensure that Hitachi Command Device is presented to the array control host. You can designate any host, including the Hyper-V Server, as the array control host.
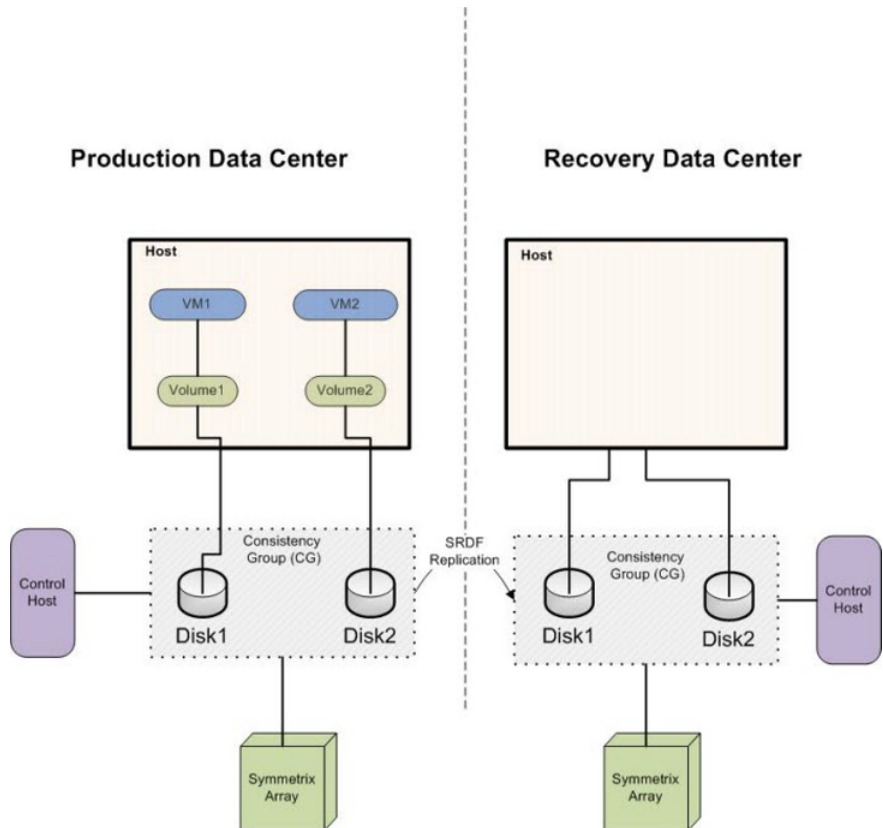
Note: The TrueCopy/HUR PVOLs & SVOLs must mapped to different hosts from different data centers.

■ Ensure to enable the Hyper-V and Failover Cluster roles on the Windows Server 2012 R2 hosts at the production and recovery (DR) data centers.

■ Ensure that you have created Microsoft failover cluster using the required nodes at the production and recovery data centers.

■ Ensure that you have created the cluster shared volume (CSV) on the replicated shared disk (R1) at the primary data center's Hyper-V Server.

■ Create virtual machines on the primary data center's Microsoft failover cluster with their data disks (.vhdx) on the replicated CSVs. In order to share the virtual machine configurations between cluster nodes, make sure to have another CSV (non-replicated). The user must set default virtual machine location to point to the non replicated CSV.

■ Ensure that you have all the integration services enabled for these virtual machines.

- Ensure to create virtual machines in the Microsoft failover cluster at the production data center.

- Ensure that the Hyper-V virtual network switch name that is used by the replicated virtual machines, is same across all the Hyper-V hosts.

Once you have performed the necessary configurations, proceed with Resiliency Platform specific tasks.

Veritas Resiliency Platform configurations:

- Add Microsoft Windows 2012 R2 host using Add Hosts option under Infrastructure Management Server (IMS).

- Add the array control host where the TrueCopy/HUR device groups are configured, to the each IMS using the Add Hosts operation.

- Add HiCommand Device Manager using the Add Enclosure option. Provide the discovery host name.

- Perform add host and add enclosure operations for the IMS at the disaster recovery data center as well.

Limitations:

- Raw Device Mapping (RDM) on replicated disks to virtual machine is not supported.

Hitachi TrueCopy/Universal Replicator without Microsoft Failover Clustering

This section lists the pre-requisites to enable data replication using Hitachi TrueCopy/Universal Replicator when the hosts are not part of a Microsoft failover cluster.

Prerequisites:

- Ensure that Hitachi Command Device is presented to the array control host. You can designate any host, including the Hyper-V Server, as the array control host.

Note: The replicated and primary LUNs must be mapped to different hosts from different data centers.

- Ensure that you have created the volumes on the primary Hyper-V host where the LUNs are read/write enabled. Veritas Resiliency Platform supports only one volume per replicated disk. Veritas Resiliency Platform does not support Windows Storage Space Storage Pool.

- Ensure that you have created virtual machines at the primary data centre under the Hyper-V Manager and kept their data files (.vhdx) on the replicated volumes. Do not keep their configuration files (.xml) on any replicated drive. Also, make sure that the default virtual machine configuration location under Hyper-V Manager Settings is not on a replicated drive.

- Ensure the respective remote disks (TC SVOL disk and ShadowImage SVOL) are in the offline state on the Hyper-V server at the DR data center. And also verify that no drive letter is assigned to the volumes present on these offline disks.

Note: For the support of rehearse operation, you must add the snapshot devices to the DR data center's SRDF device group, and thereafter map them to the DR data center's Hyper-V hosts.

Once you have performed the necessary configurations, proceed with Resiliency Platform specific tasks.

Veritas Resiliency Platform configurations:

- Add the host where the HTC/HUR device groups are configured, to the Infrastructure Management Server (IMS) using the Add Hosts operation.

- Add HiCommand Device Manager using the Add Enclosure option. Provide the discovery host name.

- Perform add host and add enclosure operations for IMS at the disaster recovery data center as well.

Limitations:

- Logical grouping of disks (Windows Server Storage space storage pool) is not supported.

- Raw Device Mapping (RDM) on replicated disks to virtual machine is not supported.

Configuring Hyper-V virtual machines for disaster recovery using HPE 3PAR Remote Copy

3PAR Remote Copy with Microsoft Failover Clustering

This section lists the pre-requisites to enable data replication using HPE 3PAR Remote Copy when the hosts are part of a Microsoft failover cluster. For HPE 3PAR Remote Copy based replication, all virtual machines consuming storage from a HPE 3PAR Remote Copy group must belong to the same resiliency group.

Prerequisites:

- Ensure that the discovery host is able to communicate with HPE 3PAR array using SSH.

- Confirm that HPE 3PAR array user has all the necessary permissions to perform HPE 3PAR RemoteCopy operations.

- HPE Remote Copy groups are set up for replication between the primary and the remote arrays. Ensure that the group names are unique across all data centers. Group name on the recover data center is auto generated by HPE. Do not modify the name.

- Ensure to enable the Hyper-V and Failover Cluster roles on the Windows Server 2012 R2 hosts at the production and recovery (DR) data centers.

- Ensure that you have created Microsoft failover cluster using the required nodes at the production and recovery data centers.

- Ensure that you have created the cluster shared volume (CSV) on the replicated shared disk at the primary data center's Hyper-V Server.

- Create virtual machines on the primary data center's Microsoft failover cluster with their data disks (.vhdx) on the replicated CSVs.

- Ensure that you have all the integration services enabled for these virtual machines.

- Ensure to create virtual machines in the Microsoft failover cluster at the production data center.

■ Ensure that the Hyper-V virtual network switch name that is used by the replicated virtual machines, is same across all the Hyper-V hosts.

Once you have performed the necessary configurations, proceed with Resiliency Platform specific tasks.

Veritas Resiliency Platform configurations:

■ Add the 3PAR enclosure to the IMS using the Add enclosure operation.

Note: Any managed host can be designated as the array discovery host. This operation returns the list of 3PAR arrays (local and remote) accessible to the discovery host. To configure disaster recovery for the virtual machines, select one or more local arrays only. Ensure that the enclosure discovery is complete before proceeding with adding the FOC cluster hosts.

- Add Microsoft Windows 2012 R2 host using Add Hosts option under Infrastructure Management Server (IMS).

- Perform add host and add enclosure operations for the IMS at the disaster recovery data center as well.

Limitations:

- Raw Device Mapping (RDM) on replicated disks to virtual machine is not supported.

- HPE 3PAR Remote Copy synchronous replication is not supported.

- 3PAR storage connectivity via iSCSI is not supported.

3PAR Remote Copy without Microsoft Failover Clustering

This section lists the pre-requisites to enable data replication using HPE 3PAR Remote Copy when the hosts are not part of a Microsoft failover cluster.

Prerequisites:

- Ensure that the discovery host is able to communicate with HPE 3PAR array using SSH.

- Confirm that HPE 3PAR array user has all the necessary permissions to perform HPE 3PAR RemoteCopy operations.

- HPE Remote Copy groups are set up for replication between the primary and remote arrays.

- Ensure that you have created the volumes on the primary Hyper-V host where the LUNs are read/write enabled. Veritas Resiliency Platform supports only one volume per replicated disk. Veritas Resiliency Platform does not support Windows Storage Space Storage Pool.

- Ensure that you have created virtual machines at the primary data center under the Hyper-V Manager and kept their data files (.vhdx) on the replicated volumes.

- Ensure the respective remote disks are in the offline state on the Hyper-V server at the DR data center. And also verify that no drive letter is assigned to the volumes present on these offline disks.

Once you have performed the necessary configurations, proceed with Resiliency Platform specific tasks.

Veritas Resiliency Platform configurations:

■ Add the 3PAR enclosure to the IMS using the Add enclosure operation.

> **Note:** Any managed host can be designated as the array discovery host. This operation returns the list of 3PAR arrays (local and remote) accessible to the discovery host. To configure disaster recovery for the virtual machines, select one or more local arrays only. Ensure that the enclosure discovery is complete before proceeding with adding the Hyper-V hosts.

- Perform add host and add enclosure operations for IMS at the disaster recovery data center as well.

Limitations:

- Logical grouping of disks (Windows Server Storage space storage pool) is not supported.

- Raw Device Mapping (RDM) on replicated disks to virtual machine is not supported.

- HPE 3PAR Remote Copy synchronous replication is not supported.

- 3PAR storage connectivity via iSCSI is not supported.

Configuring Hyper-V virtual machines for disaster recovery using IBM SVC Global Mirror

IBM SVC Global Mirror without Microsoft Failover Clustering

This section lists the pre-requisites to enable data replication using IBM SVC Global Mirror when the hosts are not part of a Microsoft failover cluster

Prerequisites:

- Ensure that you have created the volumes on the primary Hyper-V host where the LUNs are read/write enabled. Veritas Resiliency Platform supports only one volume per replicated disk. We do not support Windows Storage Space Storage Pool.

- Ensure that you have created virtual machines at the primary data center under the Hyper-V Manager and kept the data files (.vhdx) on the replicated volumes. Do not keep the configuration files (.xml) on any replicated drive. Also, make sure that the default virtual machine configuration location under Hyper-V Manager Settings is not on a replicated drive.

- Ensure the respective remote disks (Read only - R2 remote disk and snapshot) are in the offline state on the Hyper-V server at the DR data center. And also

verify that no drive letter is assigned to the volumes present on these offline disks.

- Ensure that the IBM SVC Global Mirror replicated LUNs are assigned to the respective HyperV Servers.

Once you have performed the necessary configurations, proceed with Resiliency Platform specific tasks.

Veritas Resiliency Platform configurations

Using the Resiliency Platform console **Infrastructure** settings, you add the asset infrastructure for each data center (the production and recovery data centers). The following is a summary of the steps.

For more information, refer to the *Deployment Guide*.

- Add IBM SVC enclosure using the **+ IBM Enclosure** option. Provide the IBM SVC server, username, and password. Ensure that the enclosures are discovered successfully.

- Add Microsoft Windows 2012 R2 host using the **+ Hyper-V Server** option in **Settings** > **Infrastructure** > **Virtualization**.

- Perform add discovery host and add enclosure (optional) operations at the disaster recovery data center as well.

Limitations

Raw Device Mapping (RDM) on replicated disks to virtual machine is not supported.

Configuring Hyper-V virtual machines for disaster recovery using IBM XIV Remote Mirror

IBM XIV Remote Mirror without Microsoft Failover Clustering

This section lists the pre-requisites to enable data replication using IBM XIV Remote Mirror when the hosts are not part of a Microsoft failover cluster.

Prerequisites:

- Ensure that you have created the volumes on the primary Hyper-V host where the LUNs are read/write enabled. Veritas Resiliency Platform supports only one volume per replicated disk. We do not support Windows Storage Space Storage Pool.

- Ensure that you have created virtual machines at the primary data center under the Hyper-V Manager and kept the data files (.vhdx) on the replicated volumes.

Do not keep the configuration files (.xml) on any replicated drive. Also, make sure that the default virtual machine configuration location under Hyper-V Manager Settings is not on a replicated drive.

■ Ensure the respective remote disks (Read only - R2 remote disk and snapshot) are in the offline state on the Hyper-V server at the DR data center. And also verify that no drive letter is assigned to the volumes present on these offline disks.

Complete the following tasks before you proceed with the Resiliency Platform specific tasks.

■ Install IBM XIV Command Line Interface (XCLI) on a Windows or Linux host. This host acts as a discovery host for array IBM XIV enclosures.

■ Using the XCLI command, verify if the IBM XIV array is accessible and command execution is successful.
For more information, refer to the *Deployment guide*.

■ Ensure IMS is reachable from the host having IBM XCLI installed.

■ Ensure that the IBM XIV array user is having the role of an 'Administrator'.

Once you have performed the necessary configurations, proceed with Resiliency Platform specific tasks.

Veritas Resiliency Platform configurations

Using the Resiliency Platform console, **Infrastructure** settings, you add the asset infrastructure for each data center (the production and recovery data centers). The following is a summary of the steps.

For more information, refer to the *Deployment guide*.

■ Add the host having XCLI installed as a discovery host using the Resiliency Platform console.
To add a Windows host as a discovery host, you need to configure a Windows control host.

■ Add IBM XIV enclosures to the appropriate data centers using the **+ IBM Enclosure** option. Select the discovery host, provide the enclosure name or IP address, username, and password. Enter the full path of the folder having XCLI binary in CLI Location text box.
Ensure that the enclosures are discovered successfully.

■ Add Microsoft Windows 2012 R2 host using the **+ Hyper-V Server** option in **Settings** > **Infrastructure** > **Virtualization**.

Managing disaster recovery network mapping

This chapter includes the following topics:

- Viewing and configuring network settings for a data center

- Editing network settings for a data center

- Removing network settings for a data center

- Configuring DNS server settings for a data center

- Setting up network mapping between production and recovery data centers

Viewing and configuring network settings for a data center

Using the Resiliency Platform console, you can view the details of the discovered subnets, V-Switches, and VLANs and also add new subnets.

Information of the discovered or added networks such as name, IP address of the gateway, vServer name, type, purpose etc is displayed in the table.

While adding a new subnet you need to choose from one of the following purposes:

- **Production**: Lets you perform the DR activities such as migrate and take over.

- **Rehearsal**: Lets you perform the rehearsal operation.

The add subnet wizard lets you create a subnet pair, but only if you choose the purpose as Production. This is optional. You need to enter a name, IP address of the network and the gateway for the Rehearsal network as well.

To configure network settings for a data center

1   Navigate

⚙   **Settings** (menu bar) > **Infrastructure** > **Details View**

Expand the data center > **Access Profile**

Click **+ Add Subnet**.

2   Select the purpose, enter a name, IP address of the network and the gateway.

To create a pair, you can either choose a subnet from the list, or click **+Add new**.

3   Select **Next** to review and confirm the selection.

Editing network settings for a data center

In the web console, you can edit the details of the discovered subnets, V-Switches, and VLANs.

You can create a network pair if you edit the purpose from test to production.

To edit the network settings for a data center

1   Navigate

⚙   **Settings** (menu bar) > **Infrastructure** > **Details View**

Expand the data center > **Access Profile**

Right-click the subnet, V-Switch, or VLAN and select **Edit**.

2   Edit the name and the purpose as required.

3   Select **Next** to review and confirm the selection.

Removing network settings for a data center

In the web console, you can remove the subnets that you have added. Subnets, V_Switches, and VLANs that are discovered cannot be deleted.

To remove the network settings for a data center

1 Navigate

⚙ **Settings** (menu bar) > **Infrastructure** > **Details View**

Expand the data center > **Access Profile**

Right-click the subnet and select **Remove**.

2 Review the selection and click **Submit**.

Configuring DNS server settings for a data center

Using the Resiliency Platform console, you can configure the DNS settings for the data center.

You can add DNS servers for the data center or remove the settings for servers that were previously added.

To configure DNS server settings for a data center

1 Prerequisites

You must have the following information:

- The IP address of the DNS server

- The name of the domain, and associated credentials.
  Linux Bind: For TSIG authentication, you need the TSIG key and TSIG private files.
  Windows DNS: For GSSAPI authentication, you need the user name and keytab file.

- A test host name and IP address for performing a test operation. The test operation validates the specified DNS configuration.

2 Navigate

⚙ **Settings** (menu bar) > **Infrastructure** > **Details View**

Expand the data center > **Access Profile**

Click the **Windows DNS** or **Bind** tab.

DNS servers already added for the data center are listed in the table. You can remove or add a new DNS server.

3 To add a new DNS server for the data center click **+ Add New DNS**.

4 Specify the IP address for the DNS server and select the purpose, either Rehearsal or Production.

5 Add one or more domains for the DNS server:

- Fill in the domain name and the authentication type. For TSIG, browse to the key and private files. For GSSAPI, enter the user name and browse to the keytab file.

- Enter a test host name and IP address and select **Test**. If the test is successful, that is the DNS configuration is validated, the **Add** button is enabled.

- Select **Add**.

6 If you are done adding domains, select **Next**.

7 To remove a DNS server, right-click the required DNS server in the table and select **Remove**.

Setting up network mapping between production and recovery data centers

The network mapping operation eliminates the need to manually apply an IP address for each virtual machine at the recovery (DR) data center. After you have mapped the networks successfully, the IP addresses are computed programmatically, and applied to the virtual machines.

For Hyper-V virtual machines, ensure that the mapping of all the concerned virtual switches across the data centers is configured before performing migrate, takeover, or rehearsal operations.

Else, network adapters of the virtual machines are not connected to any network after the operation. Similarly, ensure that the subnets are mapped across the data centers when IP customization is required.

If subnets are mapped and IP customization option is selected during the DR operation and if the port groups are not mapped, then IP customization fails for the concerned network adapters, causing the DR operations to fail.

This is not applicable if the recovery data center is Amazon Web Services (AWS). For AWS, subnet to subnet mapping is sufficient.

Note that the subnets are discovered only when the virtual machines are running.

If the recovery data center is AWS, then ensure that the production subnet and the rehearsal subnet are in the same virtual private cloud (VPC).

> **Note:** When you clone your virtual machines, ensure that you assign the appropriate host name and IP address to the cloned virtual machines.

To set up network mapping between production and recovery data centers

1   Navigate

 Disaster Recovery Settings (navigation pane)

Do one of the following:

- On **Overview** tab, click **+ New Network Pair**.
  On **Network** tab, click **+ Create Pair**.

Previously created network pairs are listed in the table. You can create a new pair or delete an existing pair.

2   In the **Network Mapping** page, select the source and the target data centers, and the network types that should be the part of your network pair.

If recovery is in AWS, select subnets.

3   Click **Choose selected** or drag and drop the selections in the drag area at the bottom.

4   Click **Next** to submit your selections.

5   To remove a network pair, right-click the pair and select **Delete Pair**.

You cannot edit a network pair. Instead you need to delete the pair and create another.

Section 3

Working with resiliency groups

Managing resiliency groups

This chapter includes the following topics:

About resiliency groups

Resiliency groups are the unit of management and control in Veritas Resiliency Platform. After assets are added to Resiliency Platform, you organize related assets into a resiliency group that you can protect and manage as a single entity.

For example, you can organize several applications into a resiliency group and name it `SQL_Server_Group`. Then, when you perform an operation on `SQL_Server_Group` from the Resiliency Platform console, all the applications in the group are affected. For example, if you start `SQL_Server_Group`, all the applications

in the group start. Similarly, you can organize virtual machines into a resiliency group and perform operations that affect all the virtual machines in the group.

Note: A resiliency group must contain similar types of objects, either all applications or all virtual machines. It cannot contain a mix of the two.

The operations available for a resiliency group depend on how it is configured. During the configuration of a resiliency group, you apply a service objective that identifies the objective or intent for that group of assets. If you apply a service objective that supports remote recovery, the resiliency group supports operations like migrate and take over.

You can optionally use a service objective that only monitors the assets and provides only basic operation capabilities like start and stop operations and no remote recovery operations.

See "About service objectives" on page 50.

See "Managing virtual machines for basic monitoring" on page 52.

Guidelines for organizing resiliency groups

Resiliency groups are most useful when the assets in the group share common characteristics.

While creating a resiliency group of virtual machines, follow these guidelines for selecting virtual machines:

- Ensure that all the virtual machines that are to be grouped in a single resiliency group are from a single hypervisor or virtualization server (if not clustered) or a single cluster.

- Ensure that they consume storage from the same Veritas Replication Set. E.g. EMC SRDF device group, NetApp Volume, 3PAR replication group, and so on.

About service objectives

Service objectives define the type of protection to be applied to a group of data center assets. For example, an option for remote recovery which allows assets being managed by a resiliency group to be recovered at a remote location (DR) using a service objective can include operations such as migrate or take over. Whereas the monitor assets service objective lets you start or stop your assets within the resiliency group.

The remote recovery service objective includes tunables such as Recovery Point Objective (RPO) for assets being managed in that resiliency group and you would be required to select the recovery data center.

Service objectives are provided as templates that must be activated before use. A set of pre-activated service objectives with default settings are provided.

Following is the list of service objective templates:

- Remote recovery of applications - provides recovery operations as well as the start and stop operations for applications.

- Remote recovery of hosts - provides recovery operations as well as the start and stop operations for hosts.

- Monitor assets - provides only monitoring, that is start and stop operations.

For virtual machines you have the following two options for data availability.

- Copy: The available technology is NetBackup. This option is available only for VMware virtual machines.
  This option is available only if the acceptable RPO is 240 minutes (4 hours) and above.

- Replication: The available technologies are SnapMirror, SRDF, VRP Data Mover, RemoteCopy 3PAR, RecoverPoint, Hyper-V Replication, and Hitachi True Copy.

Note: Authorization to activate a template and edit the settings depends on the permissions that are assigned to users and groups in Resiliency Platform.

Following is the list of pre-activated service objectives:

- Recover hosts

- Recover applications

- Monitor assets

- Recover hosts using data copies (remote)

You can view the details of both the activated service objectives and the templates in the web console. You can also delete any pre-activated service objective that you do not want to use in your environment, provided that it is not in use by any resiliency group.

The default pre-activated service objectives do not monitor an RPO. You can change this setting by activating a service objective so that Resiliency Platform can alert you of a failure to meet an RPO.

For more information on customizing service objectives, refer to the Deployment Guide.

When you create a resiliency group of assets in Veritas Resiliency Platform, you select a service objective to apply to that group of assets. The wizard then prompts you for any additional information that is needed to prepare the resiliency group for the supported operations.

Managing virtual machines for basic monitoring

When you create a resiliency group, you select a service objective that specifies the operations supported for that resiliency group.

There are two types of pre-activated service objectives:

■ Monitor assets - provides only monitoring, start, and stop operations

■ Recover hosts - provides recovery operations as well as the start and stop operations

This topic explains how to configure a resiliency group for basic monitoring.

Configuring a resiliency group for remote recovery has additional prerequisites and steps and is described in a separate topic.

To manage virtual machines for basic monitoring

1   Prerequisites

    The asset infrastructure must be added to Resiliency Platform and asset discovery must be complete.

    For more information on adding asset infrastructure, refer to the *Deployment Guide*.

2   Navigate

    Assets (navigation pane) > **Resiliency Groups** tab > **Manage & Monitor Virtual Machines or Applications**

    You can also launch the wizard from the **Unmanaged** or **Overview** tabs.

3   Select the virtual machines:

    ■ Select **Host** as the asset type, select the data center, and select other filters as needed to display a list of virtual machines.

    ■ Drag and drop virtual machines to **Selected Instances**.

4   The next page displays the environment for the selected assets.

5   Select the service objective that provides monitoring, start, and stop operations only.

6   Supply a name for the resiliency group.

7   Verify that the new resiliency group is added to the **Resiliency Groups** tab.

Optionally, use **Recent Activities** (bottom pane) > **Details** to view the details of this task in a graphical representation.

Starting a resiliency group

When you start a resiliency group, you start all the underlying assets in it.

To start a resiliency group

1   Navigate

        Assets (navigation pane) > **Resiliency Groups** tab

2   Locate your resiliency group. Use filters or Search as needed.

3   On the row for the resiliency group, select the vertical ellipsis > **Start**.

You can also perform operations from the Details page.

4   On the **Start Resiliency Group** wizard, select the data center in which to start the group, and submit.

If you have applied update 2.0.0.100 on Veritas Resiliency Platform 2.1, you can select the checkbox on the **Start Resiliency Group** wizard to start the post-replication operations of migrate or takeover workflow on the production data center such as refreshing storage, network, compute, and customization.

To display a record and a graphic representation of what you did, select the **Recent Activities** at the bottom of the page, find your task, and select **Details**.

5   If necessary, notify users after you start the resiliency group.

Stopping a resiliency group

When you stop a resiliency group, you stop all the assets that make up the group.

A typical reason for stopping a resiliency group would be to update or perform maintenance in one or more of the underlying assets.

To stop a resiliency group

1   Prerequisites

- Make sure that you are aware of all the assets in the resiliency group, and the potential effect on users if you shut them down.

- Choose a time for stopping the resiliency group that minimizes any disruption of service.

- If necessary, notify users before you stop the resiliency group.

2 Navigate

Assets (navigation pane) > **Resiliency Groups** tab

3 Locate the resiliency group. Use filters or Search as needed.

4 On the row for the resiliency group, select the vertical ellipsis > **Stop**.

You can also perform operations from the Details page.

5 On the **Stop Resiliency Group** screen, select the data center in which to stop the resiliency group, and submit.

To display a record and a graphic representation of what you did, select the **Recent Activities** at the bottom of the page, find your task, and select **Details**.

Displaying resiliency group information and status

You can display resiliency group information and status in the following ways:

Table 5-1   Displaying resiliency group information and status

Location	Level of detail	Useful for
Resiliency Platform Dashboard	Lowest. Displays the number of resiliency groups under Resiliency Platform control and the total number of groups in error, at risk, and healthy.	Getting a quick overview of the resiliency group population and health throughout Resiliency Platform. See "About the Resiliency Platform Dashboard" on page 91.

Table 5-1   Displaying resiliency group information and status *(continued)*

Location	Level of detail	Useful for
 Assets > **Resiliency Groups** tab	Medium. Lists all your resiliency groups in one place.	Seeing what is in each of your data centers, the state of the groups, and so on.
Resiliency group-specific screen	Highest. Lists each asset in the resiliency group, their type, and state.	Getting detailed information on a resiliency group and its underlying assets, including disaster recovery status. This screen lists available operations for the group. See "Viewing resiliency group details" on page 57.

This section discusses the second method of displaying resiliency group information and status: using the **Assets** page. The **Assets** page gives you a quick overview of all your resiliency groups.

To display resiliency group information and status

1   Navigate

   Assets (navigation pane) > **Resiliency Groups** tab

2   Review information and status

For a quick health check of your resiliency groups, review the colored boxes above the table. Select a box to show only the resiliency groups in that category; for example, select the green square to display only the resiliency groups that are healthy.

Blue	The total number of resiliency groups
Yellow	The number of resiliency groups at risk
Green	The number of resiliency groups that are healthy

By default, the table lists all resiliency groups. Use the drop-list and search field to filter your results, and select a table heading to sort the groups.

In the table, the key fields are **State**, **Service Objective**, and **Data Availability**. Possible states are:

Status	**Normal** - the assets within the resiliency group are normal.
	At Risk - the assets within the resiliency group are at risk.
State	**Online** - The assets within the resiliency group are running.
	Partial - One or more of the assets in the resiliency group are offline.
	Offline - The assets in the resiliency group are powered off or not running.
Active DC	Name of the active data center.
Type	Application Group: The resiliency group comprises of applications.
	Virtual Machine Group: The resiliency group comprises of virtual machines.
Service Objective	Service objective selected for the resiliency group.
Data Availability	Resiliency Platform supports several replication technologies.
	If no replication type is shown, consider configuring replication.

Viewing resiliency group details

Using the Resiliency Platform console, you can view detailed information on each of your resiliency groups. The overall health of the resiliency group, its underlying assets and their current state is displayed.

Resiliency group for which disaster recovery (DR) operation is configured successfully, you can view information which includes the state of the replication for the resiliency group (for example, synchronized), used replication technology, associated alerts, the details about the applications or the virtual machines in the resiliency group, replication lag, recovery time, and so on.

Note that for virtual machines, the recovery time is available only after the rehearse operation is complete.

To view details of a resiliency group

1 Navigate

Assets (navigation pane) > **Resiliency Groups** tab

2 Locate your resiliency group. Use filters and search as needed.

3 On the row for the resiliency group, select the vertical ellipsis > **Details**.

You can also double-click the row to view details.

The details page includes the following:

- Menu options for operations that you can perform on the resiliency group.

- Details of how the resiliency group is configured.

- Status information.

- A list of the resiliency group assets and their state.

See "Displaying resiliency group information and status" on page 54.

Editing a resiliency group

You can edit the resiliency group information including the group name as well as change the underlying assets on which the resiliency group is based when the resiliency group is configured for basic monitoring using the Monitor assets service objective.

If the resiliency group is already protected for DR, then the wizard proceeds with the DR configuration letting you make any changes if required.

If you add, remove, or grow a disk of a virtual machine that belongs to a resiliency group (which is DR protected), then the Resiliency Platform raises a risk. You then need to edit the resiliency group to first remove the virtual machine and then edit again to add the virtual machine.

To edit the resiliency group information

1   Navigate

              Assets (navigation pane) > **Resiliency Groups** tab

2   Locate the resiliency group. Use filters or Search as needed.

3   On the row for the resiliency group, select the vertical ellipsis > **Edit**.

              You can also edit the resiliency group from its Details page.

The steps for editing the resiliency group are the same as creating it.

Deleting a resiliency group

When you delete a resiliency group from Resiliency Platform management, you can no longer monitor, manage, or protect it using Resiliency Platform. Deleting the resiliency group from Resiliency Platform has no effect on the underlying assets.

If the recovery data center is in Amazon Web Services (AWS) cloud, then you can delete the resiliency group only when it is online on the production data center. Which means you can delete the group after successful completion of the configure for remote recovery operation or after you have migrated the assets from AWS cloud to the production data center. You cannot delete the group if you have migrated the assets from the production data center to AWS cloud and after running the take over operation.

To successfully complete the delete operation ensure the following:

■ The assets on the production data center are running and accessible.

■ The xprtld daemon on the virtual machines is running.

On successful completion of the delete operation, you will notice the following:

■ During the operation, replication is stopped and Veritas Replication Sets are deleted on gateways and on-premises virtual machines.

■ Journal disks are removed from the virtual machines on the production data center and cloud virtual machines instances are deleted.

- All the cloud virtual machines disks that are attached to the cloud Replication Gateway are deleted.

Note: Replication Gateway pairs are not deleted during the delete operation. If required you can delete the pair from the **Gateway Pair** details page.

To delete a resiliency group

1 Navigate

    Assets (navigation pane) > **Resiliency Groups** tab

2 Locate the resiliency group. Use filters or Search as needed.

3 On the row for the resiliency group, select the vertical ellipsis > **Delete**.

    You can also perform operations from the Details page

4 Confirm the deletion.

Configuring resiliency groups for remote recovery

This chapter includes the following topics:

- Understanding the role of resiliency groups in disaster recovery operations

- How Resiliency Platform configures disaster recovery protection for virtual machines

- Prerequisites for configuring Hyper-V virtual machines for disaster recovery

- Limitations for virtual machine disaster recovery

- Managing virtual machines for remote recovery (DR) using 3rd party replication technology

- Managing virtual machines for remote recovery (DR) in Amazon Web Services

Understanding the role of resiliency groups in disaster recovery operations

To perform disaster recovery (DR) operations on virtual machines or applications, they must be configured for disaster recovery as part of a resiliency group, which is the unit of management and control in Veritas Resiliency Platform.

In the configuration wizard for resiliency groups, you apply a service objective to a resiliency group. When you apply the recover hosts service objective, the wizard prompts you for the additional information required for Resiliency Platform to configure the resiliency group for disaster recovery operations.

After disaster recovery configuration on a resiliency group is complete, you can proceed with DR-specific tasks on the resiliency group, such as migrate and take over.

A Virtual Business Service (VBS) lets you further group these resiliency groups in a multi-tier grouping mechanism, and lets you perform controlled start, stop and recovery operations on these resiliency groups.

How Resiliency Platform configures disaster recovery protection for virtual machines

During the wizard configuration process, Resiliency Platform searches the complete storage stack from the virtual machines to the replicated volumes.

It also detects the complete network settings of each member of the resiliency group. If network mapping has been configured, it applies the mapping details to the network settings that need to be applied in the recovery data center after migration. The IP addresses for the virtual machines at the recovery data center are applied based on the subnet mappings. Resiliency Platform stores and uses this configuration at the time of disaster recovery operations, such as, Migrate, Takeover, or Rehearse. This network customization is applicable only if DHCP is not configured for the data center.

The wizard validates the DR configuration and displays the results. For example, the wizard can display the number of virtual machines that are needed at the recovery data center to match the number of virtual machines at the production data center.

When you configure a set of virtual machines in a resiliency group for DR, the Resiliency Platform saves some extra information about the virtual machines on the replicated storage. For VMware, the Resiliency Platform saves additional copies of the virtual machine configuration in the same folder as the original virtual machine configuration. For Hyper-V, the Resiliency Platform creates a folder with name "`vrp`" on the replicated mount point and stores additional copies of the virtual machine configuration in it. The Resiliency Platform maintains separate copies of the virtual machine configuration per data center, thus allowing you to have separate virtual machine configurations across data centers. These copies are used during the DR operations such as Migrate, Takeover, Rehearsals, etc. These files are maintained by the Resiliency Platform and should not be edited or deleted.

Note: If there are any changes to the storage stack or network settings in any of the resiliency group members, re-run the wizard so that the latest storage and network configuration snapshots are recorded.

For Hyper-V virtual machines, after the configuration snapshot is recorded and stored, Resiliency Platform copies the virtual machine configurations into a folder (.vrp) on the replicated datastore. This folder hosts data center-specific copies of the virtual machines. This allows the user to have separate configurations across the two data centers for the same virtual machines.

Prerequisites for configuring Hyper-V virtual machines for disaster recovery

Before you run the wizard to configure disaster recovery protection for a resiliency group of Hyper-V virtual machines, ensure that you have met the following prerequisites for virtual machine configuration:

- The virtualization servers for the virtual machines must be added to Resiliency Platform at the production and recovery data center.
  For more information on adding asset infrastructure, refer to the *Deployment Guide*.

- The virtual machines, that you use to create the resiliency group, should be online on one of the data centers (production or recovery). But the replication disk should be attached to both the data centers.

- If you add new disks, ensure that they are visible from the guest operating system.

- Ensure that the integrations services are installed and running inside the virtual machines.

- Ensure that you have disabled the 'Quick removal' policy for disks in Windows Server 2008 R2 and disabled the 'write-cache' policy for disks in Windows Server 2012 R2.

- If you want to do IP customization, ensure that the Hyper-V virtual machines are updated with the latest Hyper-V integration services.

- If the recovery data center is Amazon Web Services (AWS), then ensure the following:

  - The Paravirtual (PV) Drivers must be installed. Follow the documentation of AWS for more information.

  - Ensure that the Replication Gateways have sufficient storage to handle the replication for the planned number of protected virtual machines.
    Both the on-premises gateway and the cloud gateway must have external storage equivalent to 12GB for each asset protected by the gateway pair.
    For example, if a gateway pair supports 10 virtual machines, the on-premises gateway and the cloud gateway must each have 120 GB of external storage.

A maximum of 40 volumes or disks can be attached to the cloud Replication Gateway.

- On Windows hosts, initialize and reboot the disks that are attached to IDE controller.

- A bucket must be created in S3 and a policy must be created that assigns the following permissions:

  - s3:GetBucketLocation and s3:GetObject permissions on the bucket

  - ec2:ImportSnapshot, ec2:DescribeSnapshot, and ec2:CopySnapshot permissions on all the resources

  One role named ImportSnapshotRole must be created and the above policy is associated with this role. The service vmie.amazonaws.com should be able to assume this Role.

For more information about the permissions required, refer to the AWS documentation.

Additional prerequisites depend on the replication technology you are using.

See "Protecting Hyper-V virtual machines using array-based replication - an overview" on page 20.

Limitations for virtual machine disaster recovery

The following table lists the limitations of virtual machines disaster recovery using Resiliency Platform:

Table 6-1

Limitations	Descriptions
Replication limitations	For more information on replication-based limitations of virtual machines, refer to the Hardware and Software Compatibility List (HSCL).
The hypervisor should not be added as a host in certain environments	If the applications are installed inside virtual machines running on Microsoft Hyper V technology and the applications are having data that is replicated using EMC SRDF, and these applications are to be configured for DR, then you should not add the hypervisor itself as a host to the IMS.

Managing virtual machines for remote recovery (DR) using 3rd party replication technology

To provide disaster recovery protection, you organize virtual machines into a resiliency group and apply the remote recovery for hosts service objective. The wizard prompts for the inputs that are needed for the selected service objective and for the replication technology. The wizard then implements the configuration that is required for the DR operations.

To manage virtual machines for remote recovery (DR) using 3rd party replication technologyTo manage assets for remote recovery (DR) using 3PAR Remote Copy replication technology

1   Prerequisites

Ensure that you have completed the configuration prerequisites for your virtualization and replication environment.

See "Protecting Hyper-V virtual machines using array-based replication - an overview" on page 20.

2   Navigate

	Assets (navigation pane) > **Resiliency Groups** tab > **Manage & Monitor Assets**

You can also launch the wizard from the **Unmanaged** or **Overview** tabs.

3   Select the virtual machines:

■   Select **Host** as the asset type, select the data center, and select other filters as needed to display a list of virtual machines.

■   Drag and drop virtual machines to **Selected Instances**.

4   The next page displays the environment for the selected assets.

5   The next page lists the service objectives that are available for the selected asset type. You can expand the service objective to view details. Select the service objective that provides disaster recovery operations.

6   Select the target (recovery) data center.

7   Select the target Cluster or ESXi host. This panel is displayed only if the replication technology is NetApp SnapMirror. But if the storage is mounted using NFS protocol on both the data centers, then this panel is not displayed.

See "Target asset selection options" on page 65.

8   Complete the network customization steps for the virtualization technology.

See "Network customization options" on page 65.

9   Verify the summarized information and enter a name for the resiliency group.

When you finish the wizard steps, Resiliency Platform invokes a workflow which initializes the DR operation. You can view the progress or ensure that this operation is successfully completed on the **Activities** page.

See "Viewing activities" on page 105.

Verify that the new resiliency group is added to the **Resiliency Groups** tab.

See "Viewing resiliency group details" on page 57.

Target asset selection options

This panel is displayed only if the replication technology is NetApp SnapMirror or Resiliency Platform Data Mover.

Select the following in this panel:

- Select the target cluster.
  Each resiliency group must map to only one ESX cluster. The wizard validates which clusters on the target data center can meet the required number of virtual machines and disks.

- Review the hosts within the cluster.
  Once you select a cluster, the associated ESX hosts are displayed, and below them the datastores that are accessible from the ESX hosts.

- Select the datastore
  Storage that can be provisioned must be available in the selected datastore in order for the wizard to create the replicated disks on the target data center.
  Review the total disk size and compare this to the value in the **Free (GB)** column for the selected datastore.
  If there is insufficient memory, you can continue with the wizard and update the resources later.

Network customization options

You can do the following in this panel:

- Customize the static IP for virtual machines on the production and the recovery data center.

- Choose between Production and Rehearsal DNS customization.

- Manage PTR records

- Choose to continue with DR operations even if DNS updates fail.

Customizing the IPs of a virtual machine overrides the default IP settings when the virtual machine starts at the recovery data center. You can assign the static IPs to the protected virtual machine from site-specific subnets. The computation of projected static IP is done based on the subnet mappings.

Select the **Apply IP customization** option if you want to customize the IPs. You can choose to continue with the DR operation if the IP customization fails. Note that this is possible only if the virtual machines have static IPs. You need to double click on the IP that you want to edit.

Since only IPv4 is supported, you may see a warning if there are IPv6 address: *Unable to apply IP customization some of the workloads.* Ignore this warning.

If the recovery data center is in Amazon Web Services (AWS), then ensure that the IPs used for network customization are not already in use in AWS.

If you choose to apply DNS customization, then you can add a host name to the virtual machine.

See "Prerequisites for network customization" on page 66.

Prerequisites for network customization

Ensure the following prerequisites are met before you customised the IP addresses and the DNS settings.

- Linux and Windows virtual machines which are running on Hyper-V platform ensure that Hyper-V integration services are installed and running.

- IP, Gateway, Netmask, DNS, Domain Name, Mac address etc. information should present in the respective files of each network interface for which you want to customize the IP and DNS.

- If multiple network interfaces (NICs) are assigned to a virtual machine, then you need to apply IP customization to all the NICs.

- For virtual machines that are running on Linux ensure that NetworkManager and libvirtd service is in off state.

- Virtual machines which are running on Hyper-V platform ensure that FQDN is set for each virtual machine.

- The mac address configuration should be set as Manual/Static so that it does not change after the DR operation is performed.

Managing virtual machines for remote recovery (DR) in Amazon Web Services

Using the Resiliency Platform console, you can organize virtual machines into a resiliency group, apply the remote recovery for hosts service objective, and configure them for remote recovery in Amazon Web Services (AWS).

The wizard prompts for the inputs that are needed for the selected service objective and replication technology.

To manage virtual machines for remote recovery in AWS

1 Prerequisites

See "Prerequisites for configuring Hyper-V virtual machines for disaster recovery" on page 62.

2 Navigate

▦	**Assets** (navigation pane) > **Resiliency Groups** tab > **Manage & Monitor Assets** You can also launch the wizard from the **Unmanaged** or **Overview** tabs.

3 Select the virtual machines:

- Select **Host** as the asset type, select the data center, and select other filters as needed to display a list of virtual machines.

- Drag and drop virtual machines to **Selected Instances**.

4 The next page displays the environment for the selected assets.

5 The next page lists the service objectives that are available for the selected asset type. You can expand the service objective to view details. Select the service objective that provides disaster recovery operations.

6 Select the target (recovery) data center.

7 Continue through the wizard to configure Resiliency Platform Data Mover for replication.

8 Select the target volume type for each disk. Enter the IOPS required if the volume type is Provisioned IOPS SSD.

Refer to AWS documentation for more information on IOPS permitted for specific volume type and size.

9 Complete the customization for AWS.

See "AWS Customization options panel" on page 68.

10 Complete the network customization steps for the virtualization technology.

See "Network customization options" on page 65.

11 Verify the summarized information and enter a name for the resiliency group.

12 When you finish the wizard steps, Resiliency Platform invokes a workflow which initializes the DR operation. You can view the progress or ensure that this operation is successfully completed on the **Activities** page.

See "Viewing activities" on page 105.

Verify that the new resiliency group is added to the **Resiliency Groups** tab.

See "Viewing resiliency group details" on page 57.

If the operation fails and you want to delete the resiliency groups, and while deleting the resiliency group if you face any issues, you need to perform manual cleanup steps.See "Troubleshooting delete resiliency group operation" on page 115.

AWS Customization options panel

This wizard panel is displayed when you are configuring the assets for remote recovery in Amazon Web Services (AWS).

The following table summarizes the information that you must supply or verify in the wizard to complete the customization.

Table 6-2 Options for AWS customization

Wizard steps and options	Description
Select Attributes	Review the availability zone for the selected virtual machines.
	Select a security group and the instance type. You can choose multiple security groups for a virtual machine. Ensure that the security group is in the same VPC as the subnet that is selected during network mapping.
	After selecting the security group and the instance type you can choose to apply the selection to each virtual machine or to all.

Table 6-2   Options for AWS customization *(continued)*

Wizard steps and options	Description
Network Details	Displays the source and the target network mapping details that are applicable for the selected resiliency groups.
	If network mapping is not done, then the page is empty.

Section 4

Managing disaster recovery

- Chapter 7. Rehearsing DR operations to ensure DR readiness
- Chapter 8. Performing disaster recovery operations

Rehearsing DR operations to ensure DR readiness

This chapter includes the following topics:

- About ensuring the disaster recovery readiness of your assets
- Rehearse operations for Hyper-V virtual machines
- Performing the rehearsal operation
- Performing cleanup rehearsal

About ensuring the disaster recovery readiness of your assets

Resiliency Platform provides a rehearse operation to help you ensure the disaster recovery readiness of the assets in your protected resiliency groups.

A disaster recovery rehearsal is an operation to verify the ability of your configured resiliency group to fail over on to the target (recovery) data center during disaster. A rehearsal is a zero-downtime test that mimics the configuration, the application data, the storage, and the failover behavior of your resiliency group.

When you are satisfied with the testing of the simulated failover to the target data center, you can use the cleanup rehearsal operation to clean up any temporary objects created during the rehearsal.

Rehearse operations for Hyper-V virtual machines

The rehearse operations for Hyper-V virtual machines depend on the replication type.

Rehearse operations with EMC SRDF-based replication

- The device group should be associated with the snapshot LUNs. Resiliency Platform supports TimeFinder Snap and TimeFinder Mirror (BCV).

- Rehearsal operations for the resiliency groups that are replicated using EMC SRDF technology in Asynchronous mode cannot be performed using TimeFinder Snap technology (VDEV devices). You need to configure TimeFinder Mirrors (BCV devices) to perform the rehearsal operations on such resiliency groups.

- When the rehearse operation is initiated, Resiliency Platform creates point-in-time snapshots, since rehearsal cannot work with existing snapshots. The volumes present on the snapshot device are mounted on the DR host.

- When the rehearse operation is initiated, the DR virtual machines are cloned and disconnected from the network and Resiliency Platform starts the virtual machines.

Rehearse operations with Hyper-V Replica

- When the rehearse operation is initiated, the DR virtual machines are cloned and Resiliency Platform starts the virtual machines.

Performing the rehearsal operation

Use the **Rehearsal** option on the Resiliency Platform console to perform the disaster recovery rehearsal, which verifies the ability of your configured resiliency group to fail over to the disaster recovery (DR) data center during disaster. A rehearsal is a zero-downtime test that mimics the configuration, application data, storage, replication, and the fail over behavior of your resiliency group.

Note: You can perform the Rehearsal operation only on the recovery data center.

The time taken to complete the Rehearsal operation depends on the size and the number of volumes. If the recovery data center is in AWS cloud, then to reduce the time taken to complete the snapshot creation task during Rehearsal, you may take a snapshot of the volumes manually before running the Rehearsal operation. Before taking a snapshot, ensure that the replication state is Consistent. Since, in AWS the subsequent snapshots are only incremental, the time taken to create snapshots during Rehearsal is significantly reduced. Which reduces the overall time taken to complete the operation.

To perform the rehearsal operation

1 ■ Each type of replication has prerequisites and limitations for the rehearsal operation.

   ■ It is recommended to stop or disable NetworkManager on RHEL hosts having multiple NICs.

   ■ If the recovery data center is in AWS, then configure a rehearsal subnet in the cloud. The rehearsal and production subnet should be in the same VPC.

2 Navigate

    Assets (navigation pane)

   Resiliency Groups

3 Double-click the resiliency group to view the details page. Click **Rehearsal**.

4 Select the target data center and then click **Next**.

Before you perform the rehearsal operation again, you need to ensure that the previous rehearsal is cleaned up by running the Cleanup Rehearsal operation.

Performing cleanup rehearsal

After you have performed the rehearsal operation successfully to verify the ability of your configured resiliency group to fail over on to the disaster recovery data center, you can use the cleanup rehearsal operation to clean up the rehearsal virtual machines or applications in the resiliency group. All temporary objects created during the rehearsal operation are now deleted.

A few examples of these temporary objects on Hyper-V servers are:

■ A separate copy of virtual machine when you use Hyper-V Replica for data replication.

■ A new registered virtual machine that has its virtual machine data files (VHDX) residing on snapshot LUNs when array-based replication (for example, EMC SRDF) is used for data replication.

To perform cleanup rehearsal

1   Navigate

             Assets (navigation pane)

                  Resiliency Groups

2   Double-click the resiliency group to view the details page. Click **Cleanup Rehearsal**.

3   Select the target data center, and then click **Next**.

If the replication technology used is 3PAR Remote Copy, then refresh the 3PAR enclosure after successfully completing the rehearsal cleanup operation.

See "Performing the rehearsal operation" on page 72.

8

Performing disaster recovery operations

This chapter includes the following topics:

- Migrating a resiliency group of virtual machines
- Taking over a resiliency group of virtual machines
- Performing the resync operation

Migrating a resiliency group of virtual machines

Migration refers to a planned activity involving graceful shutdown of virtual machines at the production data center and starting them at the recovery data center. In this process, replication ensures that consistent virtual machine data is made available at the recovery data center. In Veritas Resiliency Platform, the migration of virtual machines is achieved by grouping them in a resiliency group, configuring disaster recovery for the resiliency group, and thereafter performing the migrate operation on this resiliency group.

If you perform the takeover operation, then you must perform the resync operation before you migrate back to the production data center.

If the **Enable reverse replication** option is not selected while configuring for remote recovery, then you need to run the Resync operation before migrating the virtual machines back to the production data center.

See "Performing the resync operation" on page 77.

To migrate virtual machines

1 ■ Ensure that the replication is in Active state and the data is consistent.

- It is recommended to stop or disable NetworkManager on RHEL hosts having multiple NICs.

- For Hyper-V virtual machines, ensure that the network mapping of all the required virtual switches across the data centers is complete.
See "Setting up network mapping between production and recovery data centers" on page 46.

- If the recovery data center is in AWS, then ensure that the network mapping of all the required subnets across the data centers is complete.

2 Navigate

    ▥   **Assets** (navigation pane)

        Resiliency Groups

3 Double-click the resiliency group to view the details page. Click **Migrate**.

4 Select the target data center and click **Next**.

If the Migrate operation fails, check **Recent Activities** to know the reason and fix it. You can then launch the **Retry** operation. The **Retry** operation restarts the migrate workflow, it skips the steps that were successfully completed and retries those that had failed.

Do not restart the workflow service while any workflow is in running state, otherwise the **Retry** operation may not work as expected.

Taking over a resiliency group of virtual machines

Takeover is an activity initiated by a user when the production data center is down due to a natural calamity or other disaster, and the virtual machines need to be restored at the recovery data center to provide business continuity. The user starts the virtual machines at the recovery data center with the available data. Since it is an unplanned event, the data available at the recovery data center may not be up to date. You need to evaluate the tolerable limit of data loss, and accordingly take the necessary action - start the virtual machines with the available data, or first use any other available data backup mechanism to get the latest copy of data, and thereafter start the virtual machines. The takeover operation brings up the virtual machines at the recovery data center using the last available data.

Perform the resync operation after successful completion of takeover operation.

If the recovery data center is in Amazon Web Services (AWS), then takeover operation from cloud data center to production (on-premises) data center is not supported.

To perform takeover operation on virtual machines

1 Prerequisites

- It is recommended to stop or disable NetworkManager on RHEL hosts having multiple NICs.

- For Hyper-V virtual machines, ensure that the network mapping of all the required virtual switches across the data centers is complete.
  See "Setting up network mapping between production and recovery data centers" on page 46.

- If the recovery data center is in AWS, then ensure that the network mapping of all the required subnets between the production and recovery data center is complete.

2 Navigate

Assets (navigation pane)

Resiliency Groups

3 Double-click the resiliency group to view the details page. Click **Takeover**.

4 Select the target data center and click **Next**.

If the Takeover operation fails, check **Recent Activities** to know the reason and fix it. You can then launch the **Retry** operation. The **Retry** operation restarts the migrate workflow, it skips the steps that were successfully completed and retries those that had failed.

Do not restart the workflow service while any workflow is in running state, otherwise the **Retry** operation may not work as expected.

Performing the resync operation

When disaster strikes on a production data center, the takeover operation is invoked to start the resiliency groups on the recovery data center.

Since the production data center is not working, the data replication between the two sites does not happen. After the production site is back up and running, you need to prepare the production site for the next failover or for a migration operation. This preparation includes cleaning up any residue and resuming the replication from the recovery to the production site.

Use the Resync operation on the Resiliency Platform console to automate these steps for the required resiliency groups. This operation cleans up the residue which includes stopping applications and virtual machines, unregistering virtual machines,

unmounting file systems, datastores, etc. If the recovery data center is Amazon Web Services, then the virtual machines are not unregistered.

In Microsoft Failover Cluster environments, the Resync operation may fail in the first step to cleanup the virtual machine residue. You can manually cleanup the virtual machine residue and proceed.

See "Manually cleaning up virtual machines" on page 115.

If the recovery data center is Amazon Web Services (AWS), and if while configuring the resiliency group for remote recovery, the **Enable reverse replication** option was selected, then do not run the Resync operation till the replication is complete and the state is IN Sync.

Performing the resync operation

1   Prerequisites

    If the recovery data center is not in cloud, then restart the ESX servers on primary site before performing resync operation. Restarting the ESX servers ensures that all stale references to virtual machines, disks, or datastores are released so that resync can work properly.

2   Navigate

Assets (navigation pane)

Resiliency Groups

3   Double-click the resiliency group to view the details page. Click **Resync**.

4   In the **Resync** panel, select the production data center name from the drop-down list, and click **Next**.

If the Resync operation fails, check **Recent Activities** to know the reason and fix it. You can then launch the **Retry** operation. The **Retry** operation restarts the migrate workflow, it skips the steps that were successfully completed and retries those that had failed.

Do not restart the workflow service while any workflow is in running state, otherwise the **Retry** operation may not work as expected.

Managing resiliency plans

This chapter includes the following topics:

About resiliency plans

Using the Veritas Resiliency Platform console you can create customized resiliency plans. A resiliency plan is a customized set of tasks that you can run as a single operation. You add each task and the particular assets on which to run the task. If you intend to use the same sequence of tasks on different assets, you can create

a resiliency template. You can save the template and use it to create multiple resiliency plans.

For example, you can create a resiliency plan template to migrate a resiliency group. Then you can add a resiliency group to the template to create a plan. You can create multiple plans using the same template.

You can create customized resiliency plans for performing all the disaster recovery operations such as migrate, takeover, rehearsal, cleanup rehearsal, and resync. You can also create customized resiliency plans for executing a manual task or a custom script.

You do not have to create a template in order to create a resiliency plan. Resiliency plans can be created using blank templates.

Note: To create a plan for migrate, takeover, rehearsal, or cleanup rehearsal operation, configure disaster recovery task must be successful on the selected resiliency group.

You can schedule the resiliency plan to run at a particular time.

Using these predefined templates, you can create resiliency plans by adding assets to the template. You can then run these plans on a later date.

See "Creating a new resiliency plan template" on page 80.

See "Creating a new resiliency plan" on page 85.

Creating a new resiliency plan template

Using the Veritas Resiliency Platform console, you can create a customized resiliency plan template for the following operations:

- Start and stop a resiliency group.

- Rehearsal and cleanup rehearsal of a resiliency group.

- Migrate and takeover a resiliency group.

- Manual task
  See "About manual task" on page 81.

- Run a custom script
  See "About custom script" on page 82.

To create a template, you need to drag and drop the required operation from the stencil into the canvas below. The arrow lets you connect various operations in the canvas.

For example, if you want to create a template to perform the Start Resiliency Group task, drag the operation from the top bar into the canvas. Now click on the arrow on the **Start** action box and drag the mouse to the **Start Resiliency** In addition to the above listed tasks, you can also add a custom script Manual task in the resiliency plan. This task temporarily pauses the operation letting you perform a task before proceeding further.

Group action box. Similarly you can drag the arrow from the **Start Resiliency Group** action box to the **End** action.

To create a new resiliency plan template

1   Navigate

     Automation Plans (menu bar) > **Resiliency Plans** or **Quick Actions** > **Resiliency Plans**

2   In the **Templates** section, click **New**.

3   In the **Create New Template** wizard panel, enter a name and a description for the template.

4   Drag and drop the required operation into the canvas. Connect the **Start** and **Stop** actions to the operation.

5   Click **Create**.

See

About manual task

Using the Resiliency Platform console, you can add a manual task in the resiliency plan. The purpose of including this task in resiliency plan is to temporarily pause the operation of the resiliency plan to perform a task or validate a step before you proceed further.

You can specify a timeout for the manual task. After the specified timeout expires, the manual task in the resiliency plan is marked as complete and the resiliency plan proceeds further.

Alternatively, you can opt for manually resuming the process. In this case, the resiliency plan enters into a pause state. You need to go to the **Inbox** in Resiliency Platform console and click **Resume** on the corresponding entry in the **Inbox**. You can also resume the resiliency plan by right-clicking the corresponding entry in **Activities > Current Activities** and selecting **Resume**.

Using manual tasks in resiliency plans

Using the Resiliency Platform console, you can add a manual task in the resiliency plan.

To use a manual task in a resiliency plan

1   You can add a manual task to a resiliency plan template or to a resiliency plan.

See "Creating a new resiliency plan template" on page 80.

See "Creating a new resiliency plan" on page 85.

2   Drag and drop **Manual Task** into the canvas. Click the pencil icon in the action box to add the task details.

3   Provide a name for the manual task.

4   Describe the reason why you want to add this manual task to the resilient plan.

5   Select your choice for resuming the process manually or automatically. If you select the option for automatically resuming the process after a timeout, enter the duration of timeout in minutes. Click **Save**.

About custom script

Using the Resiliency Platform console, you can add a custom script execution task in the resiliency plan. You can use the custom script execution task to perform customized operations before executing the next step of the resiliency plan such as repurposing capacity on the recovery site, orchestrate network changes, or any kind of post-processing.

Custom Script execution requires Resiliency Platform deployed on the Resiliency Manager, Infrastructure Management Server (IMS) and the hosts executing custom scripts. In addition, if you are using Resiliency Platform with Veritas InfoScale, the Veritas Resiliency Platform Enablement add-on has to be manually installed on applicable hosts.

The custom script can be in any format that can be directly executed on a shell on the target host. For the Linux hosts, it may be an executable or a script that specifies the interpreter on the shebang line such as a shell or a Perl script. For Windows hosts, it may be an executable or a script with known extension such as a bat file or a PowerShell script. The Script is executed as root user on a UNIX host or as Local System on a Windows host. You may use `sudo` or `RunAs` commands to execute some other scripts from these custom scripts.

Before you can execute the script as part of the resiliency plan, you need to manually copy the script to the *VRTSsfmh InstallDir*/vrp/scripts directory on the host.

Where, *VRTSsfmh InstallDir* is /opt/VRTSsfmh on the Unix/Linux hosts and *SystemDrive*/Program Files/VERITAS/VRTSsfmh on the Windows hosts. Copying the script to these specific folders enforces the security policy for running a custom script since these folders can be accessed only by a root user or a Local System.

Exit code from script execution determines the success or failure of the task in the resiliency plan workflow. An exit code of zero means the script execution was successful while a non-zero exit code means the script execution failed. If you select the option to ignore the exit code, the script task is always marked as successful after completion of the script. You can select this option, if your script does not return any exit code. You can view the output of the script in activity details for the resiliency plan in Resiliency Platform console.

If you uninstall the host package from the host where you have copied your custom script, the custom script is removed from the host as part of the uninstallation process.

Using custom scripts in resiliency plans

Using the Resiliency Platform console, you can add a custom script execution task in the resiliency plan.

To use a custom script execution task in a resiliency plan

1 You can add a custom script execution task to a resiliency plan template or to a resiliency plan.

See "Creating a new resiliency plan template" on page 80.

See "Creating a new resiliency plan" on page 85.

2 Drag and drop **Custom Script** into the canvas. Click the pencil icon in the action box to add the task details.

3 Enter a name for the custom script.

4 Select the data center and the host where you want to execute the script. Click **Next**.

5 Enter the following details:

- The relative path of the script on the specified host. The script path that you enter is taken as relative to the *VRTSsfmh InstallDir*/vrp/scripts/ directory path.
  For example, if you enter the path of the script as myscripts/backup_scripts/*script_name*, then the complete path considered by the system will be *VRTSsfmh InstallDir*/vrp/scripts/myscripts/backup_scripts/*script_name*.

- Command-line arguments to the script. This is an optional input field.

- Timeout for the script. By default, there is no timeout for the script execution. You can specify a timeout for the script execution. After the specified timeout expires, the script execution task in the resiliency plan is marked as failure but the script execution task is not stopped. The script execution may

continue in the background. If you do not specify any timeout, the task will wait till the script is not completed.

6   Click **Save**.

Editing a resiliency plan template

Using the Veritas Resiliency Platform console, you can edit an existing resiliency plan template.

You can add assets to these templates and create a customized resiliency plan. Any changes to the template do not affect the existing resiliency plans that you created from the template.

To edit a resiliency plan template

1   Navigate

   Automation Plans (menu bar) > **Resiliency Plans** or **Quick Actions** > **Resiliency Plans**

2   In the **Templates** list, place your cursor on the row which you want to edit. Do one of the following:

   ■   Right click your mouse and click **Edit**.

   ■   Click on the vertical ellipsis and select **Edit**.

3   In the **Edit Template** wizard panel, edit the required actions and click **Save**.

   The steps for editing the plan are the same as creating it.

See

Deleting a resiliency plan template

Using the Veritas Resiliency Platform console you can delete an existing resiliency plan template.

Deleting the template does not affect the existing resiliency plans that you created from the template.

To delete a resiliency plan template

1   Navigate

   Automation Plans (menu bar) > **Resiliency Plans** or **Quick Actions** > **Resiliency Plans**

2   In the **Templates** list, place your cursor on the row which you want to delete. Do one of the following:

- Right click your mouse and click **Delete**.

- Click on the vertical ellipsis and select **Delete**.

3 In the **Delete Template** panel click **Delete**.

See "Creating a new resiliency plan template" on page 80.

Viewing a resiliency plan template

Using the Veritas Resiliency Platform console, you can view the details of a resiliency plan template. To view the details of the resiliency plan templates, you need to have at least guest persona assigned to you.

To view a resiliency plan template

1 Navigate

Automation Plans (menu bar) > **Resiliency Plans** or **Quick Actions** > **Resiliency Plans**

2 In the **Templates** list, do one of the following:

- Double click the row that you want to view.

- Select the row that you want to view, right click and select Details.

- Select the row that you want to view, click on the vertical ellipsis and select Details.

3 You can now view the details of the resiliency plan template.

Creating a new resiliency plan

Using the Veritas Resiliency Platform console, you can create a new resiliency plan for the following operations. Resiliency plans can be created using an existing template or with a blank template. When you create a plan using a blank template, you need to create the plan and add the assets at the same time.

- Start and stop a resiliency group.

- Rehearsal and cleanup rehearsal of a resiliency group.

- Migrate and takeover a resiliency group.

- Manual task
  See "About manual task" on page 81.

- Run a custom script
  See "About custom script" on page 82.

Note: To create a plan for migrate, takeover, rehearsal, or cleanup rehearsal operation, disaster recovery must be configured successfully on the selected resiliency group or the VBS.

To create a new resiliency plan using blank template

1   Navigate

Automation Plans (menu bar) > **Resiliency Plans** or **Quick Actions** > **Resiliency Plans**

2   In the **Saved Plans** section, click **New**.

3   In the **Create Saved Plan - Select Template** wizard panel, select **Blank Template**, and click **Next**.

4   In the **Add Assets** panel, enter name and description.

5   Drag and drop the required operation into the canvas. Connect the **Start** and **Stop** actions to the operation.

6   Click the pencil icon in the action box to add relevant assets. Select the data center whose assets you want to add to the template. Click **Add**.

7   Click **Submit**.

To create a new resiliency plan using predefined template

1   Navigate

Resiliency Plans (menu bar) or **Quick Actions** > **Resiliency Plans**

2   In the **Saved Plans** section, click **New**.

3   In the **Create Saved Plan - "Select Template"** wizard panel, select **Pre-defined Template**.

4   Select a template from the list and click **Next**.

5   In the **Add Assets** panel, name and description are pre-populated.

6   Click the pencil icon in the action box to add relevant assets. Select the data center whose assets you want to add to the template. Click **Add**.

7   Click **Submit**.

See "About resiliency plans" on page 79.

See "Deleting a resiliency plan" on page 87.

See "Executing a resiliency plan" on page 87.

Editing a resiliency plan

Using the Veritas Resiliency Platform console, you can edit a resiliency plan.

To edit a resiliency plan

1   Navigate

Automation Plans (menu bar) > **Resiliency Plans** or **Quick Actions** >
Resiliency Plans

2   In the **Saved Plans** list, place your cursor on the row which you want to edit.
Do one of the following:

- Right click your mouse and click **Edit**.

- Click on the vertical ellipsis and select **Edit**.

3   In the **Edit Saved Plan** wizard panel, edit the required actions and click **Submit**.

The steps for editing the plan are the same as creating it.

See "Creating a new resiliency plan" on page 85.

Deleting a resiliency plan

Using the Veritas Resiliency Platform console, you can delete a resiliency plan.

To delete a resiliency plan

1   Navigate

Automation Plans (menu bar) > **Resiliency Plans** or **Quick Actions** >
Resiliency Plans

2   In the **Saved Plans** list, place your cursor on the row which you want to delete.
Do one of the following:

- Right click your mouse and click **Delete**.

- Click on the vertical ellipsis and select **Delete**.

3   In the **Delete Saved Plan** panel click **Delete**.

See "Creating a new resiliency plan" on page 85.

Executing a resiliency plan

Using the Veritas Resiliency Platform console, you can execute a resiliency plan.
After executing the resiliency plan, you can navigate to the **Activities** page to view
the progress of the plan.

To execute a resiliency plan

1 Navigate

Automation Plans (menu bar) > **Resiliency Plans** or **Quick Actions** >
Resiliency Plans

2 In the **Saved Plans** list, place your cursor on the row which you want to execute.
Do one of the following:

- Right click your mouse and click **Execute**.

- Click on the vertical ellipsis and select **Execute**.

3 In the **Execute Saved Plan** panel click **Execute**.

See "Creating a new resiliency plan" on page 85.

Viewing a resiliency plan

Using the Veritas Resiliency Platform console, you can view the details of a resiliency
plan. To view the details of the resiliency plans, you need to have at least guest
persona assigned to you.

You can also launch operations such as edit a resiliency plan or delete a resiliency
plan from this view.

See "Editing a resiliency plan" on page 87.

See "Deleting a resiliency plan" on page 87.

To view a resiliency plan

1 Navigate

Automation Plans (menu bar) > **Resiliency Plans** or **Quick Actions** >
Resiliency Plans

2 In the **Saved Plans** list, do one of the following:

- Double click the row that you want to view.

- Select the row that you want to view, right click and select **Details**.

- Select the row that you want to view, click on the vertical ellipsis and select
Details.

3 You can now view the details of the resiliency plan. Click the watch icon to see
the details of the components of a resiliency plan such as a custom script or
a manual task.

Creating a schedule for a resiliency plan

Using the Veritas Resiliency Platform console, you can create a schedule for a resiliency plan.

To create a schedule for a resiliency plan

1   Navigate

Automation Plans (menu bar) > **Resiliency Plans** or **Quick Actions** > **Resiliency Plans**

2   In the **Saved Plans** list, do one of the following:

- Double click the row for which you want to create a schedule. In the **Schedule** section of details page, click **New**.

- Select the row for which you want to create a schedule, right click and select **Create Schedule**.

- Select the row for which you want to create a schedule, click on the vertical ellipsis and select **Create Schedule**.

Editing a schedule for a resiliency plan

Using the Veritas Resiliency Platform console, you can edit a schedule for a resiliency plan.

To edit a schedule for a resiliency plan

1   Navigate

Automation Plans (menu bar) > **Resiliency Plans** or **Quick Actions** > **Resiliency Plans**

2   In the **Saved Plans** list, do one of the following:

- Double click the row for which you want to edit a schedule. In the **Schedule** section of details page, click **Edit**.

- Select the row for which you want to create a schedule, right click and select **Edit Schedule**.

- Select the row for which you want to create a schedule, click on the vertical ellipsis and select **Edit Schedule**.

Deleting a schedule for a resiliency plan

Using the Veritas Resiliency Platform console, you can delete a schedule for a resiliency plan.

To delete a schedule for a resiliency plan

1   Navigate

    Automation Plans (menu bar) > **Resiliency Plans** or **Quick Actions** >
    Resiliency Plans

2   In the **Saved Plans** list, do one of the following:

- Double click the row for which you want to delete a schedule. In the
  Schedule section of details page, click **Delete**.

- Select the row for which you want to edit a schedule, right click and select
  Delete Schedule.

- Select the row for which you want to edit a schedule, click on the vertical
  ellipsis and select **Delete Schedule**.

Viewing a schedule for a resiliency plan

Using the Veritas Resiliency Platform console, you can view a schedule for a
resiliency plan. To view the details of the resiliency plans, you need to have at least
guest persona assigned to you.

You can also launch operations such as edit a schedule or delete a schedule from
this view.

See "Editing a schedule for a resiliency plan" on page 89.

See "Deleting a schedule for a resiliency plan" on page 89.

To view a schedule for a resiliency plan

1   Navigate

    Automation Plans (menu bar) > **Resiliency Plans** or **Quick Actions** >
    Resiliency Plans

2   In the **Saved Plans** list, do one of the following:

- Double click the row for which you want to view a schedule.

- Select the row for which you want to view a schedule, right click and select
  Details.

- Select the row for which you want to view a schedule, click on the vertical
  ellipsis and select **Details**.

3   In the **Schedule** section of details page, view the details of the schedule.

Monitoring risks, reports, and activites

This chapter includes the following topics:

About the Resiliency Platform Dashboard

The Resiliency Platform Dashboard gives you an overview of your resiliency domain. Use the Dashboard to answer questions such as the following:

- Which of my data centers have Resiliency Platform managed assets?

- What is the mix of my assets by type and platform?

- Which assets are configured for disaster recovery?

The Dashboard has the following areas:

Global View	A world map that identifies the data centers that contain Resiliency Platform managed assets.
	Lines between data centers indicate that replication takes place between the locations.
	Mouse over an icon for basic Resiliency Platform configuration and asset configuration information for that data center. Click **More** for detailed information and recent activity.
Resiliency Groups and **Virtual Business Services** summaries	The upper right section of the dashboard displays total number of resiliency groups and virtual business services in the resiliency domain, as well as those at risk and normal.
	Click a square in either the **Resiliency Groups** or **Virtual Business Services** summary to display a tab of detailed information.
	The **Activity Summary** provides details of the DR activities such as average time taken, failed and successful runs.
Virtual Machines by Platform and OS	Displays a summary of virtual machines in all data centers or information on a single data center. Use the drop-down list to filter your results. The summary lists the virtual machine types by percentage and the platform types by number.
Risks Summary	Displays a summary of errors and warning in all data centers. Click **View Details** to view additional information.
Application environment	Displays the number of applications and the application types. The chart shows the number of applications that are managed by InfoScale and those that are not managed by InfoScale.

Applications by Type	Displays a summary of application types in all data centers or in a single data center. Use the drop-down list to filter your results.
Top Resiliency Groups by Replication Lag	Ranks the resiliency groups according to how long it takes the recovery data center to be in sync with the active data center.
By Service Objective	Displays the percentage of virtual machines and applications that are unprotected or unmanaged.
	Use the drop-down list to filter your results.

See "Displaying resiliency group information and status" on page 54.

Understanding asset types

On the Resiliency Platform console Assets page, assets are classified as follows.

Asset	Description
Resiliency Group	A group of applications or virtual machines under Resiliency Platform control. You can use Resiliency Platform to start and stop the resiliency group, as well as protect and manage it.
	The Overview tab identifies whether or not resiliency groups are protected. An unprotected resiliency group is one that is configured to support monitoring and start and stop operations only. A protected resiliency group supports data recovery operations as well.
Virtual Business Service	A collection of resiliency groups logically grouped for a specific business purpose.
Unmanaged	An application or virtual machine that Resiliency Platform discovers in your environment, but that is not under Resiliency Platform management. You cannot use any Resiliency Platform features with these assets until they become a part of a resiliency group.

Displaying an overview of your assets

The **Assets** page gives you an overview of all your resiliency groups and virtual business services (VBSs). You can also click links on the page to create resiliency groups and VBSs.

To access the **Assets** page, go to the navigation pane on the left side of the screen, and click:

The **Assets** page is organized into the following categories:

■ Unprotected resiliency groups, are groups under Resiliency Platform control, but that do not have disaster recovery configured.
See "Managing virtual machines for basic monitoring" on page 52.

For unprotected and protected resiliency groups, the screen also displays the following:

■ The number of resiliency groups that are based on virtual machines and the number that are based on applications

■ The number of unmanaged virtual machines or applications; that is, the assets that Resiliency Platform is aware of but that are not managed or protected in resiliency groups.

For VBSs, the screen displays the following:

■ The number of VBSs that are created from virtual machines and the number that are created from physical assets.

■ The number of resiliency groups within the VBSs that are protected and the number that are only managed (not protected).

About risk insight

The objective of the Risk Insight feature is to notify you about the vulnerabilities that might impact the recoverability or continuity of your protected assets.

Risk Insight detects the changes to the state and configuration of your protected assets. It identifies if there is a risk to the recoverability or continuity of your protected assets.

Veritas Resiliency Platform also enables you to set up the replication lag threshold or service level threshold. Risk insight alerts you when the replication lags beyond the threshold that you specified.

Risk insight generates two types of reports:

■ **Current risk reports**: Provides the summary and detail information about all the current risks in your data center.

■ **Historical risk reports**: Provides a summary and a detailed analysis of information about the risks in your environment during the specified period.

These reports help you take actions to prevent such risks. The historical risk data is purged after a period of two years.

The risks covered by risk insight can be classified into three main categories:

Table 10-1

Risk Type	Description
Recoverability	Risks that may impact the ability to recover and run the application on the recovery site.
Continuity	Risks that may impact the ability to run your applications without disruption either on your production site or on your recovery site.
SLA	Risks that may impact the ability to fulfill the service level agreements (SLA) for your applications.

On the basis of criticality, the risks can be classified into two types:

Table 10-2

Risk type	Description
Error	A risk that disrupts any stated goals of the product. An error must be fixed to make the product work as expected.
Warning	A risk that jeopardizes any stated goals of the product. A warning alerts you about a potential problem in your environment.

Displaying risk information

Resiliency Platform identifies and flags several risks that may occur during data center operations. Some of these risks are transient. They are temporary and resolve themselves without your intervention. Other risks require intervention and troubleshooting to resolve.

You can display risks in the following ways:

Table 10-3   Ways to display risks

To display ...	Do the following:	
A complete list of risks across the resiliency domain	**1**	On the menu bar, select ⊞ **More Views** > **Risks**
	2	On the **Risk** page, double-click a risk in the table to display detailed information.
Risks that are associated with a specific resiliency group or virtual business service	**1**	On the navigation pane, select 🖥 (Assets) and the tab for either **Resiliency Groups** or **Virtual Business Services**.
	2	On the tab, double-click a resiliency group or virtual business service to display detailed information.
	3	On the details page, note any risks that are listed in the **At Risk** area, and double-click the risk for details.

In addition to the above mentioned views, the **More views** > **Logs** > **All** view and the **More views** > **Logs** > **Notification** view also includes the notification about the risks in your environment. You can double-click any row to view the detailed description of the error and suggested resolution for the error.

Predefined risks in Resiliency Platform

Table 10-4 lists the predefined risks available in Resiliency Platform. These risks are reflected in the current risk report and the historical risk report.

Table 10-4   Predefined risks

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Veritas Infoscale Operations Manager disconnected	Checks for Veritas Infoscale Operations Manager to Resiliency Manager connection state	1 minute	Error	All operations	Check Veritas Infoscale Operations Manager reachability Try to reconnect Veritas Infoscale Operations Manager
vCenter Password Incorrect	Checks if vCenter password is incorrect	5 minutes	Error	■ On primary site: start or stop operations ■ On secondary site: migrate or takeover operations	In case of a password change, resolve the password issue and refresh the vCenter configuration
VM tools not installed	Checks if VM Tools are not Installed. It may affect IP Customization and VM Shutdown.	Real time, when resiliency group is created	Error	■ Migrate ■ Stop	■ In case of VMWare, install VMWare Tools ■ In case of Hyper-V, install Hyper-V Integration Tools
Snapshot removed from Virtual Machine	Checks if snapshot has been removed from virtual machine.	5 minutes	Error	Resiliency Platform Data Mover replication	Edit the resiliency group to refresh configuration
Snapshot reverted on Virtual Machine	Checks if snapshot has been reverted on virtual machine.	5 minutes	Error	Resiliency Platform Data Mover replication	Remove and re-add the virtual machine to the Resiliency group by editing Resiliency group

Table 10-4   Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Data Mover Daemon Crash	Checks if VM Data Mover filter is not able to connect to its counterpart in ESX.	5 minutes	Error	Resiliency Platform Data Mover replication	In order to continue the replication, you can move (VMotion) the VM to a different ESX node in the cluster and either troubleshoot the issue with this ESX node or raise a support case with Veritas
Snapshot created on Virtual Machine	Checks if a snapshot has been created on Virtual machine.	5 minutes	Error	Resiliency Platform Data Mover replication	Edit the resiliency group to refresh configuration
DataMover virtual machine in noop mode	Checks if VM Data Mover filter is not able to connect to its counterpart in ESX.	5 minutes	Error	Resiliency Platform Data Mover replication	In order to continue the replication, you can move (VMotion) the VM to a different ESX node in the cluster and either troubleshoot the issue with this ESX node or raise a support case with Veritas
Resiliency group configuration drift	Checks if disk configuration of any of the assets in the resiliency group has changed.	30 minutes	Error	▪ Migrate ▪ Resync	Edit the resiliency group to first remove the impacted virtual machine from the resiliency group and then add it back to the resiliency group.
Global user deleted	Checks if there are no global users. In this case, the user will not be able to customize the IP for Windows machines in VMware environment.	Real time	Warning	▪ Migrate ▪ Takeover	Edit the resiliency group or add a Global user

Table 10-4 Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Missing heartbeat from Resiliency Manager	Checks for heartbeat failure from a Resiliency Manager.	5 minutes	Error	All	Fix the Resiliency Manager connectivity issue
Infrastructure Management Server disconnected	Check for Infrastructure Management Server(IMS) to Resiliency Manager(RM) connection state.	1 minute	Error	All	Check IMS reachability Try to reconnect IMS
Storage Discovery Host down	Checks if the discovery daemon is down on the storage discovery host	15 minutes	Error	Migrate	Resolve the discovery daemon issue
DNS removed	Checks if DNS is removed from the resiliency group where DNS customization is enabled	real time	Warning	▪ Migrate ▪ Takeover	Edit the Resiliency Group and disable DNS customization
IOTap driver not configured	Checks if the IOTap driver is not configured	2 hours	Error	None	Configure the IOTap driver This risk is removed when the workload is configured for disaster recovery
VMware Discovery Host Down	Checks if the discovery daemon is down on the VMware Discovery Host	15 minutes	Error	Migrate	Resolve the discovery daemon issue
VM restart is pending	Checks if the VM has not been restarted after add host operation	2 hours	Error	Configure DR	Restart the VM after add host operation
New VM added to replication storage	Checks if a virtual machine that is added to a Veritas Replication Set on a primary site, is not a part of the resiliency group.	5 minutes	Error	▪ Migrate ▪ Takeover ▪ Rehearsal	Add the virtual machine to the resiliency group.

Table 10-4   Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Replication lag exceeding RPO	Checks if the replication lag exceeds the thresholds defined for the resiliency group. This risk affects the SLA for the services running on your production data center.	5 minutes	Warning	■ Migrate ■ Takeover	Check if the replication lag exceeds the RPO that is defined in the Service Objective
Replication state broken/critical	Checks if the replication is not working or is in a critical condition for each resiliency group.	5 minutes	Error	■ Migrate ■ Takeover	Contact the enclosure vendor.
Remote mount point already mounted	Checks if the mount point is not available for mounting on target site for any of the following reasons: ■ Mount point is already mounted. ■ Mount point is being used by other assets.	■ Native (ext3, ext4,NTFS): 30 minutes ■ Virtualization (VMFS, NFS): 6 hours	Warning	■ Migrate ■ Takeover	Unmount the mount point that is already mounted or is being used by other assets.
Disk utilization critical	Checks if at least 80% of the disk capacity is being utilized. The risk is generated for all the resiliency groups associated with that particular file system.	■ Native (ext3, ext4,NTFS): 30 minutes ■ Virtualization (VMFS, NFS): 6 hours	Warning	■ Migrate ■ Takeover ■ Rehearsal	Delete or move some files or uninstall some non-critical applications to free up some disk space.
ESX not reachable	Checks if the ESX server is in a disconnected state.	5 minutes	Error	■ On primary site: start or stop operations ■ On secondary site: migrate or takeover operations	Resolve the ESX server connection issue.

Table 10-4   Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
vCenter Server not reachable	Checks if the virtualization server is unreachable or if the password for the virtualization server has changed.	5 minutes	Error	■ On primary site: start or stop operations ■ On secondary site: migrate or takeover operations	Resolve the virtualization server connection issue. In case of a password change, resolve the password issue.
Insufficient compute resources on failover target	Checks if there are insufficient CPU resources on failover target in a virtual environment.	6 hours	Warning	■ Migrate ■ Takeover	Reduce the number of CPUs assigned to the virtual machines on the primary site to match the available CPU resources on failover target.
Host not added on recovery data center	Checks if the host is not added to the IMS on the recovery data center.	30 minutes	Error	Migrate	Check the following and fix: ■ Host is up on recovery data center. ■ Host is accessible from recovery datacenter IMS. ■ Time is synchronized between host and recovery datacenter IMS.
NetBackup Notification channel disconnected	Checks for NetBackup Notification channel connection state	5 minutes	Error	Restore	Check if the NetBackup Notification channel is added to the NetBackup master server.

Table 10-4   Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Backup image violates the defined RPO	Checks if the backup image violates the defined RPO	30 minutes	Warning	No operation	- Check the connection state of NetBackup Notification channel - Check for issues due to which backup images are not available
NetBackup master server disconnected	Checks if NetBackup master server is disconnected or not reachable	5 minutes	Error	Restore	Check if IMS is added as an additional server to the NetBackup master server
Assets do not have copy policy	Checks if the assets do not have a copy policy	3 hours	Warning	No operation	Set up copy policy and then refresh the NetBackup master server
Target replication is not configured	Checks if the target replication is not configured	3 hours	Warning	No operation	Configure target replication and then refresh the NetBackup master server

Viewing the current risk report

This report provides the summary and detail information about all the current risks in your data center. The high-level summary shows the total number of risks and its distribution by severity.

The **Distribution by type** chart displays the severity-wise distribution for recoverability, continuity, and service level agreement (SLA).

The **Unresolved risks** chart shows the risks that are unresolved for more than one month, for last one month, and for last one week. The **Recent Risks** chart shows the recent risks that are generated in the last 24 hours.

The **Current risks details by type** table provides detailed information such as the name of the resiliency group which is at risk, the name of the risk, its description, object at which the risk is generated, severity, and date and time on which the risk was generated.

To view the current risk report

1 Navigation:

Click **Reports** (menu bar).

2 In the **Risk** > **Current Risk Report** section, click **Run** or **Schedule** to receive the report on the specified email address.

Viewing the historical risk report

This report provides a summary and a detailed analysis of information about the risks in your environment during the specified period.

The high-level summary shows the total number of risks and its distribution by the time the risks have been open. The information is categorized under various headings such as **Carried forward**, **New**, **Closed**, and **Still open**.

Within these categories, you can see severity wise distribution (high or low) and category wise distribution (recoverability, continuity, and service level agreement) of the risks.

The detailed analysis is displayed in the form of various charts:

■ The various charts under **Risk by Category** display the open risks and new risks in the recoverability, continuity, and SLA categories at specific points of time within the duration specified by you.

■ The **Resolving time chart** shows the average time to resolve the risk within the recoverability, continuity, and SLA categories.

■ The **5 risks that took the longest time to resolve** chart shows the top 5 risks that took the longest time to be resolved, within the recoverability, continuity, and SLA categories. This information is displayed per resiliency group or per Virtual Business Service (VBS).

To view the historical risk report

1 Navigation:

Click **Reports** (menu bar).

2 In the **Risk** > **Risk History Report** section, click **Run** or **Schedule** to receive the report on the specified email address.

Viewing reports

Veritas Resiliency Platform provides a console for viewing the following reports:

Resiliency Groups and VBS Summary	Provides details about the resiliency groups and VBSs in the data centers across all sites.
VM Inventory	Provides the platform distribution and the OS distribution details of the virtual machines that are deployed in the data centers in the form of a pie chart.
	The **Details** table provides additional information for each virtual machine.
Virtual Infrastructure Inventory	Provides information about the virtual infrastructure inventory across data centers. A pie chart shows the platform and virtualization technology distribution of the virtual servers across all data centers.
	The **Details** table provides additional information for each virtual server.
Activity Distribution History	Provides information about tasks, such as migrate, takeover, rehearse, start, and stop, performed for a specified duration.
Recovery Activity History by RG	Provides historical information about recovery tasks, such as migrate, takeover, rehearse, and restore for each resiliency group.
Recovery Activity History by VBS	Provides historical information about recovery tasks, such as migrate, takeover, rehearse, and restore for each VBS.
Metering	Provides details of the virtualization servers that are protected for disaster recovery.
	You can view the total number of servers that are protected for disaster recovery. For these servers you can view the total memory, processor cores, and the total storage.
VBS RPO	Provides Recovery Point Objective (RPO) details for all the virtual business services (VBS) in the resiliency domain.
	The bar chart provides information on the top VBS with maximum RPO lag.
	You can view the lag in the last replication and the replication date for all the VBS in the table.

To view a report

1   Navigation

Click **Reports** (menu bar).

2   Do one of the following:

- Click **Run** to receive the report on the specified email address in HTML or PDF format, or as a comma separated (.CSV) file. You can also view the saved report on the console.

- Click **Schedule** to create a report generation schedule.

For more information on configuring email settings and scheduling reports, refer to the *Deployment Guide*.

Managing activities

Using the Veritas Resiliency Platform console, you can view the sub task information for a task or an operation that is performed on the console.

Viewing activities

Using the Veritas Resiliency Platform console, you can view the sub task information for a task or an operation that is performed on the console. You can view the details on the **Activities** page. Details such as the status of the operation (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed are displayed. You can view these details for a currently running task and for the completed tasks. On the **Current** page you can abort a running task.

Click on a currently running task, to view the details in a graphical representation. The steps that are completed are shown in green color along with the success icon. The ongoing steps are in blue color with the loader image, and the future steps are in gray. Expand **Execution Details** to view all the sub-tasks that comprise the task.

To view activities

1   Navigate

⊞   **Activities** (menu bar).

2   Choose either of the following:

- Select **Current** to view the currently running tasks.

- Select **Completed** to view the historical tasks.

To view recent activities, click **Recent Activities** on the bottom pane.

See "Aborting a running activity" on page 106.

Aborting a running activity

Using the Veritas Resiliency Platform console, you can abort a task or an operation which is currently running. You can abort an operation that is executed using a resiliency plan or from the console. When you abort an operation, the sub task which is in progress is completed and then the process is aborted. The status of the sub tasks which were already completed does not change.

For example, the migrate resiliency group operation has six sub tasks. If you abort the operation while the first sub task, Stop Virtual Machine, is in progress, then the Stop Virtual Machine sub task is completed and the remaining sub tasks are skipped. If you restart the migrate operation, it starts from the beginning.

To abort an activity

1 Navigate

Do one of the following:

Activities. Skip to 2

Recent Activities (bottom pane). Click **Abort** on the required activity.

2 In the **Current** activities page, place your cursor on the activity that you want to abort. Do one of the following:

- Right click your mouse and click **Abort**.

- Click on the vertical ellipsis and select **Abort**

See "Viewing activities" on page 105.

Managing evacuation plans

This chapter includes the following topics:

About evacuation plan

An evacuation plan lets you evacuate all the assets from the production data center to the recovery data center with a single click operation.

Using the evacuation plan template you can define the sequence in which the virtual business services (VBS) should be migrated from the production data center to the recovery data center. Resiliency groups that do not belong to any VBSs, are appended at the end of the evacuation plan workflow after the VBS.

You can create an evacuation plan using only resiliency groups also. Having a VBS is not compulsory.

An evacuation plan has Priorities. You can add the VBSs to different priority levels. Ordering of resiliency groups is done by the Resiliency Platform.

If an asset within a VBS or a resiliency group fails to recover, the evacuation plan skips the asset and continues the process for the remaining assets. To do this you need to select the **Ignore failures** check box while creating the evacuation plan.

If the check box is not selected the evacuation plan stops, enabling you to fix the problem, and proceed ahead. If you choose to restart the workflow then the already executed steps are re-executed with the same results.

Only users with **Manage Evacuation Plans** permission can create and run the evacuation plans.

Prerequisites for a VBS or a resiliency group to belong to a plan

A VBS or a resiliency group should meet the following criteria to be a part of a plan.

- The VBSs and resiliency groups must be configured for disaster recovery.
  A resiliency group that is not configured for disaster recovery, or a VBS having such a resiliency group, cannot be added to the plan.

- A resiliency group must belong to only one VBS.
  Shared resiliency groups cannot be added.

An appropriate warning to exclude these assets is shown when you generate a plan.

On completing the evacuation plan, you can perform the following operations:

- Evacuate

- Rehearse evacuation

- Cleanup evacuation rehearsal

- Regenerate a plan

An alert is raised and you need to perform the **Regenerate evacuation plan** operation in the following scenarios:

- VBSs are added, modified, or deleted.

- Resiliency groups are added and configured for disaster recovery.

- Resiliency groups which were configured for disaster recovery are deleted.

- Existing resiliency group is configured for disaster recovery.

No action is required in the following scenarios:

- Resiliency groups are modified.

- Resiliency groups which are not configured for disaster recovery are deleted.

When you run the **Evacuate**, **Rehearse evacuation**, **Cleanup evacuation rehearsal**, or the **Regenerate evacuation plan** operation, you can view the workflow details in the **Activities** view.

See "Generating an evacuation plan" on page 109.

See "Regenerating an evacuation plan" on page 110.

See "Performing evacuation" on page 111.

See "Performing rehearse evacuation" on page 111.

See "Performing cleanup evacuation rehearsal" on page 111.

Generating an evacuation plan

Using the Resiliency Platform console you can generate an evacuation plan that lets you evacuate all the assets from the production data center to the recovery data center.

Using the evacuation plan template you can define the sequence in which the virtual business services (VBS) should be migrated from the production data center to the recovery data center. Resiliency groups that do not belong to any VBSs, are appended at the end of the evacuation plan workflow after the VBS.

By default only one priority group is created. To add more priority groups, click **Change Priority** and click the **+** button. You can drag and drop the VBSs into different priority groups.

Reset to Default removes all priority groups except one. All VBSs are moved into a single priority group.

If an asset within a VBS or a resiliency group fails to recover, the evacuation plan skips the asset and continues the process for the remaining assets. To do this you need to select the **Ignore failures** check box while creating the evacuation plan.

If any VBSs and resiliency groups do not fit the evacuation plan criteria, a message is displayed. We recommend that you fix the issues before creating the plan.

Only users with **Manage Evacuation Plans** permission can create and run the evacuation plans.

To generate an evacuation plan

1   Prerequisites

    See "About evacuation plan" on page 107.

2   Navigate

    Automation Plans (menu bar) > **Evacuation Plans**

3   Select **Evacuation Plans**.

4   For the required data center click **Generate Plan**.

5   Review the message if any and click **Next**.

6   Click **Change Priority** if you want to add more priority groups. Click **Submit**.

See "Performing evacuation" on page 111.

See "Performing rehearse evacuation" on page 111.

See "Performing cleanup evacuation rehearsal" on page 111.

See "Regenerating an evacuation plan" on page 110.

Regenerating an evacuation plan

After successfully creating an evacuation plan, if any of the following scenarios occur, you need to regenerate the evacuation plan.

- VBSs are added, modified, or deleted.

- Resiliency groups are added and configured for disaster recovery.

- Existing resiliency group is configured for disaster recovery.

No action is required in the following scenarios:

- Resiliency groups are added.

- Resiliency groups are modified.

- Resiliency groups which are not configured for disaster recovery are deleted.

To add more priority groups to the plan, click **Change Priority** and click the **+** button. You can drag and drop the VBSs into different priority groups. **Reset to Default** removes all priority groups except one. All VBSs are moved into a single priority group.

To regenerate an evacuation plan

1   Navigate

   Automation Plans (menu bar) > **Evacuation Plans**

2   For the required data center click **Regenerate Plan**.

3   Review the message if any and click **Next**.

4   Click **Change Priority** if you want to add more priority groups or click **Reset to Default** if you want to have only one priority group. Click **Submit**.

See "Generating an evacuation plan" on page 109.

See "Performing evacuation" on page 111.

Performing evacuation

Using the Resiliency Platform console, you can run an evacuation plan for a data center which lets you evacuate all the assets from the production data center to the recovery data center.

To run an evacuation plan

1   Navigate

    Automation Plans (menu bar) > **Evacuation Plans**

2   For the required data center, click on the vertical ellipses and select **Evacuate** to run the evacuation plan.

Performing rehearse evacuation

Using the Resiliency Platform console, you can perform a rehearsal of an evacuation plan for a data center. This verifies whether all your assets from the production data center can evacuate to the recovery data center.

To perform a rehearsal of an evacuation plan

1   Navigate

    Automation Plans (menu bar) > **Evacuation Plans**

2   For the required data center, click on the vertical ellipses and select **Rehearse Evacuation**.

Performing cleanup evacuation rehearsal

After you have performed the rehearse evacuation operation successfully to verify if all your assets from the production data center can evacuate to the recovery data

center, you can use the cleanup evacuation rehearsal operation to clean up the rehearsal virtual machines and its volumes in the VBS or resiliency groups.

All temporary objects that are created during the rehearse evacuation operation are now deleted.

During the rehearse evacuation operation, if any virtual machines are in ERROR state, then during the cleanup evacuation rehearsal operation, these virtual machines and their volumes are not deleted. You need to manually delete them. Similarly if the recovery data center is Cloud, then manually delete the instances which are in ERROR state.

To perform the cleanup rehearsal of an evacuation plan

1 Navigate

Automation Plans (menu bar) > **Evacuation Plans**

2 For the required data center, click on the vertical ellipses and select **Cleanup Evacuation Rehearsal**.

See "Performing evacuation" on page 111.

See "Performing rehearse evacuation" on page 111.

Troubleshooting

This appendix includes the following topics:

- Viewing events and logs in the console

- Events in Hyper-V virtual machines disaster discovery

- Configure DR operation fails with an integration services error

- Manually cleaning up virtual machines

- Troubleshooting delete resiliency group operation

Viewing events and logs in the console

Veritas Resiliency Platform maintains the following types of logs that can be viewed in the web console:

System logs: System logs are typically the result of a user performing an operation in the console.

Audit logs: Audit logs are primarily used for security audits. They leave a chronological trail of activities performed on the system. They identify user, activity, affected objects, etc. They help track the individuals responsible for activities and detect security violations.

Event and notification logs: Event and notification logs are not necessarily related to user activity; they can include information such as a server going down. Events can be public or private. Rules can be configured to notify users by email of selected public events. Private events are typically unrelated to user-initiated operations. Private events are displayed in the console for troubleshooting but are not available to include in rules for notification.

By default, logs and SNMP traps are retained for 2 years. This retention period can be modified in the product settings in the console.

To view events and logs

1 Navigate

⊞ **More Views** (menu bar) > **Logs**

🔔 You can also view new notifications from the **Notifications** icon.

2 To view logs by type (System, Audit, or Notification) select the appropriate tab. You can filter by the product service and by severity (information, warning, or errors) or type (public, private), depending on the tab.

Events in Hyper-V virtual machines disaster discovery

Different events (information, warning, errors) and logs (service logs, audit logs, event logs) are generated and maintained in Resiliency Platform to track system or user-initiated changes. The solution monitors Replication State to check the current state of your data replication.

For Hyper-V Replica, the Replication State attribute comes from the Replication End-Point object. For EMC SRDF, the Replication State attribute comes from EMC Symmetrix consistency group. The replication state of a consistency group is monitored to detect any replication failure, and notify user.

Note: For EMC SRDF, the replication is supported at the consistency group-level, and all the virtual machines residing in a resiliency group must consume storage from the same consistency group.

The state of the replication is monitored and a corresponding event is generated when the replication fails. The event notification can be seen on the Resiliency Platform web console. In addition, the notification is sent by email to the recipients who are configured for SMTP. An SNMP trap is also generated, which can be used by the listener, for example, any application using the generated SNMP trap.

Configure DR operation fails with an integration services error

When the Configure DR operation fails with an error "*Hyper-V integration services either not installed, or not running on one or more virtual machines*".

Resolution: Check if the appropriate integration services are installed and running. If the services are already installed and running, and the virtual machine has been restarted recently, you need to refresh the appropriate Hyper-V host. The refresh operation detects the current state of the services inside the guest. You can then proceed with the create resiliency group for remote recovery operation.

Manually cleaning up virtual machines

In Microsoft Failover Cluster environments, the Resync operation may fail in the first step to cleanup the virtual machine residue. You can manually cleanup the virtual machine residue and re-initiate the Resync operation.

To manually cleanup virtual machines

1 Logon into the Hypervisor console.

2 Open the Failover Cluster Manager snap-in.

3 Select **Roles** in the tree view.

4 Select the appropriate virtual machine resources that were configured to the resiliency group on which you want to perform the Resync operation.

   Remove the Resources.

5 Open Hyper-V Manager and delete the virtual machines that were added to the Resiliency Group on which you want to perform the Resync operation.

Troubleshooting delete resiliency group operation

While performing the delete resiliency group operation, if any of the sub-tasks fail, you can perform the following steps to reclaim the resources.

Perform the following tasks on Replication Gateways on the production and Cloud data center

1 Check for participating consistency groups.

   ```
   /opt/VRTSsfmh/bin/xprtlc -l http://localhost:8080/ConsistencyGroup
   ```

2 Verify that replication is stopped on the gateway.

   ```
   /opt/VRTSsfmh/bin/xprtlc -l
   http://localhost:8080/ConsistencyGroup/<CG_ID>/state
   ```

3 Abort replication on gateway if replication is not stopped.

```
/opt/VRTSsfmh/bin/xprtlc -m POST -l
http://localhost:8080/ConsistencyGroup/<CG_ID>/abort
```

4 Delete the consistency groups.

```
curl -X DELETE
http://localhost:8080/ConsistencyGroup/<CG_ID>/delete
```

Perform the following tasks on the hosts on the production data center

1 Verify that the consistency groups are deleted.

Linux host:

```
/opt/VRTSitrptap/bin/vxtapinfo status
```

Windows host:

```
C:\Program Files\Veritas\VRTSitrptap\cli\vxtapinfo status
```

2 Delete the consistency groups if the state is active.

- Linux host:

  ```
  /opt/VRTSitrptap/bin/vxtapaction stop -cg <CG_ID>
  /opt/VRTSitrptap/bin/vxtapconfigure delcg -cg <CG_ID> -force
  ```

- Windows host:

  ```
  C:\Program Files\Veritas\VRTSitrptap
  \cli\vxtapaction stop -cg <CG_ID>
  C:\Program Files\Veritas\VRTSitrptap
  \cli\vxtapconfigure delcg -cg <CG_ID> -force
  ```

3 Verify that the journal disk is removed from each of the virtual machines on the production data center.

Size of the journal disk is usually 1 GB and naming format is:

[<datastore-name>] <vm-hostname>/drl-<uuid>/ITRPSRDRLDisk.vmdk

Remove the journal disk using vSphere client or Hyper-V UI.

Ensure that the file is deleted.

4 Unconfigure network setting on the virtual machines on the production data center.

- Linux host:

```
/opt/VRTSsfmh/bin/perl
 /opt/VRTSsfmh/util/reconfig_vm_settings_on_prim  -unconfigure
```

- Windows host:

```
C:\Program Files\Veritas\VRTSsfmh\bin\perl.exe
 C:\Program Files\Veritas\VRTSsfmh\util
 \reconfig_vm_settings_on_prim -unconfigure
```

5   For physical machines, you need to unsign the DRL disk:

- Linux host:
  Get the device name of DRL disk by executing the following command:

  ```
  /opt/VRTSitrptap/bin/vxtapdrlfind
  ```

  Execute the following command to unsign DRL disk:

  ```
  /opt/VRTSitrptap/bin/vxtapdrlsign clear -drl_dev device_name
  ```

  Where, *device_name* is the device name of the DRL disk.

- Windows host:
  Get the device name of DRL disk by executing the following command:

  ```
  C:\Program Files\Veritas\VRTSitrptap\cli\vxtapdrlfind.exe
  ```

  Execute the following command to unsign DRL disk:

  ```
  C:\Program Files\Veritas\VRTSitrptap\cli\vxtapdrlsign clear
   -drl_dev device_name
  ```

  Where, *device_name* is the device name of the DRL disk.

6   Refresh the virtualization server if the resiliency group contains virtual machines.

Perform the following on cloud virtual machines:

- Verify that the migrated cloud virtual machines are terminated and their volumes are deleted, including the Replication block tracking disk.
  Note that the naming convention for Replication block tracking disk is `DRLVolume_RaaS_<GUID>` whereas for other disks the name starts with `resiliency-group-name_virtual-machine-name`.

- Delete any snapshots for cloud volumes.

- Detach the cloud volumes from the cloud gateway.

■ Remove the cloud virtual machines from the cloud IMS.

See "Deleting a resiliency group" on page 58.

Glossary

Index