Veritas™ Resiliency Platform 2.1 Deployment Guide

Veritas Resiliency Platform: Deployment Guide

Last updated: 2017-03-15

Document version: Document version: 2.1 Rev 1

Legal Notice

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

Veritas Technologies LLC
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 1

Overview and planning

Chapter 1

Overview of Resiliency Platform deployment

This chapter includes the following topics:

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.

See “About Resiliency Platform features and components” on page 15.

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.

See “Resiliency domain” on page 20.

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.

See “Resiliency Manager” on page 20.

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.

See "Infrastructure Management Server (IMS)" on page 21.

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.

See "Managing InfoScale applications using Resiliency Platform" on page 93.

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.

See "Replication Gateways" on page 29.

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.

See “Resiliency Platform component diagram” on page 18.

Resiliency Platform deployment infrastructure

A typical deployment consists of an Infrastructure Management Server (IMS) reporting to a Resiliency Manager. Various physical and virtual assets are associated with the IMS. For a disaster recovery deployment, this arrangement of various components and assets exists in the recovery data center as well as in the production data center.

The following diagram depicts the deployment infrastructure and how Resiliency Manager, IMS, and various assets associated with the IMS interact with each other.

Figure 1-1 Deployment infrastructure

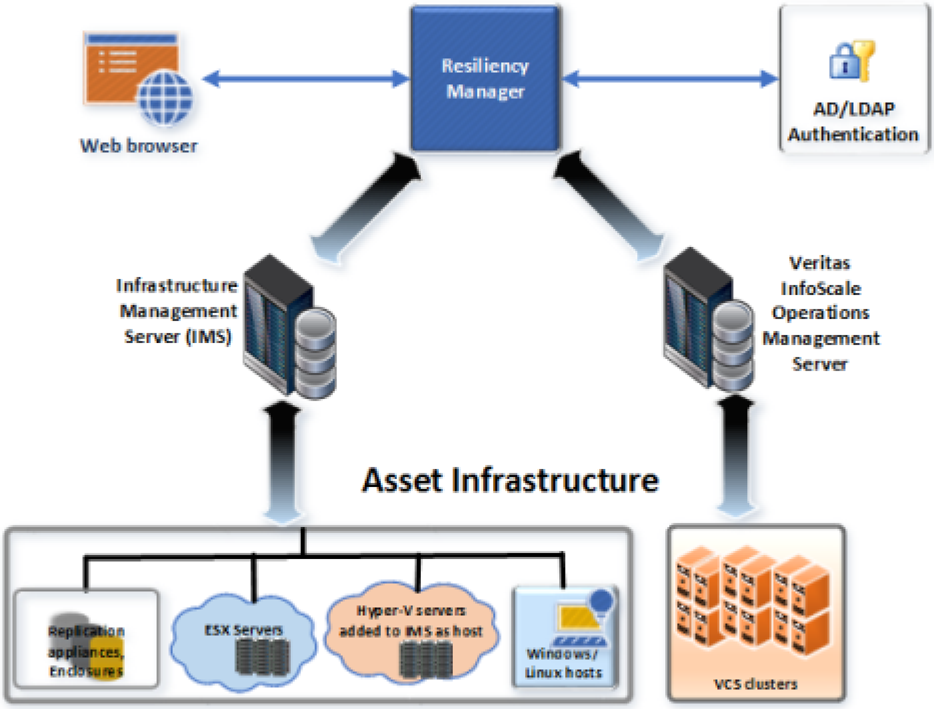
Web browser
Resiliency Manager
AD/LDAP Authentication
Infrastructure Management Server (IMS)
Veritas InfoScale Operations Management Server
Asset Infrastructure
Replication appliances, Enclosures
ESX Servers
Hyper-V servers added to IMS as host
Windows/ Linux hosts
VCS clusters

Resiliency Platform component diagram

The diagram shows a simple overview of the main components of Resiliency Platform - the Resiliency Manager, Infrastructure Management Server (IMS), and resiliency domain - and their relationships to data centers and the data center asset infrastructure.

For disaster recovery, the resiliency domain must contain at least two data centers, a production data center and recovery data center. Each data center has a Resiliency Manager and one or more Infrastructure Management Servers.

Resiliency Platform can also be implemented at a single data center for automation of workload tasks.

The asset infrastructure includes the data center assets that you add to Resiliency Platform for IMS discovery and monitoring. The asset infrastructure can include hosts (Windows or Linux servers) and virtualization servers for Hyper-V and VMware.

See “Resiliency Platform deployment infrastructure” on page 17.

Depending on the technology used for replication, the asset infrastructure can also include enclosures (storage arrays). This diagram does not show the replication details for the asset infrastructure.

See “Replication in a Resiliency Platform deployment” on page 22.

Replication, disaster recovery configuration, and disaster recovery operations are described in the solutions guides.

Figure 1-2 Resiliency Platform components

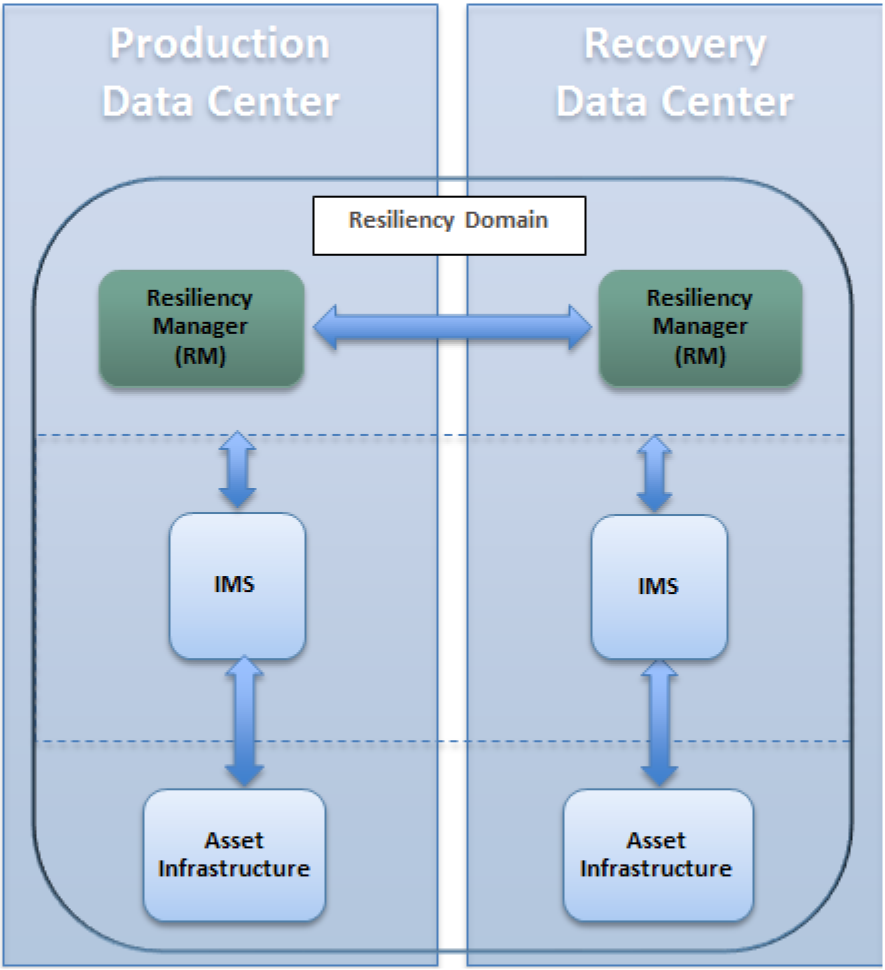

Resiliency domain

A resiliency domain is the management domain of a Veritas Resiliency Platform deployment. It represents the scope of the deployment, which can spread across multiple data centers and can include multiple Resiliency Managers and other components, along with the infrastructure that is being managed and protected. Within the resiliency domain, Resiliency Platform can protect assets, for example, virtual machines and applications, and orchestrate automation of workload tasks for the assets.

The resiliency domain is a logical object that you create from the web console after you deploy the Resiliency Manager.

Note: For disaster recovery, the resiliency domain must contain at least two data centers, a production data center and recovery data center. Each data center has a Resiliency Manager and one or more IMSs. A resiliency domain can optionally be implemented at a single data center for automation of workload tasks.

Note: For disaster recovery, the resiliency domain must contain at least two data centers, a production data center and a recovery data center. A Resiliency Domain must have at least one Resiliency Manager, preferably in the recovery datacenter. However, for efficiency and redundancy, Resiliency Managers may be deployed in more than one datacenters. A resiliency domain can optionally be implemented at a single data center for automation of workload tasks.

See “Resiliency Platform component diagram” on page 18.

See “Resiliency Manager” on page 20.

See “Infrastructure Management Server (IMS)” on page 21.

Resiliency Manager

The Resiliency Manager includes a set of loosely coupled services, a distributed data repository, and a management web console. The Resiliency Manager provides the services required for protecting assets, such as virtual machines, within a resiliency domain. It also provides workload automation services.

In a typical deployment, one Resiliency Manager is deployed in the production data center. You deploy another Resiliency Manager in a recovery data center in another geographical location.

When you deploy the first Resiliency Manager, you create the resiliency domain. When you deploy the second Resiliency Manager, you add it to the same resiliency domain (also referred to as joining the existing resiliency domain).

The Resiliency Manager discovers and manages information about data center assets from an Infrastructure Management Server (IMS), which is another required Resiliency Platform component. The Resiliency Manager stores the asset information in its data repository and displays the information in its management console.

Multiple Resiliency Managers that are part of the same domain synchronize their databases using built-in replication. Each Resiliency Manager has its own web console but because the data is synchronized, all consoles show the same data. Operations can be performed from any console and the results show in all the consoles in the resiliency domain.

See “Resiliency Platform component diagram” on page 18.

See “Resiliency domain” on page 20.

See “Infrastructure Management Server (IMS)” on page 21.

Infrastructure Management Server (IMS)

Each Resiliency Manager requires one or more Infrastructure Management Servers (IMSs). An IMS discovers and monitors assets within a data center. You use the web console to add the asset infrastructure to Resiliency Platform so that assets can be discovered and monitored by an IMS.

The asset infrastructure can include objects such as hosts, virtualization servers, and enclosures (storage arrays).

The IMS sends information about the assets to the Resiliency Manager so that the Resiliency Manager can manage the assets. Management operations on assets (for example, starting or stopping virtual machines) that you initiate from the web console are carried out by the IMS.

If there are multiple data centers in different geographical locations, a separate IMS is deployed and configured for each geographical data center location.

Each IMS connects to only one Resiliency Manager at a time. If a Resiliency Manager failure occurs, an IMS can automatically connect to another Resiliency Manager within the same domain.

You can also configure multiple Infrastructure Management Servers in the same data center. For example, to achieve scale, you can add a separate IMS for a separate business unit such as Human Resources or Finance. More than one IMS can be managed by the same Resiliency Manager.

Figure 1-3 Multiple Infrastructure Management Servers in a data center

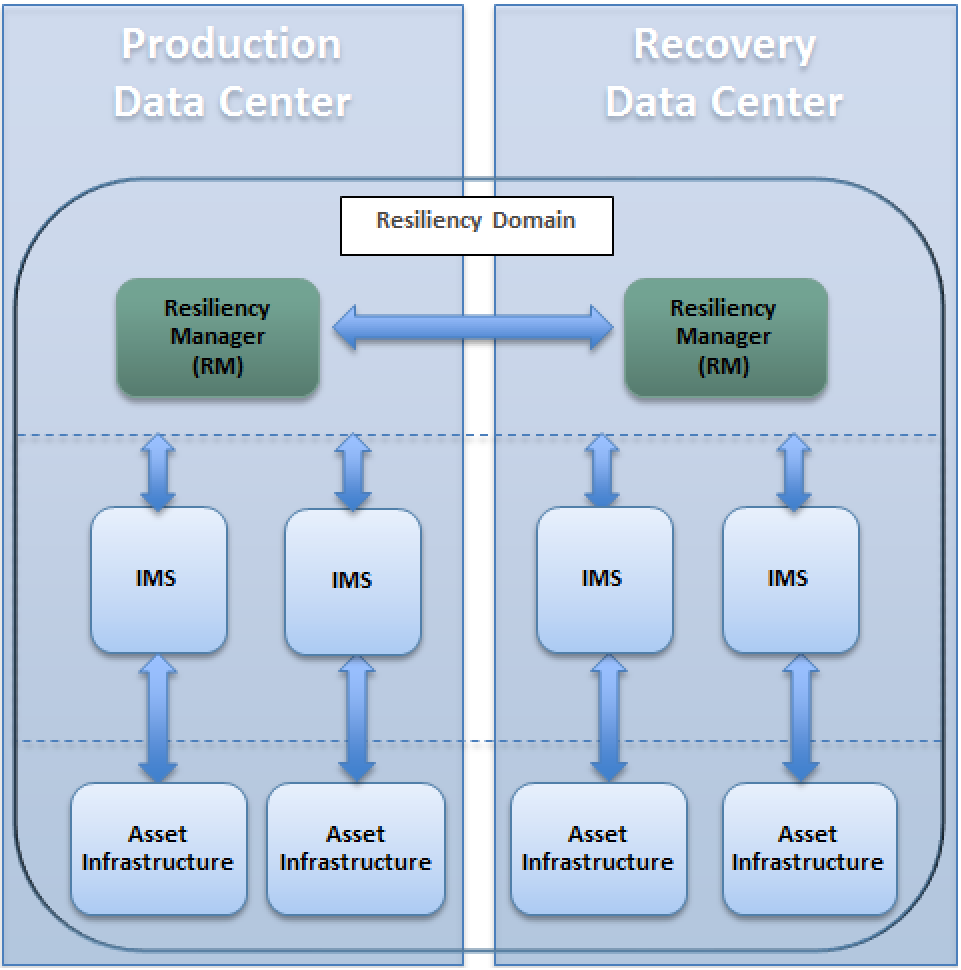

See “Resiliency domain” on page 20.

See “Resiliency Manager” on page 20.

Replication in a Resiliency Platform deployment

Veritas Resiliency Platform supports several forms of replication for data recovery from your production data center to your recovery data center.

- Array-based replication (block-based replication) using supported arrays

- Hypervisor-based replication using Hyper-V Replica
- VRP Data Mover (separately licensable feature of Resiliency Platform)
  See "About Veritas Resiliency Platform Data Mover" on page 28.

For details on supported replication hardware and software, refer to the *Hardware and Software Compatibility List*.

Figure 1-4 Array-based or hypervisor-based replication in a Resiliency Platform deployment

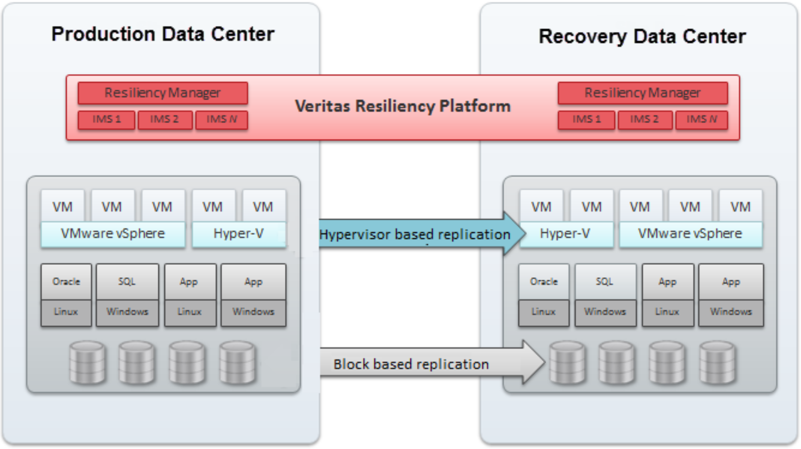

More information on how to configure replication for use with Resiliency Platform is available in the solutions guides:

Solutions for Microsoft Hyper-V.

Solutions for VMware.

Solutions for Applications

Planning a resiliency domain for efficiency and fault tolerance

Before you deploy Veritas Resiliency Platform, you should plan how to scale the deployment for efficiency and fault tolerance.

You can deploy a Resiliency Manager and Infrastructure Management Server (IMS) on the same virtual appliance. However, to meet performance requirements, production environments typically require using separate virtual appliances for the Resiliency Manager and IMS.

Therefore, the recommended minimum deployment for disaster recovery would be four virtual appliances: a Resiliency Manager and IMS in the production data center and a Resiliency Manager and IMS in the recovery data center.

The production and recovery data centers do not require a one-on-one mapping of IMSs. For example, you can have two IMSs in the production data center and one IMS in the recovery data center.

See “Resiliency domain” on page 20.

See “Resiliency Manager” on page 20.

See “Infrastructure Management Server (IMS)” on page 21.

About deploying the Resiliency Platform virtual appliance

Veritas Resiliency Platform is deployed as a virtual appliance. A virtual appliance is a virtual machine image consisting of a pre-configured operating system environment with a software application installed on it. This virtual machine image can be deployed on a hypervisor.

There are two virtual appliances available for Resiliency Platform: one is used to deploy the Resiliency Manager and the Infrastructure Management Server (IMS) and the other is used to deploy the Replication Gateway. You need to deploy the Replication Gateway only if you choose to use Resiliency Platform Data Mover for replicating your data. If you choose to use any third-party array-based or hypervisor-based replication, you only need to deploy the Resiliency Manager and IMS in your environment.

You typically deploy and configure at least one Resiliency Manager and one Infrastructure Management Server (IMS) in the production data center and at least one Resiliency Manager and one Infrastructure Management Server (IMS) in the recovery data center.

In case you plan to use Resiliency Platform Data Mover, you need to deploy at least two Replication Gateway virtual appliances. One Replication Gateway is deployed in the production data center and the other Replication Gateway is deployed in the recovery data center.

The Replication Gateway is typically not deployed on the same ESX server as the virtual machines that are to be replicated. However, the Replication Gateway on the production data center must have access to the ESX servers for the production virtual machines to be replicated. The Replication Gateway on the recovery data center must have access to the ESX servers where the recovery virtual machines will be provisioned.

You can deploy the Resiliency Platform virtual appliances using any one of the following:

- Hyper-V Manager
- VMware vSphere Client

Once the Resiliency Platform virtual appliances are deployed, you are required to configure the Resiliency Platform component through the product bootstrap.

See "Deploying the virtual appliance through Hyper-V Manager" on page 44.

See "Deploying the virtual appliance through VMware vSphere Client" on page 43.

Downloading the Resiliency Platform virtual appliances

You can download a licensed copy of the Veritas Resiliency Platform virtual appliances from MyVeritas portal.

You can download the files for deploying the virtual appliances through VMware or through Hype-V. To deploy Resiliency Platform virtual appliance through VMware, you need to download an Open Virtualization Archive (OVA) file. To deploy Resiliency Platform virtual appliance through Hyper-V, you need to download a .zip file. The .zip file contains the Virtual Hard disk (VHD) image file using which you can deploy the virtual appliance.

To download the Resiliency Platform virtual appliances

1. Log in to MyVeritas portal:

   https://my.veritas.com

2. Select **Licensing** tab, select the account and the entitlement ID that you want to use for downloading the Resiliency Platform virtual appliance.
3. In the list of products, click **Download** button next to Resiliency Platform.
4. Select the file or files to download.
   - To deploy Resiliency Manager or Infrastructure Management Server (IMS):
     - For VMware:

       `Veritas_Resiliency_Platform_VMware_Virtual_Appliance_2.1.0.100_IE.ova`
     - For Hyper-V:

       `Veritas_Resiliency_Platform_Hyper-V_Virtual_Appliance_2.1.0.100_IE.zip`
   - To deploy Resiliency Platform Data Mover:
     - For VMware:

    `Veritas_DataMover_VMware_Virtual_Appliance_2.1.0.0_IE.ova`

  - For Hyper-V:

    `Veritas_DataMover_Hyper-V_Virtual_Appliance_2.1.0.0_IE.zip`

- To deploy Resiliency Platform YUM server:
  - For VMware:

    `Veritas_Resiliency_Platform_Repository_Server_VMware_Virtual_Appliance.2.1.0.0_IE.ova`

  - For Hyper-V:

    `Veritas_Resiliency_Platform_Repository_Server_Hyper-V_Virtual_Appliance_2.1.0.0_IE.zip`

Note: For Veritas Resiliency Platform version 2.1.0.100, Resiliency Platform Data Mover version 2.1.0.0 is compatible.

You can also download a trial version of the product from the following URL:

go.veritas.com/try-vrp

Resiliency Platform deployment process overview

The following steps must be completed before you can start managing and monitoring your assets and performing disaster recovery operations with Veritas Resiliency Platform.

Table 1-1 Deployment process overview

Step	Description	More information
1	Deploy the Resiliency Platform virtual appliances	Deployment Guide See "About deploying the Resiliency Platform virtual appliance" on page 24.
2	Configure the virtual appliances as Resiliency Platform components	Deployment Guide See "About configuring the Resiliency Platform components" on page 51.
3	Set up the resiliency domain using the Getting Started wizard in the web console	Deployment Guide See "Getting started with a new Resiliency Platform configuration" on page 62.

Table 1-1 Deployment process overview *(continued)*

Step	Description	More information
4	Finish configuring the settings for the resiliency domain	Deployment Guide See "Adding an IMS " on page 70. See "Managing user authentication and permissions" on page 175. See "Managing settings for alerts and notifications and miscellaneous product settings" on page 208. See "Viewing or customizing service objectives" on page 198.
5	Add the asset infrastructure	Deployment Guide See "Adding the asset infrastructure" on page 106.
6	Create resiliency groups for the virtual machines or applications to be managed	Solutions Guides: VMware Hyper-V Applications
7	(Optional) Implement custom resiliency plans	Solutions Guides: VMware Hyper-V Applications
8	(Optional) Configure virtual business services	Solutions Guide for Virtual Business Services

Chapter 2

Overview of Resiliency Platform Data Mover support

This chapter includes the following topics:

About Veritas Resiliency Platform Data Mover

Veritas Resiliency Platform Data Mover is a licensable feature of Veritas Resiliency Platform.

Resiliency Platform Data Mover is a replication solution that is built using APIs provided by the VMware API I/O filtering (VAIO) framework. This framework is available for partners to create their own replication or caching data service for customers. Resiliency Platform Data Mover solution is certified by VMware.

Veritas Resiliency Platform Data Mover allows replication of only VMware virtual machines. Veritas Resiliency Platform Data Mover provides data replication between geographically separated data centers facilitating an effective disaster recovery solution.

Features of Veritas Resiliency Platform Data Mover include the following:

- Replicates virtual machines including its boot and data disks from source data center to target data center over any IP network in a LAN or a WAN environment.

- Enables easy recovery of virtual machines in the target data center.
- Ensures virtual machine data consistency.
- Recovers virtual machines protected by Data Mover at the Resiliency Group level.
- Enables non-disruptive testing of recovery at target data centers.

Replication Gateways

The Replication Gateway component of Veritas Resiliency Platform is a staging server that aggregates and batches data from multiple virtual machines during replication. The Gateway also performs data optimization like write cancellation. The Gateway on production data center is always paired with a Gateway on recovery data center. The recovery data center Gateway is a staging server that applies the data on the recovery data center storage.

Each Replication Gateway includes the following components:

- I/O receiver

  Receives the application I/Os that were tapped and sent by the application host in a continuous fashion.
- Transceiver

  Transfers and receives data over the WAN link periodically.
- Applier

  Applies the data to the storage after it is received on the cloud Gateway.
- Scheduler

  Manages the jobs and policies in the Gateway.
- Engine

  Maintains the state of replication and also coordinates with all other components.

During the deployment of the Gateway, the Gateway is registered as an asset to the respective IMS.

Protecting VMware virtual machines using Resiliency Platform Data Mover - an overview

To protect VMware virtual machines using Resiliency Platform Data Mover replication, ensure that the virtual machines are not already configured with another replication technology.

Before you can perform disaster recovery (DR) operations such as rehearse, migrate, and takeover on virtual machines using Resiliency Platform Data Mover, you must prepare the replication environment, and configure assets for disaster recovery. This section lists the key steps required to configure disaster recovery of VMware virtual machines using Veritas Resiliency Platform Data Mover.

Table 2-1 Configuring disaster recovery of VMware virtual machines using Data Mover

Action	Description	Refer to
Deploy the Replication Gateways	Deploy and configure the Veritas Replication Gateway virtual appliances on both data centers	See "About deploying the Resiliency Platform virtual appliance" on page 24.
Create the gateway pairs	Configure the Replication Gateway pairs to be used for replication	See "Creating a Veritas Replication Gateway pair" on page 102.
Add the VMware asset infrastructure	Prepare the VMware virtualization servers for the Resiliency Platform environment Add the VMware virtualization servers to Resiliency Platform at both data centers	See "Prerequisites for adding VMware virtualization servers" on page 125. See "Adding VMware virtualization servers " on page 134.

Table 2-1 Configuring disaster recovery of VMware virtual machines using Data Mover *(continued)*

Action	Description	Refer to
Prepare the virtual machines	Ensure that VMware Tools, storage requirements for Replication Gateway and other prerequisites are configured on virtual machines	Refer to the VMware Solutions Guide.
Configure your assets for disaster recovery	Group the required virtual machines in a resiliency group and choose the appropriate service objective to configure disaster recovery for the resiliency group	Refer to the VMware Solutions Guide.
Rehearse DR operations	Test your disaster recovery environment to ensure readiness	Refer to the VMware Solutions Guide.
Perform DR operations	Migrate or take over your virtual machines and resync the replicated data	Refer to the VMware Solutions Guide.

Overview of Amazon Web Services (AWS) support

This chapter includes the following topics:

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 3-1 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	See "Downloading the Resiliency Platform virtual appliances" on page 25. See "About deploying the Resiliency Platform virtual appliance" on page 24. See "Deploying the Resiliency Platform virtual appliances in AWS" on page 46. See "Deploying the virtual appliance through Hyper-V Manager" on page 44. See "Deploying the virtual appliance through VMware vSphere Client" on page 43.
Configure the virtual appliances as Resiliency Platform components	See "About configuring the Resiliency Platform components" on page 51.
Set up the resiliency domain and add cloud configuration using the Getting Started wizard in the web console	See "Getting started with a new Resiliency Platform configuration" on page 62.
Configure the settings for the resiliency domain	See "Adding an IMS " on page 70. See "Managing user authentication and permissions" on page 175. See "Managing settings for alerts and notifications and miscellaneous product settings" on page 208. See "Viewing or customizing service objectives" on page 198.
Add the asset infrastructure: ■ Add virtualization servers (vCenter server, Hyper-V server) ■ Prepare host for replication ■ Create gateway pair ■ Network customization (optional)	See "Adding the asset infrastructure" on page 106.
Create resiliency groups for the virtual machines to be managed	For more information refer to the VMware and Hyper-V solutions guides.
(Optional) Implement custom resiliency plans	For more information refer to the VMware and Hyper-V solutions guides.

Table 3-1 Process overview *(continued)*

Step	More information
(Optional) Configure virtual business services	For more information refer to the Solutions Guide for Virtual Business Services.
Perform disaster recovery operations.	For more information refer to the VMware and Hyper-V solutions guides.

Deployment process overview - AWS

The following is an overview of Veritas Resiliency Platform deployment infrastructure in case of AWS environment as recovery site:

Figure 3-1 Overview of Veritas Resiliency Platform deployment infrastructure

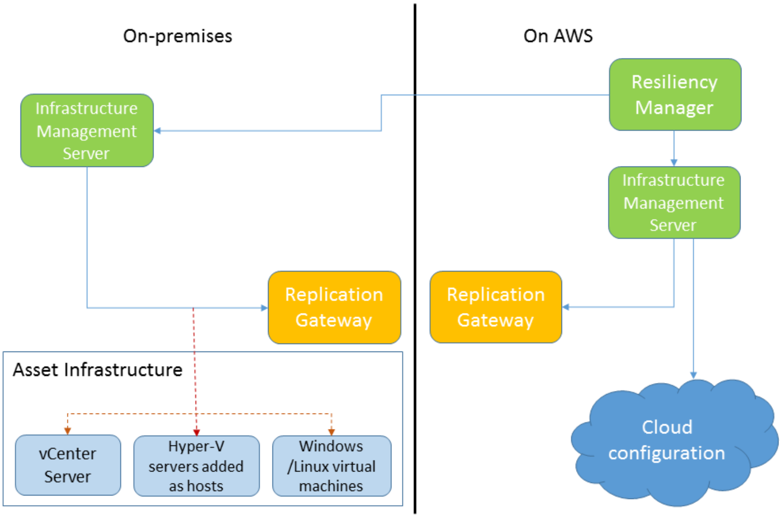

Table 3-2 describes the various steps involved in deploying and configuring Veritas Resiliency Platform virtual appliance components in case of AWS environment as recovery site:

Table 3-2 Deploying and configuring Veritas Resiliency Platform virtual appliance components

Step	Action	Description
1	Deploy and configure Resiliency Manager in cloud Deploy and configure Infrastructure Management Server (IMS) in cloud	See "Deploying the Resiliency Platform virtual appliances in AWS" on page 46.
2	Deploy and configure IMS on-premises	See "Deploying the virtual appliance through Hyper-V Manager" on page 44. See "Deploying the virtual appliance through VMware vSphere Client" on page 43.
3	Deploy and configure Replication Gateway in cloud	See "Deploying the Resiliency Platform virtual appliances in AWS" on page 46.
4	Deploy and configure Replication Gateway on-premises	See "Deploying the virtual appliance through Hyper-V Manager" on page 44. See "Deploying the virtual appliance through VMware vSphere Client" on page 43.
5	Set up the resiliency domain using the web console	See "Setting up the resiliency domain" on page 59.

Chapter 4

System requirements

This chapter includes the following topics:

- Supported hypervisors for deploying Resiliency Platform virtual appliance
- System resource requirements for Resiliency Platform
- Network and firewall requirements

Supported hypervisors for deploying Resiliency Platform virtual appliance

This section lists the hypervisor versions that are supported for Resiliency Platform virtual appliance.

Microsoft Hyper-V:

- Windows Server 2012 with Hyper-V
- Windows Server 2012 R2 with Hyper-V

VMware:

- ESXi 5.1, 5.5, 6.0
- vCenter Server 5.1, 5.5, 6.0

Note: The lists of supported platforms for deployment of virtual appliance and for disaster recovery of virtual machines are different. For information on platform support for disaster recovery of virtual machines, see the *Veritas Resiliency Platform Hardware and Software Compatibility List*.

See “About deploying the Resiliency Platform virtual appliance” on page 24.

System resource requirements for Resiliency Platform

The amount of virtual CPUs, memory, and disk space that Veritas Resiliency Platform requires are listed in this section.

The minimum configuration that is recommended for a virtual appliance for Resiliency Manager, Infrastructure Management Server (IMS), Replication Gateway, and YUM repository server:

Table 4-1 Minimum configurations

Component	Minimum configuration
Resiliency Manager	Disk space 60 GB RAM 16 GB Virtual CPU 8
Infrastructure Management Server (IMS)	Disk space 60 GB RAM 16 GB Virtual CPU 8
Gateway	Disk space 40 GB RAM 16 GB Virtual CPU 8 Additional external disk of 50 GB

Note: You need to reserve the resources for Resiliency Manager and IMS to ensure that these resources do not get swapped in case of the hypervisor getting overloaded.

If the virtual appliance does not meet the minimum configuration, you get a warning during the bootstrap of the virtual appliance and you are required to confirm if you want to continue with the current configuration.

If you plan not to use the YUM virtual appliance, you need an Linux server with a minimum of 50-GB disk space, to be configured as the repository server. Provisioning for the repository server is optional, it is required to install the Veritas Resiliency Platform patches or updates in the future.

See “Setting up the repository server ” on page 224.

If you want to enable dynamic memory on Hyper-V, make sure that the following prerequisites are met:

- Startup memory and minimal memory should be equal to or greater than the amount of memory that the distribution vendor recommends.
- If you are using dynamic memory on a Windows Server 2012 operating system, specify Startup memory, Minimum memory, and Maximum memory parameters in multiples of 128 megabytes (MB). Failure to do so can lead to dynamic memory failures, and you may not see any memory increase in a guest operating system.

  Even if you are using dynamic memory, the above mentioned minimum configuration should be met.

See "About deploying the Resiliency Platform virtual appliance" on page 24.

Network and firewall requirements

The following are the network requirements for Veritas Resiliency Platform:

- Before you use the hostname and the IP address in the **Network settings**, you need to register them with the DNS server.
- The hostname or the IP address which is used for product configuration, should not have multiple entries in the DNS server. For example, the IP address should not be associated with multiple hostnames, or the hostname should not be associated with multiple IP addresses.
- The hostname that you use for a virtual appliance must not start with a digit and must not contain the underscore (_) character.
- Veritas Resiliency Platform supports only Internet protocol version (IPV) 4.
- If you plan to use the DHCP server, the DHCP server should be in the same subnet where you plan to deploy the product.

The following ports are used for Veritas Resiliency Platform:

Table 4-2 Ports used for Resiliency Manager

Ports used	Purpose	For communication between	Direction	Protocol
443	Used for SSL communication	Resiliency Manager and web browser	Browser to Resiliency Manager	TCP

Table 4-2 Ports used for Resiliency Manager *(continued)*

Ports used	Purpose	For communication between	Direction	Protocol
14176	Used for communication between the Resiliency Manager and Infrastructure Management Server (IMS)	Resiliency Manager and IMS	Bi-directional	TCP
7001	Used for database replication	Resiliency Manager and IMS	Bi-directional	TCP
389	Used for communication with LDAP/AD server	Resiliency Manager and LDAP/AD server	Bi-directional	TCP
636	Used for communication with LDAP/AD server	Resiliency Manager and LDAP/AD server	Bi-directional	TCP
3268	Used for communication with LDAP/AD server	Resiliency Manager and LDAP/AD server	Bi-directional	TCP
3269	Used for communication with LDAP/AD server	Resiliency Manager and LDAP/AD server	Bi-directional	TCP
22	Used for communication between remote host to the appliance klish access	Appliance and the hosts	Bi-directional	TCP
123	Used for NTP synchronization	Appliance and the NTP server	Bi-directional	TCP

Table 4-3 Ports used for on-premises IMS and in-cloud IMS

Ports used	Description	For communication between	Direction	Protocol
14176	Used for communication between the Resiliency Manager and Infrastructure Management Server (IMS)	Resiliency Manager and IMS	Bi-directional	TCP

Table 4-3 Ports used for on-premises IMS and in-cloud IMS *(continued)*

Ports used	Description	For communication between	Direction	Protocol
5634	Used for IMS configuration	IMS and the hosts	Bi-directional	TCP
14161	Used for running the IMS console	Resiliency Manager and IMS	Resiliency Manager to IMS	TCP
22	Used for communication between remote host to the appliance klish access Used for remote deployment of the packages on remote UNIX host from IMS	IMS and the hosts	Bi-directional	TCP
135	Used for remote deployment on client computer (inbound)	Host and remote Windows hosts	Bi-directional	TCP
123	Used for NTP synchronization	Appliance and the NTP server	Bi-directional	TCP

Table 4-4 Ports used for on-premises Replication Gateway and in-cloud Replication Gateway

Ports used	Description	For communication between	Direction	Protocol
33056	Used for replication	On-premises virtual machine and Replication Gateway/Storage Proxy	Bi-directional	TCP
5634	Used for communication with IMS	IMS and Replication Gateway/Storage Proxy	Bi-directional	TCP
8089	Used for replication	in-cloud component and on-premises component	Bi-directional	TCP

Table 4-5 Ports used for target Gateway in resync operation

Ports used	Description	For communication between	Direction	Protocol
67	BOOTP server	Target Gateway enabled with DHCP role and physical host	Uni-directional	UDP
68	BOOTP client	Target Gateway enabled with DHCP role and physical host	Uni-directional	UDP
69	TFTP protocol	Target Gateway enabled with PXE role and physical host	Uni-directional	TCP/UDP

Table 4-6 Ports used for virtual machines

Ports used	Description	For communication between	Direction	Protocol
22	Used for communication between remote host to the appliance klish access Used for remote deployment of the packages on remote UNIX host from IMS	IMS and the hosts	Bi-directional	TCP
5634	Used for communication with IMS	IMS and the hosts	Bi-directional	TCP
33056	Used for replication	On-premises virtual machine and Replication Gateway	Bi-directional	TCP

See “About deploying the Resiliency Platform virtual appliance” on page 24.

Section 2

Deploying and configuring the virtual appliances

Deploying the virtual appliances on-premises

This chapter includes the following topics:

- Deploying the virtual appliance through VMware vSphere Client
- Deploying the virtual appliance through Hyper-V Manager

Deploying the virtual appliance through VMware vSphere Client

You can deploy Veritas Resiliency Platform virtual appliance through VMware vSphere Desktop Client or VMware vSphere Web Client using the Open Virtualization Archive (OVA) file that you have downloaded.

To deploy Resiliency Platform through VMware vSphere Desktop Client

1. In the VMware vSphere Desktop Client, click **File** and select **Deploy OVF Template**.
2. Select the source location of the Resiliency Platform virtual appliance OVA file.
3. Specify a name for the virtual machine and location for the deployed template.
4. Select the host or cluster on which you want to deploy the template.
5. Select a destination where you want to store the virtual machine files.
6. Select the format in which you want to store the virtual disks.
7. If you have multiple networks configured, select the appropriate destination network.
8. Review the virtual machine configuration and click **Finish**.

9 If you want to use DHCP as your network, enter the MAC address of the appliance in the DHCP server. For information on obtaining the MAC address of the appliance, see the documentation of VMware vSphere client.

10 Power on the virtual machine.

To deploy Resiliency Platform through VMware vSphere Web Client

1 In the VMware vSphere Web Client, click **vCenter Servers** and select a vCenter Server. Click **Actions > Deploy OVF template**.

2 Select the source location of the Resiliency Platform virtual appliance OVA file.

3 Specify a name and location for the deployed template.

4 Select a cluster, host, vApp, or resource pool in which to run the deployed template.

5 Select a location to store the files for the deployed template.

6 Configure the networks the deployed template should use.

7 Review the virtual machine configuration and click **Finish**.

8 If you want to use DHCP as your network, enter the MAC address of the appliance in the DHCP server. For information on obtaining the MAC address of the appliance, see the documentation of VMware vSphere client.

9 Power on the virtual machine.

You can now configure the Resiliency Platform component.

See “About configuring the Resiliency Platform components” on page 51.

Deploying the virtual appliance through Hyper-V Manager

You can deploy Veritas Resiliency Platform virtual appliance through Hyper-V Manager using the Virtual Hard Disk (VHD) files that you have downloaded. There are two VHD files used for deploying the Resiliency Platform virtual appliance.

To deploy Resiliency Platform through Hyper-V Manager

1 Download the Hyper-V supported VHD file for the Resiliency Platform virtual appliance on a system where Hyper-V Manager is installed.

See “Downloading the Resiliency Platform virtual appliances” on page 25.

2 In the Hyper-V Manager console, right-click the Hyper-V server and select **New Virtual Machine**.

3 Provide a name for the virtual machine.

4 Select **Generation 1** while specifying generation.

5 Assign minimum 16 GB RAM.

6 Select a network adapter for the virtual machine.

7 Select the option **Attach a virtual hard disk later** while specifying option to connect virtual hard disk.

8 Review the virtual machine configuration details and click **Finish**.

9 Go to **Settings**, and increase the number of virtual processors as **8**.

10 Add both the VHD files of the Resiliency Platform virtual appliance as **IDE Controller 0**.

11 Click **Apply**, and then click **OK**.

12 If you want to use DHCP as your network, enter the MAC address of the appliance in the DHCP server. For information on obtaining the MAC address of the appliance, see the documentation of Hyper-V Manager.

13 Right-click the name of the virtual machine and select **Start** to power on the virtual machine.

You can now configure the Resiliency Platform component.

See “About configuring the Resiliency Platform components” on page 51.

Chapter 6

Deploying the virtual appliance in AWS

This chapter includes the following topics:

Deploying the Resiliency Platform virtual appliances in AWS

Following is an overview of the key steps that are performed for deploying the Resiliency Platform virtual appliances in Amazon Web Services (AWS):

Table 6-1 Overview of deployment process in AWS

Step	Action	Description
1	Ensure that the prerequisites for deploying virtual appliances in AWS are met.	See "Prerequisites for deploying the Resiliency Platform virtual appliances in AWS" on page 47.

Table 6-1 Overview of deployment process in AWS *(continued)*

Step	Action	Description
2	Upload the OVA files to Amazon S3	See "Uploading the OVA file using web-based method" on page 48. See "Uploading the OVA file using command-line method" on page 48.
3	Create AMI using EC2	See "Creating Amazon Machine Image" on page 49.
4	Launch the instances of virtual appliances to deploy Resiliency Manager, Infrastructure Manager (IMS), and Replication Gateway	See "Launching the instances of virtual appliances" on page 50.

Prerequisites for deploying the Resiliency Platform virtual appliances in AWS

Following are the prerequisites for deploying the Resiliency Platform virtual appliances in AWS:

- Follow the documentation of AWS to create the required security groups. make sure that the security groups meet the network and port requirements mentioned in the Resiliency Platform documentation are open.
- Create a role and grant the permissions required to import an image. Follow the documentation of AWS to know about the permissions required to import an image.
- Ensure that there is direct communication between the premise network and the AWS network. It is recommended to use VPN for AWS environment.
- The virtual appliances such as Resiliency Manager, IMS, and Replication Gateway should be placed in a private subnet in AWS, and not in a public subnet.
- Ensure that Resiliency Manager and Infrastructure Management Server (IMS) have outgoing internet access enabled. It is recommended to restrict incoming internet access on these virtual appliances.

See "Deploying the Resiliency Platform virtual appliances in AWS" on page 46.

Uploading the OVA file using web-based method

You can create a S3 bucket and upload the ova file to that bucket using a web-based method.

To upload the OVA file using web-based method

1. Log in to the AWS console and go to **Services**.
2. Go to **S3**, and click **Create a bucket**.
3. Enter a name for the bucket and select the appropriate region. Click **Create**.
4. Once the bucket gets created, open the bucket and click **Upload**. Click **Add files** and then select the OVA file from your local disk.
5. Click **Start Upload**.

See “Deploying the Resiliency Platform virtual appliances in AWS” on page 46.

Uploading the OVA file using command-line method

You need to first create a S3 bucket in AWS and then upload your ova file to that bucket.

To upload the OVA file using command-line method

1. Download and install the AWS Command Line Interface.
2. Use the `aws s3 mb` command to create a new bucket. Bucket names must be unique and should be DNS compliant:

   ```
   aws s3 mb s3://my-bucket --region my-region
   ```

   where, *my-bucket* is the name that you provide for your bucket and *my-region* is the region that you provide.

   If you do not use the region option of the command, the bucket is created in the region specified in your configuration file.
3. Upload the OVA file by running the following command:

   ```
   aws s3 cp my-ova_file s3://my-bucket/my-ova-key
   ```

   Where, *my-ova_file* is the path and name of your local ova file, *my-bucket* is the name of your bucket on S3 storage and *my-ova-key* is the key or alias name for the ova file in your bucket.

See “Deploying the Resiliency Platform virtual appliances in AWS” on page 46.

Creating Amazon Machine Image

Once you upload the OVA files to Amazon S3 bucket, you need to use the AWS command line interface (CLI) to create an Amazon Machine Image (AMI) from the OVA files that you have uploaded. This AMI can be later used to launch the instances for deploying Resiliency Manager, Infrastructure Manager (IMS), and Replication Gateway in AWS.

To create Amazon Machine Image

1. Go to the Command prompt and then go to AWS CLI.
2. Refer to the AWS documentation for instructions on how to enter your AWS credentials and region and create a json file in the following format:

```
[
         {
           "Description":"my description",
           "Format": "ova",
           "UserBucket":{
             "S3Bucket":"my-bucket",
             "S3Key": "my-ova-key"
            }
       }]
```

   Where, *my-bucket* is the name of your bucket and *my-ova-key* is the alias name that you provided for the OVA file.

3. Run the following command to create an AMI:

```
  aws ec2 import-image --description "my description"
 --disk-containers file://Mycontainers.json_with_path
```

   Where, *Mycontainers.json_with_path* is the path and name of the json file that you have created.

4. The above command displays a number of parameters and their values. Note down the value of **ImportTaskId** parameter.
5. Run the following command to verify that the import task is complete and the AMI is ready to be used:

```
aws ec2 describe-import-image-tasks
--import-task-ids MyImportTaskID
```

   Where, *MyImportTaskID* is the task ID that you receive from the command described in the prior step.

See “Deploying the Resiliency Platform virtual appliances in AWS” on page 46.

Launching the instances of virtual appliances

Once an Amazon Machine Image (AMI) gets created, you can use the AMI to launch instances to deploy the Resiliency Manager and any number of Infrastructure Management Servers (IMS), and Replication Gateways in AWS.

To launch the instances of virtual appliances

1. Go to the command prompt and open the AWS console. Go to **Services** and then go to the EC2 console.
2. In the left hand side pane, under **IMAGES**, click **AMIs** and you can see the list of AMIs created.
3. Select the AMI that you want to use and click **Launch**. Make sure to select an instance type that matches with the system resource requirements mentioned in the documentation:

   See “System resource requirements for Resiliency Platform” on page 37.

   For example, you can select instance types m3.2xlarge and m4.2xlarge. Network Optimization should be high for the instance.

See “Deploying the Resiliency Platform virtual appliances in AWS” on page 46.

Chapter 7

Configuring the Resiliency Platform virtual appliance

This chapter includes the following topics:

About configuring the Resiliency Platform components

After the Veritas Resiliency Platform virtual appliance deployment, you are expected to configure the Resiliency Platform component that you have deployed, through the bootstrap process. The bootstrap process is automatically invoked when you log in to the virtual appliance console for the first time using the admin user login. The following settings are configured as part of this process to set up the component:

- **Host Network settings:** Settings such as hostname, IP address, subnet mask, default gateway, and DNS server. Before using the hostname and the IP address, you need to register them with the DNS server.
- **Appliance settings:** Settings such as NTP server.
- **Product settings:** Whether you want to configure a Resiliency Manager, Infrastructure Management Server (IMS), or Resiliency Manager and IMS both.

Note: The hostname and the IP address that you use for product configuration, must not be changed later.

This configuration is done through the bootstrap process only for the first time. After the successful configuration, the bootstrap process is disabled. The subsequent admin user logins to the virtual appliance will automatically start with Command Line Interface SHell (klish) menu. If you want to change these settings later, you can use klish menu for changing these settings.

See “Configuring the Resiliency Manager or IMS” on page 52.

Prerequisites for configuring Resiliency Platform components

Before configuring the component through product bootstrap, make sure that following prerequisites are met:

- In case of multiple Resiliency Managers, make sure that the NTP servers that are used for configuration of Resiliency Managers are properly synchronized.
- Make sure that you have disabled the dynamic or the automatic MAC address change for your hypervisor. Follow the documentation of your hypervisor to set the MAC address manually or to disable the setting for automatic MAC address change.
- To use DHCP network, you need to reserve an IP address for the appliance in the DHCP server along with the corresponding MAC address.
- Before you use the hostname and the IP address in the Network settings, make sure that the reverse lookup for that IP works.
- In case of a Replication Gateway, make sure to attach an extra disk of at least 50 GB before configuring the Gateway.

Configuring the Resiliency Manager or IMS

After Veritas Resiliency Platform (Resiliency Platform) deployment, when you log into the virtual appliance console for the first time using the admin user credentials, the bootstrap process is automatically invoked. This bootstrap process is used to set up or configure the Resiliency Platform component for the first time.

The default network protocol for virtual appliance is Dynamic Host Configuration Protocol (DHCP). If the appliance detects DHCP during the first boot or before the completion of bootstrap process, the appliance network automatically gets configured. After the network configuration, you can either use the virtual appliance console or Secure Shell (SSH) to log in and complete the bootstrap process.

If DHCP is not configured in your environment, you have an option to use a static IP for the appliance. Since the appliance network is not automatically configured in this case, you can only use the console to log into the virtual appliance.

To configure the Resiliency Manager or IMS

1. Log in to the virtual appliance console or SSH using the following credentials:

   - **Username:** admin
   - **Password:** P@ssw0rd

   After a successful login, you are prompted to change the password of the admin user.

   See “Password policies for Resiliency Platform virtual appliance” on page 57.

   If you are logged in to SSH, you will be logged off the SSH session after the password change and you need to again log in to complete the rest of the steps of the bootstrap process. If you are logged in to the virtual appliance console, you can continue and complete the rest of the steps of the bootstrap process.

2. Accept the End User License agreement (EULA) to proceed with the configuration.

3. In the **Host Network Settings** section, you can configure the appliance network by using DHCP or static IP.

   If the DHCP configuration is working in the environment, the details are printed on the screen. Confirm if you want to proceed with these settings.

   In case of DHCP, you need to ensure that a Dynamic Host Configuration Protocol (DHCP) server is working in the subnet where the virtual appliance is deployed. In case of static IP, you need to respond to the following additional prompts:

   - **Enter the fully qualified hostname:**
   - **Enter the IP address:**
   - **Enter the Subnet mask:**
   - **Enter the Default Gateway:**
   - **Enter the DNS server (space separated if more than one DNS, maximum 2 DNS entries):**

4. Confirm if you are in Network Address Translation (NAT) environment and want to verify the NAT configuration using an external IP and hostname. You need to respond to the following additional prompts:

- **Enter the fully qualified external hostname:**
- **Enter the external IP address:**

Note: The external hostname or IP that you provide are set as the appliance hostname or IP.

For recovery to AWS, if the VPN is configured between the sites using the private network, then NAT option will not be applicable.

5 In the **Appliance Settings** section, do the following:

- Press the Enter key to confirm the use of an NTP server for configuring the date and time.
- You are required to select the time zone. Follow the instructions as displayed on the virtual appliance console or SSH session to select the correct time zone.
- Enter the FQDN or IP address of the NTP server. The appliance verifies the NTP server details. If there are any issues, details are displayed on the screen and you are prompted to enter the details again.

You can reset the timezone and NTP server at a later point of time using klish menu. Changing the system settings can affect the product functionality if incorrect values are set.

6 In the **Product Settings** section, enter your choice for configuring the virtual appliance as a Resiliency Manager, Infrastructure Management Server (IMS), or both. Type **1** for configuring the role of Resiliency Manager, **2** for configuring the role of IMS, and **3** for both. For production environment, it is recommended not to configure Resiliency Manager and IMS together on one appliance, but for test or evaluation purposes, you can configure a Resiliency Manager and Infrastructure Management Server (IMS) on the same virtual appliance.

7 After a successful product configuration, a message is displayed. If you have configured Resiliency Manager on the virtual appliance, a URL for the Resiliency Platform web console login is provided. You can type the URL in a web browser and log in to the web console.

See “About configuring the Resiliency Platform components” on page 51.

Configuring the Replication Gateways

After the virtual appliance deployment, when you log into the virtual appliance console for the first time using the admin user credentials, the bootstrap process is automatically invoked. This bootstrap process is used to set up or configure the Resiliency Platform component for the first time.

The default network protocol for virtual appliance is Dynamic Host Configuration Protocol (DHCP). If the appliance detects DHCP during the first boot or before the completion of bootstrap process, the appliance network automatically gets configured. After the network configuration, you can either use the virtual appliance console or Secure Shell (SSH) to log in and complete the bootstrap process.

If DHCP is not configured in your environment, you have an option to use a static IP for the appliance. Since the appliance network is not automatically configured in this case, you can only use the console to log into the virtual appliance.

To configure a Replication Gateway

1. Prerequisites
   - The Replication Gateway configuration requires an extra disk of at least 50 GB. You can attach this before or during the configuration. The default disk size of 50GB lets you protect up to 48 virtual machines and each additional virtual machine requires a staging disk of 12GB size. You can increase the size of the disk by using klish commands:
     See "Using klish" on page 256.
   - The configuration requires that you add the Replication Gateway to an Infrastructure Management Server (IMS) on the same data center. Ensure that you have the hostname or IP address and admin login for the IMS. The Replication Gateway can be added to only one IMS. An IMS can have more than one Replication Gateway added to it.
   - The On-premise Gateway and the virtual machines to be protected need to be on a shared data store.
2. Log in to the virtual appliance console or SSH using the following credentials:
   - **Username:** admin
   - **Password:** P@ssw0rd

   After a successful login, you need to change the password of the admin user.

   See "Password policies for Resiliency Platform virtual appliance" on page 57.

   If you are logged in to SSH, you will be logged off the SSH session after the password change and you need to again log in to complete the rest of the steps of the bootstrap process. If you are logged in to the virtual appliance console, you can continue and complete the rest of the steps of the bootstrap process.
3. Accept the End User License agreement (EULA) to proceed with the configuration.

4 In the **Host Network Settings** section, you can configure the appliance network by using DHCP or static IP.

If the DHCP configuration is working in the environment, the details are printed on the screen. Confirm if you want to proceed with these settings.

In case of DHCP, you need to ensure that a Dynamic Host Configuration Protocol (DHCP) server is working in the subnet where the virtual appliance is deployed. In case of static IP, you need to respond to the following additional prompts:

- **Enter the fully qualified hostname:**
- **Enter the IP address:**
- **Enter the Subnet mask:**
- **Enter the Default Gateway:**
- **Enter the DNS server (space separated if more than one DNS, maximum 2 DNS entries):**

5 Confirm if you are in Network Address Translation (NAT) environment and want to verify the NAT configuration using an external IP and hostname. You need to respond to the following additional prompts:

- **Enter the fully qualified external hostname:**
- **Enter the external IP address:**

Note: The external hostname or IP that you provide are set as the appliance hostname or IP.

---

6 In the **Appliance Settings** section, do the following:

- Press the Enter key to confirm the use of an NTP server for configuring the date and time.
- You are required to select the time zone. Follow the instructions as displayed on the virtual appliance console or SSH session to select the correct time zone.
- Enter the FQDN or IP address of the NTP server. The appliance verifies the NTP server details. If there are any issues, details are displayed on the screen and you are prompted to enter the details again.

You can reset the timezone and NTP server at a later point of time using klish menu. Changing the system settings can affect the product functionality if incorrect values are set.

7 **Product Settings**:

Configure the role of Replication Gateway on this virtual appliance.

8 You are prompted to attach an extra disk to the appliance. If you have already attached the extra disk, press **Enter** to confirm. Otherwise, attach the extra disk and then confirm or select the extra disk to be used.

While attaching the extra disk to the Replication Gateway appliance in AWS, use the full device path and use the format xvdb[a-z] instead of sd[a-z]. For example use /dev/xvdba instead of just xvdba.

9 You are prompted to enter the hostname or IP address of the IMS that you want to connect to this Replication Gateway. Enter the required information and the admin password for the IMS. Typically, a gateway reports to an IMS server that is part of the same data center.

10 After a successful product configuration, a confirmation message will be displayed and you will be logged out of the virtual appliance console.

Password policies for Resiliency Platform virtual appliance

The following is a list of rules that you need to follow for the password of an admin user login for virtual appliance:

- Must be at least 8 characters long.
- Must contain at least one uppercase letter (A-Z), one lowercase letter (a-z), one numeric (0-9), and one special character such as @&%.
- Cannot contain the user name or its characters in reversed order.
- Cannot contain same character used consecutively for more than 2 times.
- Cannot contain 5 or more characters from the previous password.
- Cannot be the same as your previous 6 passwords.
- Can be changed after a minimum of 15 days since the last password change. The password can be changed only through klish menu.
- Expires in 90 days. you get an error message when you are not able to login using admin user credentials.
- 7 days before the password expiry date, a warning is provided to change the password. This warning is not displayed in the Resiliency Manager console. You get to see this warning only when you log in to the virtual appliance console or in the SSH session using the admin user credentials.

Section 3

Setting up and managing the resiliency domain

Chapter 8

Setting up the resiliency domain

This chapter includes the following topics:

Setting up the resiliency domain

The following tasks are required to set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.

Table 8-1 Tasks for setting up the resiliency domain

Task	More information
Set up the Resiliency Manager and create the resiliency domain	When you first log in to the console after deploying the Resiliency Manager virtual appliance, the Getting Started wizard requires you to set up the Resiliency Manager and create the resiliency domain. Other steps on the Getting Started wizard are optional and can also be done from the console **Settings** page after you complete the Getting Started wizard. See "Connecting to the Resiliency Platform web console" on page 61.
Add Infrastructure Management Servers (IMS)	If you did not add the IMSs using the Getting Started wizard, you can add them later. See "Adding an IMS " on page 70.
Add the asset infrastructure to the IMS	See "Adding the asset infrastructure" on page 106.
Set up user authentication and personas	See "Managing user authentication and permissions" on page 175.
Configure settings for notifications and other general product settings	See "Managing settings for alerts and notifications and miscellaneous product settings" on page 208.

About the web console

Once you have finished configuring the virtual appliance settings for the Resiliency Manager on the bootstrap menu, you log in to the Veritas Resiliency Platform web console to continue with setting up the resiliency domain.

See "Connecting to the Resiliency Platform web console" on page 61.

Note: For the best console experience, use a minimum resolution of 1280x1040.

You must complete the basic configuration of the Resiliency Manager and the resiliency domain using the Getting Started wizard before you have access to the remainder of the web console.

See "Getting started with a new Resiliency Platform configuration" on page 62.

From that point, any time you log in, you can view the full web console screen and menus.

See “Tour of the Resiliency Platform web console screen” on page 266.

Connecting to the Resiliency Platform web console

Once the Resiliency Manager virtual appliance is deployed and configured, you can connect to the web console.

To connect to the Resiliency Platform web console

1. Prerequisites

   - A supported web browser on a system that has a network connection to the Resiliency Manager

     For information on web browser requirements, refer to the Hardware and Software Compatibility List (HSCL).

     Your browser must be configured to accept cookies and enabled for JavaScript. If you use pop-up blockers, either disable them or configure them to accept pop-ups from the Resiliency Manager or Infrastructure Management Server (IMS) host.

   - Login credentials, including the domain

     The initial credentials that are required are for the Admin user of the Resiliency Manager.

     The default domain for the Admin user is vrp.local.

     Once the Admin user configures Resiliency Platform to use an LDAP or Active Directory authentication broker and configures user access, users can login with their credentials for that authentication domain.

2. Type the URL as follows:

   https://*hostname*

   Example: https://myhost.example.com

3. Enter your login credentials, including the domain, and select **Login**.

   A configurable setting determines whether the login screen shows a list of domains for user selection. By default, domains are not listed and users must enter a fully qualified username as username@domain or domain\username.

   For example, for the Admin login, enter:

   admin@vrp.local

See “Getting started with a new Resiliency Platform configuration” on page 62.

See “Showing domains on login screen ” on page 188.

Getting started with a new Resiliency Platform configuration

When you first log in to the web console on a new Resiliency Manager, a Getting Started wizard helps you to set up a basic Resiliency Platform configuration.

The following table shows the steps involved in getting started with the first Resiliency Manager and creating a new resiliency domain.

The procedure for adding a Resiliency Manager to an existing resiliency domain is covered in a separate topic.

See “Adding a Resiliency Manager to an existing resiliency domain” on page 65.

Prerequisites:

- The basic configuration includes the Resiliency Manager and Infrastructure Management Server (IMS).
  See “About Resiliency Platform features and components” on page 15.
- If the IMS is on a separate virtual appliance from the Resiliency Manager, ensure that you have the fully qualified host name and login credentials for the IMS virtual appliance. Optionally, you can add the IMS later.

Table 8-2 Getting Started wizard

Wizard step	Details
1. Set up Resiliency Manager	Specify the data center location of the Resiliency Manager, the data center friendly name, and the Resiliency Manager name. Default entries are shown if the Resiliency Manager has external Internet access to determine the geographical location. Click **Confirm & Continue**.
2. Create or Join a Resiliency Domain	For a new Resiliency Platform deployment, select the option to create a resiliency domain and supply a name for the domain. You can choose whether to allow collection of product usage information. Click **Create**. Wait for the message showing that the domain is successfully created. This process may take several minutes. More information is available about telemetry collection. See “Enabling or disabling telemetry collection ” on page 214.

Table 8-2 Getting Started wizard *(continued)*

Wizard step	Details
3. Enable Solutions Licenses	You can select a license file to apply or enable the trial license. See "About licenses" on page 171.
4. Set up Authentication Domain	Optional. By default the Admin user on the virtual appliance has the Super admin persona. Personas are user roles with access to a predefined set of operations. The Super admin persona has full access to all operations in the console. If you want to assign a different user as Super admin you must first set up an LDAP or Active Directory authentication domain. Then, on the next step, you can add a user or group from that identity provider as Super admin and optionally reassign the virtual appliance Admin user to a more limited persona. Otherwise, you can skip this step and set up authentication and assign personas later using the console **Settings** page.
5. Set up Users and Personas	Optional. If you set up an authentication domain in the previous step, you can specify the user or user group to which you want to assign the Super admin persona. Optionally, you can also reassign the virtual appliance Admin to the more limited Resiliency Platform Deployment admin persona, with permission to perform deployments and updates only. The user with the Super Admin persona can add other users and groups and assign them personas later using the **Settings** page. See "Managing user authentication and permissions" on page 175.

Table 8-2 Getting Started wizard *(continued)*

<table>
<tr><th>Wizard step</th><th>Details</th></tr>
<tr><td><b>6. Add Infrastructure Management Servers</b></td><td>Optional.
Add an Infrastructure Management Server (IMS). Optionally you can add more than one.
You can also add an IMS later from the <b>Settings</b> page.
■ Choose one of the following
■ To add an IMS co-located with the Resiliency Manager, choose <b>Enable internal IMS</b>.
■ If you deployed an IMS separate from the Resiliency Manager, choose <b>Connect to IMS</b>.
■ Fill in the following information and click <b>Add</b>:
<table>
<tr><td>Data Center Location</td><td>Select the data center location, for example, the city.
To specify a new data center, select <b>New</b> and then specify the location and name. For the location, enter location identifier, such as city, and the location list populates with potential matches for you to select.</td></tr>
<tr><td>Server Name</td><td>If you are adding an IMS separate from the Resiliency Manager, enter the fully qualified host name.
For the login, use the Admin login credentials for the IMS virtual appliance.</td></tr>
<tr><td>Friendly Name</td><td>Enter a user-friendly name for the IMS. This name helps identify the IMS in the console.</td></tr>
</table></td></tr>
<tr><td><b>7. Set up Cloud Configuration</b></td><td>Optional.
This step is enabled only if you select <b>Cloud data center</b> in the Getting Started wizard. You can skip this step and add the cloud configuration later from the console.
If you have added an IMS to the cloud data center, The IMS is listed in this wizard. Enter the cloud configuration name, AWS access key, AWS secret key, region, and name of the S3 bucket that can be accessed using the AWS credentials that you have provided.
The wizard verifies the information you enter and notifies you if the information is invalid.</td></tr>
</table>

Table 8-2 Getting Started wizard *(continued)*

Wizard step	Details
	Exit the Getting Started wizard. The Dashboard is displayed. You can complete any steps that you skipped from the **Settings** page. See "About settings in the web console" on page 270.

Adding a Resiliency Manager to an existing resiliency domain

If you are using Resiliency Platform for disaster recovery, you deploy a Resiliency Manager on both a production data center and a recovery data center. When adding the first Resiliency Manager, you create a resiliency domain. You must add the second Resiliency Manager to the existing resiliency domain.

To add a Resiliency Manager to an existing resiliency domain

1 Prerequisites:

- Deploy a new Resiliency Platform virtual appliance node. During deployment, specify the node as either Resiliency Manager only or both Resiliency Manager and Infrastructure Management Server (IMS).
- Ensure that you have the fully qualified host name/IP address and the Admin login credentials for an existing Resiliency Manager virtual appliance in the resiliency domain.

2 Log in to the web console on the new Resiliency Manager. The Getting Started wizard is displayed.

3 In **Set up Resiliency Manager**, specify the data center location, the data center friendly name, and Resiliency Manager friendly name. Click **Confirm & Continue**.

4 In **Create or Join a Resiliency Domain**, select **Join resiliency domain**.

Enter the fully qualified host name or IP address of a Resiliency Manager in the domain you want to join, and click **Verify**.

5 Once the name or address has been verified as a Resiliency Manager, the login fields are available. Enter the credentials for that Resiliency Manager and click **Join**.

A confirmation message shows the name of the resiliency domain that you are joining. Wait for the message that shows that the domain has been joined.

6. You have completed the Getting Started steps that are required for the new Resiliency Manager. Optionally you can add an Infrastructure Management Server, or you can do so later from the **Settings** page.

   See “Adding an IMS ” on page 70.

7. If you refresh the page in the web console of the new Resiliency Manager, the information for the domain that you joined is shown in the Dashboard

   Each Resiliency Manager in the domain has its own web console but the data that is shown is synchronized with other Resiliency Managers in the domain.

Removing a Resiliency Manager from a resiliency domain

A Veritas Resiliency Platform resiliency domain typically contains two Resiliency Managers. You can remove a Resiliency Manager from the domain as long as another remains online.

Removing a Resiliency Manager is necessary, for example, if you need to do the following:

- Change the host name or IP address of the Resiliency Manager virtual appliance.
- Change a virtual appliance that is both a Resiliency Manager and Infrastructure Management Server (IMS) so that it is used only as an IMS.

Caution: Ensure that you meet the prerequisites listed in the procedure.

For example, if you want to decommission a Resiliency Manager virtual appliance node, you do the following steps:

- If the virtual appliance node that you want to decommission is configured as both a Resiliency Manager and Infrastructure Management Server (IMS), first remove the IMS from the resiliency domain.
  See “Removing an IMS” on page 72.
- Remove the Resiliency Manager from the resiliency domain using the Leave Domain procedure below.
  Completing this operation ensures that the Resiliency Manager is cleanly decommissioned and that all references to it are removed from the Resiliency Manager database and no longer appear in the web console user interface.
- If there is an IMS on a separate node that was reporting to the Resiliency Manager, ensure that it is reconnected to another Resiliency Manager.

To remove a Resiliency Manager from a resiliency domain

1 Prerequisites

- Both Resiliency Managers must be online.
- Perform the operation from the Resiliency Manager that is remaining in the resiliency domain, not from the Resiliency Manager that is being removed.
- You should perform the operation during a maintenance window and send appropriate notifications in advance.
- Ensure that no activity is occurring on the Resiliency Manager that you plan to remove. For example, ensure that no workflow is in progress.

2 Navigate

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

3 Expand the data center, locate the Resiliency Manager, and do the following:

⋮ Select the vertical ellipsis icon next to the Resiliency Manager and select **Leave Domain**.

The operation can take over five minutes to complete as it is a multistep process.

4 Once the operation is successfully completed, you can remove the Resiliency Manager virtual appliance node using the appropriate hypervisor manager.

See “Troubleshooting removing a Resiliency Manager from a resiliency domain” on page 249.

Viewing the status of a Resiliency Manager in a data center

In the web console, you can view a list of data centers and any associated Resiliency Manager and IMS. Under the Resiliency Manager you can view the status, as follows:

Connected

The Resiliency Manager is up and healthy, and if there is another Resiliency Manager node in the domain, they are connected.

Disconnected	The Resiliency Manager node is down, the Resiliency Manager services are not running, or there is a connection issue between the Resiliency Manager nodes. See “Troubleshooting the connection between Resiliency Managers in a resiliency domain” on page 249.
Leaving Domain	The Leave Domain operation has been initiated and is in progress. See “Removing a Resiliency Manager from a resiliency domain” on page 66.
Leave domain failure	The Leave Domain operation did not complete successfully.

To view the status of a Resiliency Manager in a data center

1. Navigate

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

2. You can expand or contract each data center listed. Use the arrow on the right side of a data center row to expand the data center and view the information about the Resiliency Manager.

Chapter 9

Managing Infrastructure Management Servers

This chapter includes the following topics:

How Infrastructure Management Servers relate to data centers

Veritas Resiliency Platform provides for the resiliency management of virtual machines or applications by data center. Virtual machines and applications are associated with other infrastructure, such as physical hosts, virtualization servers, and storage arrays. Resiliency Platform includes an Infrastructure Management Server (IMS) to discover, monitor, and manage assets in a data center.

See “Infrastructure Management Server (IMS)” on page 21.

A Resiliency Platform domain can extend across data centers in different locations. Each data center has at least one IMS. A data center can also have more than one IMS. You determine which infrastructure assets to add to each IMS.

In the Resiliency Platform web console, you associate each IMS with a data center location and a data center name (friendly name). If a data center has more than one IMS, the best practice is to associate each IMS with the same data center location and name.

Adding an IMS

Veritas Resiliency Platform includes an Infrastructure Management Server (IMS) to discover and monitor assets. When you first configure Resiliency Platform in the web console, you set up the Resiliency Manager and resiliency domain with the Getting Started wizard. Optionally, you can also add one or more IMSs. You can also add IMSs later, after you exit the Getting Started wizard. This procedure describes how to add IMSs later.

To add an IMS

1 Prerequisites

- A Resiliency Manager and resiliency domain must be set up using the Getting Started wizard.

  See “Getting started with a new Resiliency Platform configuration” on page 62.

- The virtual appliance for the IMS must be deployed and configured.

- Information needed for adding the IMS:

  The fully qualified domain name (FQDN) or IP address.

  The admin user credentials for the IMS virtual appliance. This information is optional and you need to enter only if the server is directly accessible.

  If the server is not directly accessible, you can still initiate the process of adding an IMS by entering only the data center, friendly name, and FQDN/IP address. In this case, you get a registration URL which you have to use after logging in to the virtual appliance console of the IMS that you want to add and then the IMS is added to the data center.

- Ensure that the IP address and hostname of the IMS gets resolved from the Resiliency Manager.

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

Select **+ Infrastructure Management Server**.

You can also access this page from the **Quick Actions** menu > **Manage Asset Infrastructure**.

3 In **Add Infrastructure Management Server**, enter the information for the IMS and submit.

Tips:

You can select from a list of existing data centers or add a new data center.

See “How Infrastructure Management Servers relate to data centers” on page 69.

- To specify a new data center, select **New** in the **Data Center** field, then specify the location and name. When entering the location, enter a form of location identifier, such as city, and the location list will populate with potential matches for you to select.
- Enter a friendly name for the IMS.
- Enter the FQDN or IP address of the server.
- Enter the user name and the password. These two are optional information that you need to enter only if the IMS is directly accessible. If you provide this information, the IMS is immediately added to the data center.
  If you do not provide the username and password of the IMS, a registration URL is displayed on the screen. This URL is valid only for 30 minutes. if the URL expires, you need to regenerate the registration URL to complete the process.
  See “Regenerating registration URL for an IMS” on page 72.
  Copy the URL string and then log in to the virtual appliance console of the IMS. In the klish menu, use that URL string as a parameter to the following command:

  ```
  manage > configure ims_register registration URL
  ```

  Where, registration URL is the URL obtained after initiating the process form Resiliency Manager console.
  See “Using klish” on page 256.

4 Verify that the IMS is successfully added.

Once the IMS is successfully added, you can add the asset infrastructure to the IMS.

See “Adding the asset infrastructure” on page 106.

5 If you add an IMS to an existing data center after the DNS settings for the data center have been configured, go to the DNS settings for the data center, select the modify option for the DNS server, enter a test host name and IP address, and run a test. This ensures that this newly added IMS can be used to perform DNS updates.

Regenerating registration URL for an IMS

When you add an IMS (Infrastructure Management Server) to a data center using the registration URL, the URL is valid only for 30 minutes. If the URL expires, you can regenerate the registration URL and then complete the operation.

To regenerate registration URL for an IMS

1 Navigate

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

2 Under the data center, locate the IMS, and do the following:

⋮ Select the vertical ellipsis for the IMS > **Regenerate Registration URL**.

3 In the **Regenerate Registration URL** window, click **Submit**.

Removing an IMS

In the web console, you can remove an Infrastructure Management Server (IMS) from a resiliency domain. Removing an IMS removes the information about the IMS in the Resiliency Manager repository for the resiliency domain. It also removes the assets such as virtual machines and applications that were discovered by that particular IMS.

The remove IMS operation does not remove the asset Infrastructure configurations such as configurations of a vCenter server, enclosures, hyper-V server. These asset Infrastructure configurations can be used again for any new or re-provisioned IMS.

Caution: Remove IMS operation does not perform any cleanup of the assets configured with the IMS. It is recommended not to remove an IMS and add it back without re-provisioning it with Resiliency Manager since it may lead to incorrect configurations. You may use **Reconnect** operation in case you want to remove and then re-add the same IMS.

To remove an IMS

1 Navigate

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

2 Under the data center, locate the IMS, and do the following:

⋮ Select the vertical ellipsis for the IMS > **Remove**. Confirm the deletion.

3 Verify that the IMS is removed.

Note: Additional cleanup is optional. If resiliency groups were created for the assets managed by the IMS, and they are no longer needed, you can remove them.

Modifying an IMS

In the web console, you can modify the friendly name associated with an Infrastructure Management Server (IMS) that has been added to a resiliency domain or change the data center name for the IMS. You cannot change the information about the virtual appliance that hosts the IMS.

Modifying an IMS is a separate operation from configuring or modifying assets for an IMS.

See “Adding the asset infrastructure” on page 106.

To modify an IMS

1 Navigate

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

2 Expand the data center, locate the IMS, and do the following:

Select the vertical ellipsis for the IMS > **Edit**.

Tips:

You can select from a list of existing data centers or add a new data center.

See "How Infrastructure Management Servers relate to data centers" on page 69.

To specify a new data center, select **New** in the **Data Center** field, then specify the location and name. When entering the location, enter a location identifier, such as city, and the location list will populate with potential matches for you to select.

3 Verify the change.

Reconnecting the IMS to a Resiliency Manager

You can use the web console to reconnect an Infrastructure Management Server (IMS) to a Resiliency Manager. The operation disconnects any existing connection and the IMS then reinitiates the connection. This operation can be useful when troubleshooting or repairing a connection between an IMS and a Resiliency Manager.

For example, say the connection between an IMS and the Resiliency Manager located in the same data center (data center A) fails. Another Resiliency Manager in the same resiliency domain is online in another data center (data center B). In such a case, the IMS can automatically connect itself to the Resiliency Manager in data center B.

However, once the issue with the first Resiliency Manager is fixed, the administrator may want to reconnect the IMS back with the first Resiliency Manager. The Reconnect operation will accomplish this as follows: If there is an existing connection, the Reconnect operation disconnects it. The IMS then reconnects to a Resiliency Manager in the resiliency domain based on priority. A Resiliency Manager in the same data center as the IMS has a higher priority than a Resiliency Manager in a different data center. Therefore, in the above scenario, the IMS reconnects to the Resiliency Manager in data center A.

To reconnect an IMS to a Resiliency Manager

1 Navigate

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

2 Expand the data center, locate the IMS, and do the following:

⋮ Select the vertical ellipsis for the IMS > **Reconnect**.

Managing data centers

Resiliency Managers and Infrastructure Management Servers (IMSs) must always be associated with a data center location and name. You specify the data center information while setting up the Resiliency Manager and while adding an IMS.

See “How Infrastructure Management Servers relate to data centers” on page 69.

In the web console, you can view a list of data centers and any associated Resiliency Manager and IMS. You can edit the data center information. You can also add data centers separately for use later or delete a data center if it has no Resiliency Manager or IMS associated with it.

Details about configuring network settings for data centers are covered in a separate topic.

Adding data centers

Using Resiliency Platform console, you can add a data center.

To add a data center

1. Navigate

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

2. Select **+ Datacenter**.

3. Specify the location and a friendly name for the data center. When entering the location, enter a form of location identifier, such as city, and the location list will populate with potential matches for you to select.

   If this is a cloud data center, select the checkbox and then select the cloud type. Click **Next**.

   You need to enter details for cloud configuration on the next wizard.

   Select the checkbox if you want to use an existing cloud configuration that is already added to another cloud data center associated with the Resiliency Manager. You need to enter the secret key for verification and need to select a region other than the one already used. You can not use same region for multiple data centers.

   See “Adding a cloud server” on page 78.

   The cloud configuration you enter is validated before it gets saved in the Resiliency Manager.

Note: From the console, you can add multiple IMSs to a single cloud Data center, but this configuration is not supported in Veritas Resiliency Platform 2.1 release.

Editing data centers

Using Resiliency Platform console, you can edit a data center.

If a new virtual private cloud (VPC) is added in the AWS region after the completing the cloud configuration, then that VPC and its components (e.g. subnets, security groups, instances etc) won't be discovered and available in VRP. If you need the new VPC and its components to be discovered and available in VRP, then you need to Edit the cloud config of the DC.

To edit the configuration of a cloud server

1 Navigate

⚙ **Settings** (menu bar)

Under **Infrastructure Settings**, click **Infrastructure**

Click the vertical ellipses next to the cloud data center, and click **Edit Datacenter**.

2 In the **Edit Datacenter** wizard panel, edit the data center details.

3 Select **Is Cloud Datacenter** if the datacenter you wan to edit is a cloud datacenter.

4 If you want to edit the cloud server configuration associated with the data center, select the **Edit AWS cloud configuration** checkbox and click **Next**.

See “Editing the configuration of a cloud server” on page 79.

5 Click **Finish** to close the wizard.

Removing data centers

Using Resiliency Platform console, you can remove a data center.

To remove a data center

1 Navigate

⚙ **Settings** (menu bar)

Under **Infrastructure Settings**, click **Infrastructure**

Click the vertical ellipses next to the cloud data center, and click **Remove Datacenter**.

2 In the **Remove Datacenter** wizard panel, click **Submit**.

Note: Before you remove a Datacenter, you need to verify that there is no associated Resiliency Manager and IMS. You cannot delete a data center that is associated with a Resiliency Manager or IMS. You can edit an IMS to change its data center; however you cannot change the data center associated with a Resiliency Manager.

Managing cloud server configurations

To access the Cloud resources, you need to configure the cloud credentials. These credentials are used to interface with Cloud APIs. You can configure cloud credentials as part of the initial getting started experience or add the cloud server later to the Infrastructure Management Server (IMS) on the cloud.

Cloud APIs perform the following operations:

- Discover virtual machines and volumes in the cloud.
- Invoke operations on the cloud resources such as provision virtual machines, attach or detach disks.

If you add a new virtual private cloud (VPC) to the AWS region after completing the cloud configuration, the VPC and its components are not discovered and available in the console. For discovering the new VPC and its components, you need to edit the cloud configuration, enter your AWS access key and AWS secret key, and then save the configuration.

See “Adding a cloud server” on page 78.

See “Editing the configuration of a cloud server” on page 79.

See “Refreshing the configuration of a cloud server” on page 79.

Adding a cloud server

Using the Veritas Resiliency Platform console, you can add an existing cloud server to the Infrastructure Management Server (IMS) on the cloud.

To add a cloud server

1. In the **Add Cloud Server** wizard panel, enter the following details:
   - Enter a configuration name.
   - Enter the AWS access key.
   - Enter the AWS secret key.
   - Select the region.

- Enter the name of the S3 bucket that already exists in the region and is accessible with the user credentials that you have entered.

2 Click **Finish** to complete the configuration.

See “Managing cloud server configurations” on page 78.

Editing the configuration of a cloud server

Using the Veritas Resiliency Platform console, you can edit the configuration of a cloud server. You can edit only the configuration name and the password.

To edit the configuration of a cloud server

1 Navigate

Do one of the following:

Settings (menu bar)

Under **Infrastructure Settings**, click **Infrastructure**

Under the data center, locate the IMS and click **Manage Asset Infrastructure**.

DR Capability (left navigation bar)

Under **Disaster Recovery Capability** , click **Cloud Configuration**

Click **Manage Cloud IMS**.

2 Click **Cloud access**.

3 Right-click the configuration that you want to edit, and select **Edit**.

4 In the **Edit Configuration** wizard panel, edit the configuration name, enter the AWS access key and AWS secret key. You can not edit the region.

5 Click **Finish** to complete the configuration.

See “Managing cloud server configurations” on page 78.

Refreshing the configuration of a cloud server

Using the Veritas Resiliency Platform console, you can refresh the configuration of a cloud server. This operation re-discovers all the cloud-based objects such as regions, zones, security groups, flavors, virtual machines in the cloud.

To refresh the configuration of a cloud server

1 Navigate

⚙ **Settings** (menu bar)

Under **Infrastructure Settings**, click **Infrastructure**

Click the vertical ellipses next to the cloud data center, and click **Refresh Cloud discovery**.

2 In the **Refresh Configuration** wizard panel, click **Next**.

3 Click **Finish** to close the wizard.

See “Managing cloud server configurations” on page 78.

Chapter 10

Managing NetBackup and NetBackup Appliances

This chapter includes the following topics:

About NetBackup and NetBackup Appliances

About NetBackup

NetBackup provides a data protection solution for a variety of platforms such as Windows, UNIX, and Linux systems. NetBackup administrators can set up periodic or calendar-based schedules to perform automatic, unattended backups for clients across a network. The backups can be full or incremental: Full backups back up all indicated client files, while incremental backups back up only the files that have changed since the last backup.

During a backup or archive, the client sends backup data across the network to a NetBackup server. The NetBackup server manages the type of storage that is specified in the backup policy. When you restore virtual machine data using the Resiliency Platform console, you can choose from the available backup images.

About NetBackup Appliances

NetBackup appliances provide a simplified solution for NetBackup configuration and the daily management of your backup environment. The appliances are rack-mount servers that run on the Linux operating system. NetBackup Enterprise Server software is already installed and configured to work with the operating system, the disk storage units, and the robotic tape device.

For more information, refer to the NetBackup documentation on Services and Operations Readiness Tools (SORT).

About protecting assets with NetBackup using Resiliency Platform

Using the Resiliency Platform you can restore the virtual machine from NetBackup generated backup images to the recovery data center. To do this ensure that the required components are deployed and configured at both the production and recovery data centers. Refer to the below image for deployment information.

In the image, data center 1 is the production data center and data center 2 is recovery data center. Targeted Auto Image Replication, denoted as AIR in the below image, ensures that the backup images are available on NetBackup master server in the recovery data center. The image shows two Infrastructure Management Servers (IMS) although you can have only one IMS which discovers the vCenter and is also added as an additional server to NetBackup.

See “Using NetBackup - an overview” on page 83.

Figure 10-1 Deployment architecture for NetBackup master server

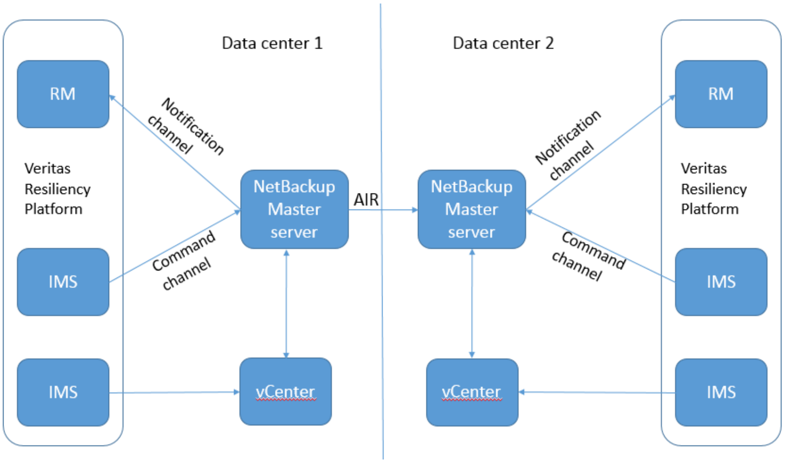

Using NetBackup - an overview

The following is a summary of the steps that are required to be performed to enable recovery of assets managed by NetBackup and where to go for more information on each step.

Table 10-1 Process overview

Step	More information
Add the NetBackup master servers to the Resiliency Manager at each data center. The NetBackup master server name is case sensitive.	See "Adding NetBackup master server" on page 85.
Add the Infrastructure Management Server (IMS) to the NetBackup master server each data center.	The IMS acts as an additional server to the NetBackup master server to discover the backup information of assets in the data center. See "Adding IMS to NetBackup master server as an additional server" on page 87.

Table 10-1 Process overview *(continued)*

Step	More information
Add the VMware virtualization servers to the IMS each data center. The vCenter Server should be configured to Veritas Resiliency Platform and NetBackup with the same name or with the same IP address. vCenter Server name is case sensitive.	See "Prerequisites for adding VMware virtualization servers" on page 125. See "Adding VMware virtualization servers " on page 134.
Customize and activate the Copy service objective for NetBackup recovery.	See "Viewing or customizing service objectives" on page 198.
Customize the network mapping between the data centers.	For more information refer to the VMware solutions guide.
Protect the assets by applying the service objective to the virtual machines.	For more information refer to the VMware solutions guide.
Restore virtual machines from backups	For more information refer to the VMware solutions guide.

Prerequisites for integrating NetBackup with Resiliency Platform

The NetBackup server administrator must perform some configuration tasks on the NetBackup master server in order to recover assets using Veritas Resiliency Platform. The following tasks must be completed before the Resiliency Platform administrator can activate backup service definitions and apply them to protect assets.

Establishing the connection to Resiliency Platform

The NetBackup administrator must help establish the connection between the NetBackup master server and Resiliency Platform. To establish the connection between Resiliency Platform and the NetBackup server, ensure the following:

- NetBackup version 8.0 or NetBackup Appliances 52XX version 3.0 is installed.
- Ensure that NetBackup Auto Image Replication is configured. It can be configured in multiple ways but only the Targeted mode is supported in Resiliency Platform. It must be configured between both the master servers across the two data centers.

Targeted Auto Image Replication ensures that the backup image is available on NetBackup master server in the recovery data center.

The correlation of NetBackup master servers at source and target data centers is through registered names. Ensure that while configuring Targeted Auto Image Replication the registered name of the NetBackup master server is used.

- The NetBackup master server is online.

See “Adding NetBackup master server” on page 85.

See “Adding IMS to NetBackup master server as an additional server” on page 87.

Adding NetBackup master server

Using the Resiliency Platform console you can add the NetBackup master server to the Resiliency Platform.

You can add the NetBackup master server in the following ways:

- Using the NetBackup webservice credentials.

  The Resiliency Platform uses these credentials to add itself as a websocket server in NetBackup. This channel is used for notification of backup activities.

  When you remove NetBackup master server using the Resiliency Platform console, then the websocket server for resiliency manager is deleted from NetBackup master server only if user has chosen to save the credentials.

  If you do not save the credentials, you need to manually delete the websocket server for resiliency manager using the NetBackup Administration Console.

  See “Adding the NetBackup master server using the webservice credentials” on page 86.

- Without the NetBackup webservice credentials.

  This is a two step task; first task is done using the Resiliency Platform console to add the NetBackup master server and a registration URL is generated on the console. The second step is done using the NetBackup Administration Console.

  Note that the URL contains a temporary token which is valid only for 30 minutes. If the operation is not completed in 30 minutes, then you need to edit the NetBackup master server to regenerate the URL.

  See “Adding the NetBackup master server without the webservice credentials” on page 86.

  NetBackup Appliances can be added using the URL registration method only.

Prerequisites

Ensure that the requirements for the NetBackup master server are met.

See “Prerequisites for integrating NetBackup with Resiliency Platform” on page 84.

Adding the NetBackup master server using the webservice credentials

1 Navigate

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

Copy Manager > **+ NetBackup Master**

2 In the **Configure NetBackup Master** panel, enter the following information and click **Next**.

- Select the Infrastructure Management Server (IMS) for command channel.
- Enter the name and the Webservice port number of the NetBackup master server.
- Enter the credentials to access the master server. You have the option to save the credentials.

3 Click **Finish** to exit the wizard.

Adding the NetBackup master server without the webservice credentials

1 Resiliency Platform administrator:

Navigate

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

Copy Manager > **+ NetBackup Master**

2 In the **Configure NetBackup Master** panel, enter the following information and click **Next**.

- Select the Infrastructure Management Server (IMS).
- Enter the name of the NetBackup master server.

User name and password are not required.

3 Click **Finish** to exit the wizard.

A URL and steps to complete the registration are provided on the wizard panel. Follow the steps using the NetBackup Aministration Console.

See “Adding Resiliency Manager to NetBackup using a URL” on page 87.

The NetBackup master server status is shown on the Resiliency Platform console as Pending.

4. After completing the registration using the URL, verify that the status of the NetBackup master server is shown as discovery pending on the Resiliency Platform console.
5. Add the Infrastructure Management Server (IMS) as an additional server to NetBackup Master.

   See "Adding IMS to NetBackup master server as an additional server" on page 87.

   After successful addition, the status of NetBackup master server is shown as Connected.

Adding Resiliency Manager to NetBackup using a URL

If NetBackup webservice credentials were not provided while adding the NetBackup master server to the resiliency domain using the Resiliency Platform console, then the NetBackup administrator needs to complete the following procedure.

While adding the NetBackup master server using the Resiliency Platform console a URL is provided. Use this URL to complete the procedure.

To add Resiliency Manager to NetBackup

1. In the NetBackup Administration Console, expand **Media and Device Management** > **Credentials**.
2. Right click **WebSocket Servers** and select **New WebSocket Server**.
3. In the **WebSocket Server** panel select URL to enter the URL along with the token that is displayed on the Resiliency Platform console.
4. Click **Validate**, review the certificate, and click **Yes**.
5. Click **Add Host**.

The NetBackup master server status is shown on the Resiliency Platform console as Connected.

See "Adding IMS to NetBackup master server as an additional server" on page 87.

Adding IMS to NetBackup master server as an additional server

After adding the NetBackup master server to the resiliency domain, you need to add the Infrastructure Management Server (IMS) to NetBackup as an additional server. The NetBackup administrator must complete the following procedure to add the IMS .

To add IMS to NetBackup master server as an additional server

1. Prerequisites

   The IMS should be accessible from the NetBackup master server.

2. In the NetBackup Administration Console, expand **NetBackup Management** > **Host Properties**.

3. Double click **Master Servers** to view the list of master servers.

4. Click on the required master server and select **Servers**.

5. On the **Additional Servers** tab, click **Add**.

6. In the **Enter Server Name** text box enter the fully qualified host name of the IMS. Click **Add**.

See "Adding NetBackup master server" on page 85.

See "Adding Resiliency Manager to NetBackup using a URL" on page 87.

Editing the NetBackup master server

Using the Resiliency Platform console, you can update the Infrastructure Management Server (IMS) and the credentials associated with the NetBackup master server.

Editing the NetBackup master server

1. Navigate

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

   Copy Manager > **NetBackup** tab

2. On the row of the NetBackup master server that you want to edit, select the vertical ellipsis > **Edit**.

3. You can edit the IMS selection and update the credentials.

   Select **Save Credentials** if you have updated them.

Removing the NetBackup master server

Using the Resiliency Platform console, you can remove the NetBackup master server from the Resiliency Manager. If you have not saved the credentials, then you need to remove the Resiliency Manager as a websocket server from NetBackup using the NetBackup Administration console.

After removing NetBackup master server, remove the Infrastructure Management Server (IMS) which is added as an additional server from the NetBackup Administration console.

See "Removing IMS from NetBackup master server" on page 89.

Note: When you remove the NetBackup master server, the Resiliency Platform cannot recover the assets protected by NetBackup.

To remove NetBackup master server using the Resiliency Manager console

1 Navigate

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

Copy Manager > **NetBackup** tab

2 On the row of the NetBackup master server that you want to remove, select the vertical ellipsis > **Remove**.

If you have not saved the credentials then manually remove the Resiliency Manager using the NetBackup Administration console.

To remove resiliency manager using the NetBackup Administration Console

1 In the NetBackup Administration Console, expand **Media and Device Management** > **Credentials**.

2 Right click **WebSocket Servers**, select the Resiliency Manager, and click **Delete**.

3 Review the selection and click **OK**.

Removing IMS from NetBackup master server

After removing NetBackup master server from the Resiliency Manager, it is required to remove the Infrastructure Management Server (IMS) from the additional server list in NetBackup console.

To remove IMS from NetBackup using the NetBackup Administration Console

1 In the NetBackup Administration Console, expand **NetBackup Management** > **Host Properties**.

2 Double click **Master Servers** to view the list of master servers.

3. Click on the required master server and select **Servers**.
4. On the **Additional Servers** tab, select the IMS, and click **Remove**.

Refreshing the NetBackup master server

Once the NetBackup master server is added successfully, the Resiliency Platform periodically (every 3 hours) discovers information about the assets protected by NetBackup master server. If there is a need to discover this information between the scheduled discoveries, perform the refresh operation.

Refreshing the NetBackup master server

1. Navigate

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

   Copy Manager > **NetBackup** tab

2. On the row of the NetBackup master server that you want to refresh, select the vertical ellipsis > **Refresh**.
3. Review the information and click **Next**.

Chapter 11

Managing Veritas InfoScale Operations Manager Server

This chapter includes the following topics:

About Veritas InfoScale Operations Manager

Veritas InfoScale Operations Manager gives you a single, centralized management console for the Veritas InfoScale products. You can use it to monitor, visualize, and manage storage and cluster resources, and generate reports about these components in the Management Server domain. Veritas InfoScale Operations Manager helps administrators centrally manage diverse data center environments.

Resiliency Platform support for InfoScale applications

A typical workflow of Veritas Resiliency Platform for InfoScale applications consists of a Veritas InfoScale Operation Manager server reporting to a Resiliency Manager. The InfoScale applications should be already configured in Veritas InfoScale Operations Management server. You can group the InfoScale applications into resiliency groups or VBSs to recover, monitor, visualize, and generate reports about these applications in the Resiliency Platform.

Note: Only the applications that are managed by InfoScale Availability (VCS) are supported in Veritas Resiliency Platform.

The following diagram depicts the general workflow of configuring the InfoScale applications using Resiliency Platform.

Figure 11-1 A typical workflow for recovering managed InfoScale applications

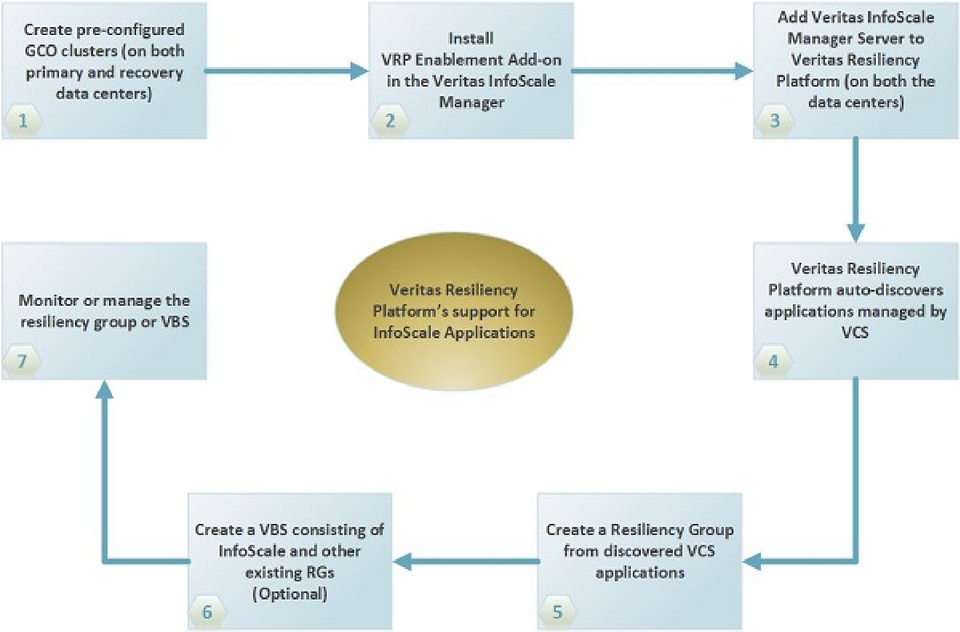
Create pre-configured GCO clusters (on both primary and recovery data centers)
1
Install VRP Enablement Add-on in the Veritas InfoScale Manager
2
Add Veritas InfoScale Manager Server to Veritas Resiliency Platform (on both the data centers)
3
Veritas Resiliency Platform auto-discovers applications managed by VCS
4
Create a Resiliency Group from discovered VCS applications
5
Create a VBS consisting of InfoScale and other existing RGs (Optional)
6
Monitor or manage the resiliency group or VBS
7
Veritas Resiliency Platform's support for InfoScale Applications

Managing InfoScale applications using Resiliency Platform

Veritas Resiliency Platform lets you manage the InfoScale applications that are already configured in Veritas InfoScale Operations Manager Management Server. The InfoScale applications are automatically discovered in the Resiliency Platform when the Veritas InfoScale Operations Manager server is added to the resiliency domain. They are listed on the **Assets** page under the **Unmanaged** tab. You can filter the InfoScale applications using the **InfoScale applications** check box under the **More Options** drop-down menu.

Veritas InfoScale Operations Manager users must download and install Veritas Resiliency Platform Enablement add-on to automatically discover the InfoScale applications. You can download the add-on from Veritas Services and Operations Readiness Tools (SORT). You cannot add or modify InfoScale applications through Resiliency Platform. They can be added or modified only by an administrator through Veritas InfoScale Operations Manager.

For information on configuring resiliency groups with InfoScale applications, refer to the *Veritas Resiliency Platform 2.0: Solutions for Applications* guide on SORT.

Support matrix for InfoScale applications

Resiliency Platform supports creating a resiliency group which contains any one of the following:

- A single VCS service group
- A group of service groups belonging to a group dependency

Following is the support matrix for InfoScale applications in Veritas Resiliency Platform2.1:

Table 11-1

Scenario	Supported service objective	Prerequisite	Firedrill support
Single local service group	Monitor	None	No

Table 11-1 *(continued)*

Scenario	Supported service objective	Prerequisite	Firedrill support
Single global service group	Monitor Recover	For recover service objective, at least one resource of replication type should be present	Yes For firedrill, the global service group name should be suffixed with _fd Example: Global service group name: global_grp Firedrill: global_grp_fd
Multiple local service groups	Monitor	None	No
Multiple global service groups	Monitor Recover	For recover service objective, at least one resource of replication type should be present	No

During disaster recovery operations, Veritas Resiliency Platform does not perform any operation on the local service groups.

Note: Support for multiple global service groups will be available only if you apply update 2.0.0.100 on Veritas Resiliency Platform2.1.

Prerequisites for managing Veritas InfoScale Operations Manager server using Resiliency Platform

The following prerequisites are required for managing InfoScale applications in Veritas Resiliency Platform:

- Veritas InfoScale Operations Manager (7.0 onwards) users should download and install the Veritas Resiliency Platform Enablement add-on on the Veritas InfoScale Operations Manager server and then on the intended managed host nodes.

  Note: You need to restart the Veritas InfoScale Operations Management server after installing the add-on.

- Veritas Resiliency Platform supports Veritas Cluster Servers (6.0 onwards).
- Global Cluster Option (GCO) must be enabled for performing DR operations.
- InfoScale applications must be configured along with replication in VCS for performing DR operations.
- Make sure the managed hosts in the GCO pair are not reporting to the IMS in the production data center and recovery data center simultaneously.

Adding a Veritas InfoScale Operations Manager server

Veritas Resiliency Platform includes Veritas InfoScale Operations Manager Management Server to discover and monitor InfoScale applications. You must download and install Veritas Resiliency Platform Enablement add-on add-on in the Veritas InfoScale Operations Manager console, and restart the Veritas InfoScale Operations Manager server after you complete the installation.

To add a Veritas InfoScale Operations Manager server in Resiliency Platform

1. Navigate

   Prerequisites

   Information needed for adding the Veritas InfoScale Operations Manager:

   - The fully qualified domain name (FQDN) or IP address of the server.
   - The username and password of the server and the domain name. These three are optional information that you need to use only if the server is directly accessible.
   - If the server is not directly accessible, you can still initiate the process of adding a Veritas InfoScale Operations Manager server by entering only the data center, friendly name, and FQDN/IP address of the server. In this case,

you get a registration URL which you need to use after logging in to the Veritas InfoScale Operations Manager server.

2. Navigate

   ⚙ **Settings** (menu bar)

   Under **Infrastructure Settings**, click **Infrastructure**

   You can also access this page from the **Quick Actions** menu > **Manage Asset Infrastructure**.

   Click **Application Cluster**.

   Under **InfoScale**, Click **Add Veritas Infoscale Operations Manager**.

3. In **Add Veritas InfoScale Operations Manager** window, enter the information for the Veritas InfoScale Operations Manager:

   - You can select from a list of existing data centers or add a new data center. To specify a new data center, select **New** in the **Data Center** field, then specify the location and name. When entering the location, enter a form of location identifier, such as city, and the location list will populate with potential matches for you to select.
   - Enter a friendly name for the server.
   - Enter the FQDN or IP address of the server.
   - Enter the user name, password, and then domain name. These are optional information that you need to enter only if the Veritas Infoscale Operations Manager is directly accessible. If you provide this information, the Veritas Infoscale Operations Manager is immediately added to the domain.
     If you do not provide the username, password, and domain name of the Veritas Infoscale Operations Manager, a registration URL is displayed on the screen. This URL is valid only for 30 minutes. if the URL expires, you need to regenerate the registration URL to complete the process.
     See "Regenerating registration URL for a Veritas InfoScale Operations Manager server" on page 97.

4. If you have applied update 2.0.0.100 on Veritas Resiliency Platform 2.1, you can select the checkbox to allow script execution on the hosts that are associated with the Veritas InfoScale Operations Manager server.

5. Log into the Veritas InfoScale Operations Manager console. Go to **Settings**> **Deployment**, and click **Add VIOM to Resiliency Manager**. In the text area, paste the URL that you had copied from Resiliency Manager console and click **Finish**.

6. Verify that the Veritas InfoScale Operations Manager server is successfully added. Once the server is successfully added, the applications that are already added in the Veritas InfoScale Operations Manager are shown in the **Unmanaged** tab under **Assets**.

   Note: You can add or manage the asset infrastructure related to the InfoScale applications only through Veritas InfoScale Operations Manager console. For managing applications in Veritas InfoScale Operations Manager, you need to log in using administrator credentials.

   After you have successfully added the Veritas InfoScale Operations Manager server, click the vertical ellipsis and select **Manage Asset Infrastructure** to launch Veritas InfoScale Operations Manager console.

Regenerating registration URL for a Veritas InfoScale Operations Manager server

When you add a Veritas InfoScale Operations Manager server to a data center using the registration URL, the URL is valid only for 30 minutes. If the URL expires, you can regenerate the registration URL and then complete the operation.

To regenerate registration URL for a Veritas InfoScale Operations Manager server

1. Navigate

   ⚙ **Settings** (menu bar)

   Under **Infrastructure Settings**, click **Infrastructure**

   You can also access this page from the **Quick Actions** menu > **Manage Asset Infrastructure**.

   Click **Application Cluster**.

   Under **InfoScale**, locate the Veritas InfoScale Operations Manager server, right click and select **Regenerate Registration URL**.

2. In the **Regenerate Registration URL** window, click **Submit**.

When you add a Veritas InfoScale Operations Manager server to a data center using the registration URL, the URL is valid only for 30 minutes. If the URL expires, you can regenerate the registration URL and then complete the operation.

Deleting a Veritas InfoScale Operations Manager server

In the web console, you can delete a Veritas InfoScale Operations Manager server from a resiliency domain by performing the following steps:

To delete a Veritas InfoScale Operations Manager server

1 Navigate

⚙ **Settings** (menu bar)

Under **Infrastructure Settings**, click **Infrastructure**

You can also access this page from the **Quick Actions** menu > **Manage Asset Infrastructure**.

Click **Application Cluster**.

Under **InfoScale**, locate the Veritas InfoScale Operations Manager server, right click and select **Delete**.

2 In the **Remove Veritas InfoScale Operations Manager server** window, click **Submit**.

3 Verify that the Veritas InfoScale Operations Manager server is removed.

Note: Additional cleanup is optional. If resiliency groups were created for the assets managed by the Veritas InfoScale Operations Manager server and VBS, and they are no longer needed, you can remove them.

Editing the details of a Veritas InfoScale Operations Manager server

In the web console, you can edit the friendly name associated with a Veritas InfoScale Operations Manager server that has been added to a resiliency domain or change the data center name.

To edit the details of a Veritas InfoScale Operations Manager server

1 Navigate

⚙ **Settings** (menu bar)

Under **Infrastructure Settings**, click **Infrastructure**

You can also access this page from the **Quick Actions** menu > **Manage Asset Infrastructure**.

Click **Application Cluster**.

Under **InfoScale**, locate the Veritas InfoScale Operations Manager server, right click and select **Edit**.

2 In the **Edit Veritas InfoScale Operations Manager server** window, make the required changed and click **Submit**.

Reconnecting the Veritas InfoScale Operations Manager to a Resiliency Manager

You can use the web console to reconnect a Veritas InfoScale Operations Manager to a Resiliency Manager. The operation disconnects any existing connection and the Veritas InfoScale Operations Manager then reinitiates the connection. This operation can be useful when troubleshooting or repairing a connection between an Veritas InfoScale Operations Manager and a Resiliency Manager.

To reconnect a Veritas InfoScale Operations Manager to a Resiliency Manager

1 Navigate

⚙ **Settings** (menu bar)

Under **Infrastructure Settings**, click **Infrastructure**

You can also access this page from the **Quick Actions** menu > **Manage Asset Infrastructure**.

Click **Application Cluster**.

Under **InfoScale**, locate the Veritas InfoScale Operations Manager server, right click and select **Reconnect**.

2 In the **Reconnect Veritas Operations Manager Server** window, click **Submit**.

Chapter 12

Managing Resiliency Platform Data Mover gateway pairing

This chapter includes the following topics:

Viewing Veritas Replication Gateways

After deployment of Veritas Replication Gateways, you can view information in the console about the gateway name, health, IP address and associated IMS.

To view Replication Gateways

1 Navigate

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

You can also access this page from the Quick Actions menu.

2 Click **Data Mover**

The **VRP Data Mover** tab lists the gateway information.

Healthy state indicates that all the required daemons are running on the gateway.

See "About Veritas Replication Gateway pairs" on page 101.

About Veritas Replication Gateway pairs

To use the Resiliency Platform Data Mover feature, you must deploy at least one Replication Gateway on both the source and target data center. The source and target Replication Gateways must be paired before replication is enabled.

Starting from version 2.1, Resiliency Platform supports the asymmetric pairing of Replication Gateways. This feature facilitates deployment of only the required number of gateways on each side, based on data transfer rate and technology specific limits.

One Gateway on production site can be paired with multiple Gateways on recovery site and vice versa. One gateway can be paired with up to 16 gateways on the peer site.

For each Gateway pair, you can choose to apply an encryption scheme to the data replication.

When you protect virtual machines using Resiliency Platform Data Mover, you select the Gateway pair to use for the replication. Once the DR configuration is complete, the Replication Gateway at the source data center starts replicating the data to the paired Gateway at the target data center.

If the recovery data center is Amazon Web Services (AWS) cloud, then ensure that the virtualization server storage (datastore and volumes) on which the virtual machines reside are available to the Replication Gateway on the production data center.

How Resiliency Platform Data Mover supports encryption for data replication

The Veritas Resiliency Platform Data Mover Replication Gateway supports encryption using OpenSSL for data transfer. When creating or modifying a Replication Gateway pair, you can choose whether to apply an encryption scheme to the data replication.

Creating a Veritas Replication Gateway pair

To protect virtual machines using Resiliency Platform Data Mover, you must create Replication Gateway pairs.

To create a Replication Gateway pair

1 Prerequisites

A Replication Gateway must be deployed and configured in each data center.

2 Navigate

Disaster Recovery Settings (navigation pane)

Replication Appliance tab > **+ Replication Gateway Pair**

3 Specify the information in the wizard:

- Select the gateways to be paired. You can filter each list by data center.
- Enter the IP address to be used by the gateway. You can specify different IP addresses for communication between gateways at the source and target data centers and for communication between an ESXi host and the gateway.
- Optionally edit the default name of the gateway pair.
- Optionally change the data encryption scheme selection.

When you submit, a message confirms that the pairing configuration is initiated. You can view the progress of the operation on the **Activities** pane.

Once the operation is complete, the gateway pair is listed on the **Replication Appliance** tab. When connection between gateways is established, the gateway pair state will be **Connected**.

See “Removing a Veritas Replication Gateway pair” on page 103.

See “About Veritas Replication Gateway pairs” on page 101.

Removing a Veritas Replication Gateway pair

Using the Resiliency Platform console, you can remove an existing Veritas Replication Gateway pair. Removing a gateway pair does not remove the Replication Gateways themselves, only the pairing configuration.

To remove a Replication Gateway pair

1 Prerequisites

If a resiliency group is configured for disaster recovery using a Replication Gateway pair, then you need to unconfigure DR for the resiliency group before you delete the Replication Gateway pair. Ensure that the replication set is removed during the unconfigure DR operation.

2 Navigate

Disaster Recovery Settings (navigation pane) > **Replication Appliance** tab

3 Select the vertical ellipses next to the pair name and select **Delete**.

See “About Veritas Replication Gateway pairs” on page 101.

Viewing Veritas Replication Gateway pairs

Using the Veritas Resiliency Platform console, you can view information about the Replication Gateway pairs. The information includes the source and target data centers, the connection state, and whether a data encryption scheme is applied. The initial state is **Disconnected**, until all the connections between the gateways are established.

To view Replication Gateway pairs

◆ Navigate

Disaster Recovery Settings (navigation pane) > **Replication Appliance** tab

See “About Veritas Replication Gateway pairs” on page 101.

Modifying encryption for a Veritas Replication Gateway pair

When you create a Replication Gateway pair, you can specify an encryption scheme for replication. You can modify this option after a gateway pair is created. When you change an encryption scheme, the Replication Gateway transceiver component is restarted. When the transceiver restarts, it resumes sending or receiving update sets from where it left off, hence full synchronization is not required. The gateway pair may be in a disconnected state temporarily during the process of restarting.

AES128-GCM-SHA256 and AES256-GCM-SHA384 are the available encryption schemes. The default scheme is None.

To modify a Replication Gateway pair

1. Navigate

   Disaster Recovery Settings (navigation pane) > **Replication Appliance** tab

2. Select the vertical ellipses next to the pair name and select **Edit**.
3. In the wizard, change the encryption scheme selection and submit.

See “About Veritas Replication Gateway pairs” on page 101.

Section 4

Adding the asset infrastructure

Chapter 13

Managing asset infrastructure

This chapter includes the following topics:

Adding the asset infrastructure

Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The Infrastructure Management Server (IMS) then discovers the asset information for monitoring and operations in the console.

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.

The asset infrastructure can include the following technology categories, depending on the use case and your environment:

- **Managed Hosts** - For application discovery on physical and virtual machines
- **Virtualization** - For VMware servers and Hyper-V servers
- **Storage** - For storage arrays and replication appliances used for replication

To add the asset infrastructure

1 Prerequisites

To verify supported assets, refer to the *Hardware and Software Compatibility List (HSCL)*.

The data center must contain at least one IMS for discovery of the data center assets.

2 Navigate

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

You can also access this page from the **Quick Actions** menu > **Manage Asset Infrastructure**.

3 Expand the data center and locate the technology category for the asset.

See “Managing host assets” on page 107.

See “Managing Hyper-V assets ” on page 108.

See “Managing VMware virtualization servers” on page 108.

See “Managing enclosure assets” on page 109.

See “Adding RecoverPoint appliance for replication” on page 166.

Managing host assets

The asset infrastructure that you must add to Veritas Resiliency Platform for discovery and monitoring by an Infrastructure Management Server (IMS) can include assets that you add as hosts. The following topics describe the types of assets you add as hosts, the prerequisites, and how to add and remove host assets.

See “About adding host assets” on page 110.

See “Prerequisites for adding hosts ” on page 112.

See “Adding application hosts ” on page 114.

See “Preparing host for replication” on page 115.

See “Removing hosts ” on page 117.

See “Refreshing host discovery information” on page 118.

See “Uninstalling the host package from a Linux host” on page 119.

See “Uninstalling the host package from a Windows host” on page 119.

See “Refreshing host discovery information” on page 118.

Managing Hyper-V assets

You can add Hyper-V servers to Veritas Resiliency Platform for discovery of Hyper-V virtual machines by an Infrastructure Management Server (IMS). Hyper-V servers are added as virtualization servers.

See “About Microsoft Hyper-V virtualization discovery” on page 121.

See “Prerequisites for Microsoft Hyper-V virtualization discovery” on page 122.

See “Adding Hyper-V virtualization servers” on page 122.

See “Removing Hyper-V virtualization servers ” on page 123.

See “Refreshing Hyper-V virtualization servers” on page 124.

Managing VMware virtualization servers

You can add VMware vCenter servers to Resiliency Platform for discovery by an Infrastructure Management Server (IMS).

The VMware discovery provides the following information:

- Information on vCenter servers
- Information on the ESX servers managed by the vCenter servers
- Information on the virtual machines that are configured on the ESX servers

See “Prerequisites for adding VMware virtualization servers” on page 125.

See “About adding a host for discovery of VMware servers” on page 133.

See “Adding a host for discovery of VMware servers ” on page 133.

See “Adding VMware virtualization servers ” on page 134.

See “Editing a VMware virtualization discovery configuration” on page 136.

See “Viewing the details of a VMware virtualization discovery configuration” on page 138.

See “Removing a VMware vCenter Server discovery configuration ” on page 138.

See “Refreshing VMware vCenter Server discovery information” on page 139.

Managing enclosure assets

You can add storage enclosures (arrays) to Resiliency Platform for discovery of storage information to monitor array-based replication. This does not apply for environments that are using other types of replication.

See “About the discovery host for enclosures” on page 141.

See “Configuration prerequisites for adding storage enclosures to an IMS” on page 145.

See “Adding storage enclosures ” on page 160.

See “Adding RecoverPoint appliance for replication” on page 166.

See “Editing the discovery configuration for an enclosure” on page 167.

See “Removing the discovery configuration for an enclosure” on page 168.

Chapter 14

Preparing and maintaining host assets

This chapter includes the following topics:

About adding host assets

You add several types of assets as hosts to Veritas Resiliency Platform for discovery and monitoring by an Infrastructure Management Server (IMS). Host assets that you add can include physical systems, virtual machines, and discovery hosts, depending on the use case, as described in the table.

Note: You must add a host for discovery only once.

Table 14-1 Use cases for adding host assets

Use case	Details
Application discovery and management	For discovery of supported applications on either physical systems or virtual machines, you must add the physical system or virtual machine as a host. **Note:** For the use case of discovering and managing virtual machines rather than applications, you do not need to add the virtual machines as hosts. For discovery of a custom application, after you add the hosts, you must also add the application on the Assets page. More information is available on adding custom applications. See *Veritas Resiliency Platform Solutions for Applications*.
VMware vCenter Server discovery (optional)	You can add a host to be used by the IMS for discovery of a VMware vCenter Server. See "About adding a host for discovery of VMware servers" on page 133.
Hardware replication	For storage array-based replication, you may need to install array-specific software on a host and add the host as a discovery host. See "About the discovery host for enclosures" on page 141. More information is available on requirements for adding enclosures for array-based replication. See "Managing enclosure assets" on page 109.

When you add hosts to Resiliency Platform, the IMS installs the host package (VRTSsfmh) on the host. On Linux hosts, the VRTSsfmh package is installed in the /opt directory. On Windows hosts, the VRTSsfmh package is installed in the system drive.

The IMS also installs several add-on packages on the host for use by the IMS discovery:

- Veritas Resiliency Platform Enablement add-on
- Applications Enablement add-on

Before you add hosts, ensure that all prerequisites are met.

See "Prerequisites for adding hosts " on page 112.

Prerequisites for adding hosts

Before you add hosts to Veritas Resiliency Platform for discovery and monitoring by an Infrastructure Management Server (IMS), ensure that the following prerequisites are met. Prerequisites include general prerequisites for all hosts and additional prerequisites for Linux or Windows systems.

General prerequisites for adding host assets:

- Ensure that the IMS can communicate with the host.
- Ensure that the time difference between the system clocks on the IMS and host is no more than 90 minutes. The managed hosts must report synchronized universal time clock time (UC/UTC).
- If a CSV file is used to add hosts, ensure that it uses the correct syntax.

Additional prerequisites for Linux systems:

- In order to install the host package while adding the Linux host, ensure that the PasswordAuthentication field is set to **yes** in the `/etc/ssh/sshd_config file` on the host.

Additional prerequisites for Windows systems:

- You must have at least one Windows Install host already added to the IMS, where you want to add the Windows host. If you do not have any Windows Install host associated with the IMS, you first need to add a Windows Install host, and then you can add any number of Windows hosts using the Windows Install host.
  See "Installing the host package on a Windows host" on page 113.
- If you install the host package using the web console, you should be a domain user having administrative privileges on the host. If you install the host package manually, then you need to be a local user having administrative privileges on the host.
- The Windows Management Instrumentation (WMI) service must be running.

More information is available about the add host operation.

See "About adding host assets" on page 110.

Adding a Windows Install host

Before you add a Windows host to any Infrastructure Management Server (IMS) for applications discovery or as a discovery host, you need to have at least one Windows Install host associated with that IMS. This Windows Install host acts as a control host that enables the process of adding a Widows host to the IMS.

To add a Windows Install host

1 Prerequisites

Ensure that the managed host package (VRTSsfmh) is installed on the host.

See "Installing the host package on a Windows host" on page 113.

2 Navigate

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

You can also access this page from the **Quick Actions** menu.

3 Under the data center, locate the IMS, and do the following:

⋮ Select the vertical ellipsis for the IMS > **Manage Windows Install Hosts**.

4 Under **Enter Host Details**, enter the host name, username, and password for the host to be added, and click **Submit.**

Installing the host package on a Windows host

Before you can use the wizard in the web console to add Windows hosts to an Infrastructure Management Server, you must first manually install the `VRTSsfmh` host package on at least one Windows host.

Note: By default, the `VRTSsfmh` package is installed in the system drive. You cannot specify a different location to install the package.

To install the host package on a Windows host

1 Log on to the target host as a user with administrator privileges.

2 Make sure that the value for environment variable PATHEXT on the target host includes the extensions .exe, .bat, and .vbs.

3 Download the host installation files bundle, and unzip it.

See "Downloading the Resiliency Platform virtual appliances" on page 25.

4 From the directory to which you unzipped the installation files bundle, open an elevated command prompt and run `VRTSsfmh_7.0.0.0_Windows_arch_x64.msi`.

5 On the welcome screen of the Installation Wizard, click **Next**.

6 On the **Ready to Install the Program** screen, click **Install** to start the installation.

7 Click **Finish** to exit the Installation Wizard.

See “Managing host assets” on page 107.

Adding application hosts

To enable the discovery of an application in Resiliency Platform, you need to add the application host to an Infrastructure Management Server (IMS).

To add an application host

1 Prerequisites

Ensure that you understand the use cases and prerequisites for adding hosts to an IMS.

See “About adding host assets” on page 110.

See “Prerequisites for adding hosts ” on page 112.

Ensure that the following Linux packages are installed: glibc (32-bit), libstdc++ (32-bit), and NetworkManager.

2 Navigate

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

You can also access this page from the **Quick Actions** menu.

3 Click **Managed Host**.

4 Under **Application**, click **+ Add application host**

5. In the wizard, select the IMS to which you want to add this host.

   Select the installation option that corresponds to the platform of the hosts. The appropriate host package is automatically installed on the hosts by the IMS if you continue with the Add Host operation.

   If you want to add the host on a Windows platform, you need to select a Control Host, that helps in the installation of the Windows host. If you do not have any Control Host associated with the IMS, then you first need to add a Windows Install host to the IMS, and then you can select the same host as a Control Host in the Add application host wizard for adding other Windows hosts.

   See "Adding a Windows Install host" on page 112.

6. Type the information in the table row. You can use the following icons for entering details of multiple hosts simultaneously or for deleting a particular row from the table:

Icon	Task
+	To add a blank table row.
⧉	To copy the details of the selected table row. You can edit the details of the newly added row.
⇥	To import the information from a CSV file. Click **Browse** to select the text file and then click **Load host details**.
✕	To delete a row.

7. Once you are done with entering the data for all the hosts, click **Submit** and verify that the hosts have been added successfully.

8. Click **Finish** to close the wizard.

You can go to the **Activities** pane to view the progress of the add host operation.

See "Managing host assets" on page 107.

Preparing host for replication

To enable the replication in Resiliency Platform using a Data Mover, you need to add the virtual machine and prepare it for replication.

To prepare a host for replication

1 Prerequisites

Ensure that you understand the use cases and prerequisites for adding hosts to an IMS.

See “About adding host assets” on page 110.

See “Prerequisites for adding hosts ” on page 112.

Ensure that you select the same IMS while adding a virtual machine as well as while adding a vCenter server or a Hyper-V server which hosts that virtual machine.

2 Navigate

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

You can also access this page from the **Quick Actions** menu.

3 Click **Data Mover**.

4 Under **VRP Data Mover**, click **+ Prepare host for replication**

5 In the wizard, select the IMS to which you want to add this host.

Select the installation option that corresponds to the platform of the hosts. The appropriate host package is automatically installed on the hosts by the IMS if you continue with the Add Host operation.

If you want to add the host on a Windows platform, you need to select a Control Host, that helps in the installation of the Windows host. If you do not have any Control Host associated with the IMS, then you first need to add a Windows Install host to the IMS, and then you can select the same host as a Control Host in the Add application host wizard for adding other Windows hosts.

See “Adding a Windows Install host” on page 112.

6 Type the information in the table row. You can use the following icons for entering details of multiple hosts simultaneously or for deleting a particular row from the table:

Icon	Task
+	To add a blank table row.
⧉	To copy the details of the selected table row. You can edit the details of the newly added row.
⇥	To import the information from a CSV file. Click **Browse** to select the text file and then click **Load host details**.
✕	To delete a row.

7 Once you are done with entering the data for all the hosts, click **Submit** and verify that the hosts have been added successfully.

See “Managing host assets” on page 107.

Removing hosts

You can remove one or more hosts that were added to Veritas Resiliency Platform for discovery and monitoring by an Infrastructure Management Server (IMS).

If the hosts contain assets that were added to a Resiliency Platform resiliency group, after you remove the hosts, the assets are no longer shown as part of the resiliency group in the console. However, removing a resiliency group does not remove related hosts from the IMS. Removing hosts and removing resiliency groups are separate operations and can be performed in either sequence.

For more information about resiliency groups, see the Solutions guides.

When you perform the remove host operation on any host, it first uninstalls all the add-ons that were installed on that host, and then removes the host from the IMS. The host is not removed in case the uninstallation of any of the add-ons fails. If the same host is being used in in any other context in add-ons such as discovery host, the add-ons that are required for that particular context are not removed from the host. If a host is used in multiple contexts, then it is removed only from the context from where you perform the remove host operation.

To remove hosts

1 Navigate

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

2 Go to the section from where you want to remove the host. For example, the **Managed Host** section or the **Storage** section.

3 On the host listing page, right-click the host and select **Remove**.

4 Confirm that you want to remove the host.

5 You can check the details of the remove host workflow in the **Recent Activities** pane.

Removing a host does not uninstall the host package (VRTSsfmh) from the host. More information is available on uninstalling the host package.

See "Uninstalling the host package from a Linux host" on page 119.

See "Uninstalling the host package from a Windows host" on page 119.

See "Managing host assets" on page 107.

Refreshing host discovery information

You can submit a refresh request to update the information displayed for the hosts that have been added to Veritas Resiliency Platform. Once the refresh operation is complete, the Assets page in the console is also updated.

To refresh a host discovery

1 Navigate

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

2 Under the data center, locate the IMS and click **Manage Asset Infrastructure**.

3 On the IMS **Settings** page, click **Host**.

4 Right-click the host and select **Refresh**.

5 Click **OK**.

The refresh operation is asynchronous. The wizard displays that the operation has triggered the refresh, but the discovery operation is in progress in the background. The Discovery State column shows a status of Refreshing. When it is complete, you can view the status change reflected in the Discovery State column.

See “Managing host assets” on page 107.

Uninstalling the host package from a Linux host

You can use an operating system command to remove the VRTSsfmh package from a Linux host. Before you uninstall the host package, remove the host from the Infrastructure Management Server (IMS).

See “Removing hosts ” on page 117.

To uninstall the host package from a Linux host

1 Open an operating system console.

2 On the managed host where you plan to uninstall the host package, log on as root.

3 At the command prompt, enter the following command to uninstall the package:

```
rpm -e VRTSsfmh
```

Uninstalling the host package from a Windows host

You can use an operating system command to remove the VRTSsfmh package from a Windows host. Before you uninstall the host package, remove the host from the Infrastructure Management Server (IMS).

See “Removing hosts ” on page 117.

To uninstall the host package from a Windows host

1 Log in to the target host as a user with administrator privileges.

2 Go to the Windows **Control Panel**, and click **Programs and Features**.

3 From the list of installed programs, select **Veritas InfoScale Operations Manager (Host Component)**.

4 Do one of the following:

- Select **Uninstall** at the top of the list.
- Right click and select **Uninstall**. Click **Yes** to confirm.

Chapter 15

Managing Hyper-V virtualization server assets

This chapter includes the following topics:

About Microsoft Hyper-V virtualization discovery

Hyper-V is a hypervisor-based virtualization technology from Microsoft. The Infrastructure Management Server (IMS) can discover Hyper-V host and virtual machine-related information if the Hyper-V role is enabled on the host. The Hyper-V WMI API and Windows PowerShell commandlets are used for the discovery.

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

- Virtual machine discovery: Discovery of the Hyper-V virtual machines and its correlation with the Hyper-V server.

  When you add the Hyper-V server to the IMS, IMS discovers all virtual machines including the virtual machines without the guest operating system installed.
- Exported storage discovery: Discovery of storage that is provisioned to the guests and its correlation with the virtual machine and Hyper-V server.

IMS discovers the storage provisioned to the guests from the host's local storage, or storage area network (SAN). The Hyper-V guest, when added to the IMS domain, provides storage mapping discovery.

See "Managing Hyper-V assets " on page 108.

Prerequisites for Microsoft Hyper-V virtualization discovery

You can add Microsoft Hyper-V servers to Veritas Resiliency Platform for virtualization discovery by an Infrastructure Management Server (IMS).

For information on supported operating system versions for the Hyper-V Server, refer to the *Hardware and Software Compatibility List (HSCL)*.

Table 15-1 Requirements for Microsoft Hyper-V virtualization discovery

Type of discovery	Requirements
Virtual machine discovery	■ The `VRTSsfmh` package must be installed on the Hyper-V Server (parent partition). This is done automatically by the IMS when you add the Hyper-V server to Resiliency Platform. ■ The Hyper-V role must be enabled. ■ The Windows Management Instrumentation (WMI) service must be running on the Hyper-V Server.
Exported storage discovery	■ The Windows Management Instrumentation (WMI) service must be running on the guest.

See "Managing Hyper-V assets " on page 108.

Adding Hyper-V virtualization servers

You can add Microsoft Hyper-V servers to Veritas Resiliency Platform for virtualization discovery by an Infrastructure Management Server (IMS).

To add Hyper-V virtualization servers

1. Prerequisites:

   See “Prerequisites for Microsoft Hyper-V virtualization discovery” on page 122.

2. Navigate

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

   Expand the data center > **Virtualization** > **Hyper-V** tab

   Launch the **+ Hyper-V Server** wizard

   Tip: You can also access **Manage Asset Infrastructure** from the **Quick Actions** menu.

3. In the wizard, select the IMS, specify the required information about the Hyper-V server, and click **Submit**.

4. The Hyper-V server that has been added is listed on the **Hyper-V** tab. Discovery of the Hyper-V virtual machines occurs in the background. You can view the progress on the **Activities** page.

If changes are made after the IMS discovery is complete, you need to refresh the discovery of the Hyper-V server.

See “Managing Hyper-V assets ” on page 108.

Removing Hyper-V virtualization servers

You can remove a Hyper-V virtualization server that has been added to Resiliency Platform.

To remove a Hyper-V virtualization server

1. Navigate

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

   Expand the data center > **Virtualization** > **Hyper-V** tab

   Tip: You can also access **Manage Asset Infrastructure** from the **Quick Actions** menu.

2. ⋮ On the row for the Hyper-V server that you want to remove, select the vertical ellipsis > **Remove**.

Refreshing Hyper-V virtualization servers

You can refresh the IMS discovery for a Hyper-V virtualization server that has been added to Resiliency Platform.

To refresh Hyper-V virtualization servers

1. Navigate

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

   Expand the data center > **Virtualization** > **Hyper-V** tab

   Tip: You can also access **Manage Asset Infrastructure** from the **Quick Actions** menu.

2. ⋮ On the row for the Hyper-V server that you want to remove, select the vertical ellipsis > **Refresh**.

See “Managing Hyper-V assets ” on page 108.

Configuring asset infrastructure for replication for Hyper-V virtual machines

You can configure Hyper-V virtual machines for replication using either Hyper-V Replica or supported hardware array-based replication.

Ensure that you add the Hyper-V host to Resiliency Platform at both the production and target data center.

If you are using array-based replication, you must add the arrays used for replication to Resiliency Platform.

See “Managing enclosure assets” on page 109.

For an overview of the steps required to set up replication with Hyper-V Replica, see the *Veritas Resiliency Platform: Solutions for Microsoft Hyper-V* guide.

Chapter 16

Managing VMware virtualization server assets

This chapter includes the following topics:

Prerequisites for adding VMware virtualization servers

Ensure that the following requirements are met to add the VMware vCenter or ESX servers to Resiliency Platform for discovery by an Infrastructure Management Server (IMS):

- Ensure that the IMS can ping the vCenter servers or the ESX servers from which it can discover the information on VMware Infrastructure.

  Optionally, you can add a separate host to act as the discovery host for the vCenter Server.

  See "About adding a host for discovery of VMware servers" on page 133.
- Ensure that you have configured near real-time discovery of VMware events.

  See "About near real-time discovery of VMware events" on page 129.
- Ensure that the vCenter Server user account that is used to add the servers to Resiliency Platform has the following privileges assigned:

  Note: See "VMware vCenter Server privileges required for implementing Resiliency Platform Data Mover" on page 127.

  - System.Anonymous
  - System.View
  - System.Read
  - Datastore.AllocateSpace
  - Datastore.Rename
  - Host.Config.Storage
  - Host.Config.Settings
  - Host.Config.Network
  - VirtualMachine.Config.AddNewDisk
  - VirtualMachine.Config.RemoveDisk
  - VirtualMachine.Config.DiskExtend
  - VirtualMachine.Interact.PowerOff
  - VirtualMachine.Interact.PowerOn
  - VirtualMachine.Inventory.Register
  - VirtualMachine.Inventory.Unregister
  - VirtualMachine.Provisioning.Customize

There are additional requirements for virtual machines if added to the IMS, depending on the use case.

See "Prerequisites for adding hosts " on page 112.

VMware vCenter Server privileges required for implementing Resiliency Platform Data Mover

To implement Veritas Resiliency Platform Data Mover with VMware vCenter Server, the following privileges are required on the VMware vCenter Server account that is used to add the vCenter Server to Resiliency Platform:

Table 16-1 VMware vCenter Server privileges required for Resiliency Platform Data Mover

Category	Privilege
System	System.View System.Anonymous System.Read
Host	Host.Config.Patch Host.Cim.CimInteraction Host.Config.Maintenance Host.Config.Storage Host.Config.Settings Host.Config.Network
Virtual machine configuration	VirtualMachine.Config.AddExistingDisk VirtualMachine.Config.AddNewDisk VirtualMachine.Config.RemoveDisk VirtualMachine.Config.Rename VirtualMachine.Config.CPUCount VirtualMachine.Config.Memory VirtualMachine.Config.EditDevice VirtualMachine.Config.DiskExtend

Table 16-1 VMware vCenter Server privileges required for Resiliency Platform Data Mover *(continued)*

Category	Privilege
Virtual machine interaction and inventory	VirtualMachine.Interact.PowerOn VirtualMachine.Interact.PowerOff VirtualMachine.Interact.ToolsInstall VirtualMachine.Inventory.Create VirtualMachine.Inventory.Register VirtualMachine.Inventory.Unregister VirtualMachine.Inventory.Remove
Virtual machine provisioning	VirtualMachine.Provisioning.CloneVirtualMachine VirtualMachine.Provisioning.Customize
Network privileges	Network.Assign Network.Configure
Storage policy	VM storage policies.Update VM storage policies.View Profile.create
Folder	Folder.delete
Datastore	Datastore.AllocateSpace Datastore.FileManagement Datastore.Browse
Alarm	Alarm.Create Alarm.Delete Alarm.Modify
vCenter	Global.Settings Global.Diagnostics
Snapshots	VirtualMachine.State.CreateSnapshot VirtualMachine.State.RemoveSnapshot
Virtual appliance deployment	VApp.Import

About near real-time discovery of VMware events

The Infrastructure Management Server (IMS) uses VMware events to discover in near real-time a change in the state of a virtual machine (for example, virtual machine powered on) and changes occurring at the vCenter Server infrastructure level (for example, virtual machine created).

The near real-time discovery of VMware infrastructure enables the partial discovery of ESX servers managed under a vCenter Server. This discovery is triggered by the event notification from the VMware vCenter Server to the IMS using SNMP traps. For example, if an SNMP trap is received for a virtual machine (VM1) hosted on ESX1, the IMS runs the discovery cycle only for ESX1. Other ESX servers under that vCenter Server are not re-discovered.

The IMS component of near real-time discovery is `xtrapd`. After you configure a vCenter Server to send the SNMP traps to the IMS, you add the vCenter Server to the IMS. The `xtrapd` daemon now detects the SNMP traps that are sent from the specified vCenter Server. The Resiliency Platform database and console are updated with the latest state of the virtual machine or infrastructure changes.

Note: SNMP version 1 (SNMPv1) and version 2 (SNMPv2) are supported.

For details on supported events, see the following table.

Table 16-2 Supported events for near-real time discovery

Discovered state	Event as shown in VMware vCenter Server
Virtual machine powered on	VM powered on
Virtual machine powered off	VM powered off
Virtual machine Distributed Resource Scheduler (DRS) powered on	DRS VM powered on
Virtual machine suspended	VM suspended
Virtual machine created	VM created
Virtual machine migrated Hot migration: A powered-on virtual machine is migrated from one ESX server to another ESX server.	VM migrated

Table 16-2 Supported events for near-real time discovery *(continued)*

Discovered state	Event as shown in VMware vCenter Server
Virtual machine relocated from one ESX server to another Cold migration: A powered-off virtual machine is migrated from one ESX server to another ESX server.	VM relocating
Virtual machine renamed	VM renamed
Virtual machine migrated to another host by VMware DRS (Distributed Resource Scheduler)	DRS VM migrated

Setting up near real-time discovery of VMware events

To set up the near real-time discovery of VMware events, complete the following steps.

Table 16-3 Setting up near real-time (NRT) discovery of VMware events

Step	Action	Description
Using VMware vCenter Server console:		
Step 1	In the vCenter Server console, provide IMS details and configure the alarm for sending the SNMP traps.	Configure the IMS address as the SNMP trap receiver URL. Also configure the alarm to send the SNMP traps when the state of the virtual machine changes. See “Configuring the VMware vCenter Server to generate SNMP traps” on page 131.
Using the Resiliency Platform console:		

Table 16-3 Setting up near real-time (NRT) discovery of VMware events *(continued)*

Step	Action	Description
Step 2	Add the vCenter Server to the IMS as a virtualization server.	After you add the vCenter Server to the IMS, the `xtrapd` daemon on the IMS starts accepting SNMP traps from the specified vCenter Server. **Note:** If you have not configured the vCenter Server as in step 1 before adding it to the IMS, a warning message is displayed. It does not affect the vCenter Server discovery. However, near real-time discovery of VMware events is not functional. To enable the near real-time discovery subsequently, first configure the vCenter Server. Then refresh the vCenter Server configuration in the IMS using the Resiliency Platform console.

By default, near real-time discovery of VMware events is enabled. To disable it, you need to remove the IMS as the SNMP receiver in the vCenter Server and refresh the vCenter Server configuration in the IMS.

Configuring the VMware vCenter Server to generate SNMP traps

In the VMware vCenter Server console, provide the following information to configure the vCenter Server to generate SNMP traps and send them to the IMS:

- Configure the Infrastructure Management Server (IMS) as the SNMP trap receiver, as follows:

  Navigate to the SNMP configuration. Enable one of the SNMP receivers and enter the following details:

Field	Description
Receiver URL	Provide the host name of the IMS which will be connected to the vCenter Server. The vCenter Server sends the SNMP traps to this IMS. Also, configure port 162 as the SNMP port. Ensure that port 162 is not used by any other application in IMS.
Community String	Provide community string. SNMP versions v1 and v2 are supported.

- Configure an alarm for generating SNMP traps when a virtual machine state changes or any virtual infrastructure-related change occurs.

  This step includes adding an alarm to monitor the changes related to virtual machine state and vCenter Server infrastructure, and then adding the appropriate action (for example, send a notification trap).

  - You can set the alarm at an individual virtual machine level, at the data center level, or at the entire VMware vCenter Server level. It is recommended to set it at the vCenter Server level.
  - For the alarm type details, make sure to specify the following
    - Set the alarm type to monitor virtual machines
    - Set the alarm to monitor for specific events occurring on this object, for example, VM powered on
    - Enable the alarm
  - Add the required triggers to monitor the states of the virtual machine. For example, VM created, VM migrated, VM powered on, VM powered off, VM suspended, DRS VM powered on (for clustered environment with DRS enabled) and so on. The values of the fields are as follows:

For the following value of an event...	Select the following status...
VM powered on	Unset
VM powered off	Unset
DRS VM powered on	Unset
VM suspended	Unset
VM created	Unset
VM migrated	Unset
VM relocating	Unset
VM renamed	Unset
DRS VM migrated	Unset

- Add the required triggers to monitor the states of the hosts. The values of the fields are as follows:

For the following value of an event...	Select the following status...
Host disconnected	Unset
Host connected	Unset

- Add a new action to send a notification trap. Specify to send the notification trap as in the following example:

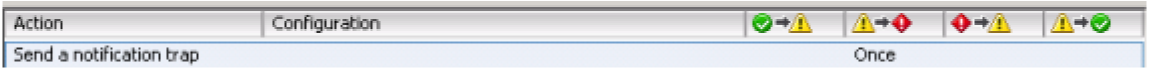

Action	Configuration	✓→⚠	⚠→◆	◆→⚠	⚠→✓
Send a notification trap			Once		

See "Setting up near real-time discovery of VMware events" on page 130.

About adding a host for discovery of VMware servers

Resiliency Platform uses a designated discovery host to discover a VMware vCenter Server. This discovery displays those ESX servers that the VMware vCenter Server manages and the virtual machines that are configured on the ESX servers.

You can use the Infrastructure Management Server (IMS) as the discovery host, or you can add a separate host to use for discovery of a vCenter server.

For example, you may want to designate a separate host to address the following situations:

- To discover a vCenter server that is behind a firewall and you do not want to install the IMS inside the firewall.
- To reduce the discovery load on the IMS for better scalability.

To designate a separate host as the discovery host for a vCenter Server:

- Add a discovery host to Resiliency Platform.
  See "Adding a host for discovery of VMware servers " on page 133.
- Select the host as the discovery host when you add the vCenter Server to Resiliency Platform.
  See "Adding VMware virtualization servers " on page 134.

Adding a host for discovery of VMware servers

The Infrastructure Management Server (IMS) can be used for discovery of VMware vCenter servers, or you can optionally add a separate host that the IMS can use for discovery of vCenter servers.

The Add Host operation for discovery of VMware virtualization installs the following add-ons on the host, in addition to the managed host package:

- Veritas Resiliency Platform Enablement add-on
- Control Host add-on

To add a host for discovery of VMware servers

1. Prerequisites
   - Ensure that the discovery host can communicate with the vCenter server and with the IMS.
   - To use a Windows host as a discovery host, a separate Windows host must first be configured as a Windows Install Host.
     See “Adding a Windows Install host” on page 112.
2. Navigate

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

   Expand the data center > **Virtualization** > **VMware** tab.

   Tip: You can also access **Manage Asset Infrastructure** from the **Quick Actions** menu.
3. Click the **Manage Discovery Host** link to add or view discovery hosts for VMware.
4. Select the **+ Discovery Hosts** button.
5. In the wizard, specify the details and click **Next**.
   - Specify the host name, username, and password
   - Select the IMS to which the host will be added
   - To add a Windows discovery host, select **This is a Windows host**.

See “About adding a host for discovery of VMware servers” on page 133.

Adding VMware virtualization servers

You can add VMware vCenter servers to Veritas Resiliency Platform for discovery by an Infrastructure Management Server (IMS). The VMware discovery provides the following information:

- Information on the vCenter Server
- Information on the ESX servers that the vCenter Server manages

When adding a vCenter Server, you have the option to automatically discover all ESX servers registered to the vCenter Server or select which of the available ESX servers to discover.

- Information on the virtual machines that are configured on the ESX servers

---

Note: If there is more than one IMS in a data center, you can add the same vCenter Server to more than one IMS. For example, you may want to split up the ESX server discovery between multiple IMSs. To accomplish this, you first add the vCenter Server to one IMS for one set of ESX servers. Then once discovery is complete, you use the Edit option on the existing vCenter Server to add it to another IMS and select a different set of ESX servers.

To add VMware virtualization servers

1. Prerequisites:

   See “Prerequisites for adding VMware virtualization servers” on page 125.

   Optionally, you can add a separate host to act as the discovery host for the vCenter Server and select it while adding the VMware server.

   See “Adding a host for discovery of VMware servers ” on page 133.

2. Navigate

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

   Expand the data center > **Virtualization** > **VMware** tab

   Launch the **+ vCenter** wizard

   Tip: You can also access **Manage Asset Infrastructure** from the **Quick Actions** menu.

3. In the wizard, specify the following information and click **Next**.

   - Specify the fully-qualified name of the vCenter Server that you want to discover along with its port number. The default port is 443.
   - When entering login credentials, an administrative vCenter Server user account is required.
   - If the data center has more than one IMS, a list of IMS names is shown. Select the IMS that you want to use to discover and monitor the vCenter Server and ESX servers.
   - If you have added a separate discovery host, select it.

4. Choose to automatically discover all ESX servers or select ESX servers to discover. If multiple clusters are available, you can use **Group By** to sort the list of ESX servers by cluster. Click **Next**. It is recommended to select all ESX servers within a cluster.

   If you choose the auto discover option, all currently available ESX servers are discovered. In addition, ESX servers later added to the vCenter Server will be automatically discovered.

5. Review the configured vCenter Server, ESX servers, and IMS on the verification screen and submit the configuration.

   The wizard notifies you of any issues.

   The vCenter Server that has been added is listed on the **VMware** tab. Discovery of the ESX servers occurs in the background. You can view the progress on the **Activities** page.

If changes are made on the virtualization servers after the IMS discovery is complete, you need to refresh the discovery of the vCenter Server.

See “Refreshing VMware vCenter Server discovery information” on page 139.

See “Editing a VMware virtualization discovery configuration” on page 136.

Editing a VMware virtualization discovery configuration

You can edit a vCenter Server discovery configuration that was added to Veritas Resiliency Platform to modify the information for an existing IMS or to add the vCenter Server to another IMS (if there is more than one in the data center). For example, you may choose to split ESX discovery between multiple IMSs.

For an existing IMS previously selected for the vCenter Server, you can modify:

- The credentials to log on to the vCenter Server
  When entering login credentials, an administrative vCenter Server user account is required.
- The ESX server discovery
  You can add or remove selected ESX servers or change between individual server discovery and auto discovery.

Note: You can only use autodiscovery if the vCenter Server is configured with only one IMS. If an existing IMS is autodiscovering the ESX servers, you are not able to add another IMS for that vCenter Server without first editing the configuration for the existing IMS to change from auto discovery.

To edit a virtualization discovery configuration

1 Navigate

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

Expand the data center > **Virtualization** > **VMware** tab

2 ⋮ On the row for the vCenter server, select the vertical ellipsis > **Edit**.

3 In the **Edit vCenter** wizard:

- To edit the vCenter configuration for the current IMS:
  If more than one IMS is listed, select the IMS that currently discovers the vCenter Server.
  Optionally, edit the credentials to log on to the vCenter.
  On the next screen, you can add or remove selected ESX servers or change between individual server discovery and auto discovery.
- To add the vCenter to a new IMS for discovery (only available if first IMS is not configured for autodiscovery):
  Select the new IMS and enter the vCenter Server logon credentials.
  On the next screen, any ESX servers not yet selected for the existing IMS are listed. Select those you want the new IMS to discover.
  To see the ESX servers already added to the first IMS, select **Show ESX in other IMS**. These are not available for selection. To return to the list of available servers, uncheck the box.

4 Review the configured vCenter Server, ESX servers, and IMS on the verification screen and proceed with the configuration.

The wizard notifies you of any issues. Discovery of the modifications occurs in the background. You can view the progress on the **Activities** page.

On the VMware tab, if you have added an IMS, both the existing and new IMS are listed on the row for the vCenter Server.

See “Managing VMware virtualization servers” on page 108.

Viewing the details of a VMware virtualization discovery configuration

You can view details for a vCenter Server that has been added to Veritas Resiliency Platform. Details include whether autodiscovery is enabled, the ESX servers that are discovered, and the number of virtual machines for each ESX server.

To view the details of a virtualization discovery configuration

1. Navigate

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

   Expand the data center > **Virtualization** > **VMware** tab

2. ⋮ On the row for the vCenter Server, select the vertical ellipsis > **Details**.

   On the Details page, if more than one IMS is configured to discover the vCenter Server, the ESX servers are grouped under the IMS that is configured to discover them.

See “Managing VMware virtualization servers” on page 108.

Removing a VMware vCenter Server discovery configuration

You can remove a VMware vCenter Server from Veritas Resiliency Platform. If the vCenter Server has been configured for discovery by more than one Infrastructure Management Server (IMS), you can choose whether to select the IMS to remove it from or remove it from all the IMSs.

To remove a VMware virtualization discovery configuration

1. Prerequisites

   Ensure that you consider how removing the virtualization discovery may affect resiliency groups. If you remove a vCenter Server (or associated ESX servers) from discovery by an IMS, any virtual machines from those ESX servers are no longer discovered and monitored by the IMS. Therefore, if any of those virtual machines are currently in a resiliency group, they will be automatically removed from the resiliency group.

   Reviewer question: Per demo feedback discussion, the decision was that if the ESX servers affected contain VMs that are part of RGs, a risk will be raised for the affected RGs - will need more info on this. Also need more info on implications for a resiliency group configured for DR -in particular. Do you need to perform any update DR operation on an affected resiliency group after removing ESX servers with VMs in the RG?

2. Navigate

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

   Expand the data center > **Virtualization** > **VMware** tab

3. On the row for the vCenter server, select the vertical ellipsis > **Remove**.

4. If more than one IMS is configured for discovery of the vCenter Server, the wizard prompts you to choose whether to remove the vCenter configuration from all IMSs or from selected IMSs. When you choose **from selected Infrastructure Management Servers**, the list of IMSs is displayed for you to make your selection.

5. Confirm that you want to remove the vCenter Server configuration from one or more IMSs. Discovery occurs in the background. You can view the progress on the **Activities** page.

See “Managing VMware virtualization servers” on page 108.

Refreshing VMware vCenter Server discovery information

You can refresh the information displayed for VMware virtualization servers that have been added to Veritas Resiliency Platform.

To refresh a virtualization discovery

1 Navigate

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

Expand the data center > **Virtualization** > **VMware** tab

2 ⋮ On the row for the vCenter Server that you want to refresh, select the vertical ellipsis > **Refresh**.

If more than one IMS is configured for discovery of the ESX servers managed by this vCenter Server, the refresh operation applies to all of the IMSs. Discovery occurs in the background. You can view the progress on the **Activities** page.

See “Managing VMware virtualization servers” on page 108.

Chapter 17

Managing enclosure assets

This chapter includes the following topics:

About the discovery host for enclosures

A Windows or Linux host is assigned the *Discovery host* role to discover the arrays in the data center. Infrastructure Management Server (IMS) is a discovery host by default. In addition to the IMS you can configure additional discovery hosts to offload the discovery and monitoring of multiple enclosures rather than discovering all of them from a single IMS.

When you add the enclosure to the IMS, you can specify a discovery host in the **Add Enclosure** wizard. The discovery host must also be added to the IMS.

The discovery host can be any virtual machine with supported operation system as required by the VRTSsfmh package. You also need the administrative credentials of the virtual machine to install the VRTSsfmh package on it. The IMS connects to this host using the 5634 port.

The discovery host also helps in situations where the IMS cannot satisfy following pre-requisites for the discovery of required enclosures:

- Specific native tools or vendor CLIs for enclosure management are required. For example, EMC Solutions Enabler for Symmetrix arrays.
- IMS does not have direct connectivity with the arrays.

For HPE 3PAR, NetApp SnapMirror, and IBM SVC Global Mirror specify the IMS as a discovery host. Separate discovery host is not required.

Note: Replication appliances are added in a different way from other types of asset infrastructure. For the RecoverPoint appliance, the IMS can be used as a discovery host or you can specify a separate discovery host.

See “Adding RecoverPoint appliance for replication” on page 166.

See “EMC Symmetrix configuration prerequisites” on page 147.

See “EMC CLARiiON configuration prerequisites” on page 152.

See “EMC VNX configuration prerequisites” on page 155.

See “NetApp configuration prerequisites” on page 155.

See “Hitachi configuration prerequisites” on page 156.

See “HPE 3PAR configuration prerequisites” on page 157.

See “Managing enclosure assets” on page 109.

Adding a discovery host

Using the Resiliency Platform console, you can add a discovery host to the Infrastructure Management Server (IMS).

To add a discovery host

1 Prerequisites

Install the vendor specific array management tools on a Windows or Linux host.

2 Navigate

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

Expand the data center > **Storage** > Enclosure tab.

Tip: You can also access **Manage Asset Infrastructure** from the **Quick Actions** menu.

3 Click the message at the bottom of the page.

4 On the **Discovery Host** page, launch the **+ Discovery Host** wizard.

5 In the wizard, specify the details and click **Next**.

See “Add Discovery host panel options” on page 143.

6 In the **Finish** panel review the information and click **Finish**.

See “Managing enclosure assets” on page 109.

Add Discovery host panel options

Use this wizard panel to configure the discovery host. Discovery host can be a Windows host or a Linux host.

Table 17-1

Field	Description
Infrastructure Management Server	Select the Infrastructure Management Server (IMS).
Host Name	Enter the host name.
User Name	Specify the user name for the host.
Password	Specify the password for the host.

Table 17-1 *(continued)*

Field	Description
Linux	Select Linux or Windows.
Windows	To add a Windows host, you need to select a Control Host. If there are no Control Host listed, you need to first install a Windows Install host. See "Adding a Windows Install host" on page 112.

Refreshing a discovery host

Using the Resiliency Platform console, you can refresh the information displayed for a discovery host.

To refresh a discovery host

1. Navigate

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

   Expand the data center > **Storage** > Enclosure tab.

   Tip: You can also access **Manage Asset Infrastructure** from the **Quick Actions** menu.

2. Click the message at the bottom of the page.
3. On the row for the discovery host that you want to refresh, select the vertical ellipsis > **Refresh**.
4. In the **Confirm and Refresh** panel, review the selection and click **Next**.
5. In the **Finish** panel review the information and click **Finish**.

Removing a discovery host

Using the Resiliency Platform console, you can remove a discovery host.

Vendor specific array management tools are installed on a discovery host that are required for discovering and managing enclosures. Before you remove the discovery host, ensure that you have installed the tools on another host, the host is added to the IMS, and the enclosures are configured using that host.

You can not remove a discovery host if any enclosure is configured using that host.

To remove a discovery host

1 Navigate

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

Expand the data center > **Storage** > Enclosure tab.

Tip: You can also access **Manage Asset Infrastructure** from the **Quick Actions** menu.

2 Click the message at the bottom of the page.

3 On the row for the discovery host that you want to remove, select the vertical ellipsis > **Remove**.

4 In the **Verify and Remove** panel, review the selection and click **Next**.

5 In the **Finish** panel review the information and click **Finish**.

Configuration prerequisites for adding storage enclosures to an IMS

For array-based replication environments, the asset infrastructure that you add to the Infrastructure Management Server (IMS) includes storage enclosures and the discovery host.

For Resiliency Platform to discover replication-related device groups, you must also add the hosts that contain those device groups to the IMS.

Table 17-2 Configuration prerequisites

For information on configuration prerequisites of	See the following
EMC Symmetrix	See the section called "Physical connection requirements" on page 147. See the section called "Device setup requirements" on page 147. See the section called "Replication requirements" on page 147. See "Configuring the remote SYMAPI server for EMC Symmetrix array discovery" on page 148. See "Verifying the configuration of a remote SYMAPI server" on page 151.
EMC Clariion	See "Using Password Security file for CLARiiON enclosure" on page 152. See "Verifying NaviSecCLI communication with CLARiiON enclosure" on page 153.
EMC VNX	See "EMC VNX configuration prerequisites" on page 155.
NetApp	See the section called "Physical connection requirements" on page 155. See the section called "Device setup requirements" on page 156.
Hitachi	See the section called "Physical connection requirements" on page 157. See the section called "Device setup requirements" on page 156. See the section called "Replication requirements" on page 157.
3PAR	See "HPE 3PAR configuration prerequisites" on page 157.

For more information on replication requirements, refer to the Solutions guides.

See "About the discovery host for enclosures" on page 141.

See "Adding storage enclosures " on page 160.

EMC Symmetrix configuration prerequisites

For the Infrastructure Management Server (IMS) to discover EMC Symmetrix storage arrays, ensure that your storage network's physical connections and device settings are properly configured.

Physical connection requirements

The physical connection requirements are as follows:

- Fibre Channel connection between each Symmetrix array and the SAN fabric.

Device setup requirements

The device setup requirements include the following:

You configure an array for discovery using the EMC Symmetrix Command Line Interface (SYMCLI). The SymCLI utilities must be configured on a discovery host. Install EMC Solutions Enabler (SYMCLI) on the discovery host.

IMS can also use the EMC Symmetrix Remote Data Facility (SRDF).

Veritas Resiliency Platform supports SYMCLI 7.x for IMS discovery of the EMC Symmetrix storage enclosures.

For the complete information on supported hardware and software, see the *Hardware and Software Compatibility List (HSCL)*.

The IMS discovers all in-band Symmetrix storage arrays with a Fibre Channel or SCSI connection to a discovery host where SYMCLI is installed.

The IMS also supports discovery of EMC Symmetrix storage arrays through remote SYMAPI servers. This discovery method does not require in-band array connectivity to the discovery host specified in the array configuration. However, the host on which the SYMAPI server is running must have in-band connectivity with the Symmetrix array.

For the IMS to discover EMC Symmetrix arrays using a remote SYMAPI server, you must specify the remote SYMAPI server while configuring the enclosure in IMS.

See "Configuring the remote SYMAPI server for EMC Symmetrix array discovery" on page 148.

See "Verifying the configuration of a remote SYMAPI server" on page 151.

Replication requirements

To discover and manage EMC SRDF replication technology, you need to create appropriate device groups on hosts which have EMC Symmetrix Gatekeeper devices mapped. Device groups need to be defined on such hosts in both the production and recovery data centers. You must add these hosts to the local IMSs so that Resiliency Platform can discover the device groups.

See “Adding application hosts ” on page 114.

For more information on requirements for array-based replication, see the Solutions guides.

See “Managing enclosure assets” on page 109.

Configuring the remote SYMAPI server for EMC Symmetrix array discovery

The Infrastructure Management Server (IMS) supports the discovery of EMC Symmetrix arrays with a remote SYMAPI server mechanism. This discovery method does not require in-band array connectivity to the host from which the EMC Symmetrix array is discovered.

For the IMS to discover EMC Symmetrix arrays using a remote SYMAPI server, you need to configure the SYMAPI server. To configure the remote SYMAPI server in your environment, you need to perform two tasks:

- Ensure that the remote SYMAPI server daemon is running on the server where the EMC Solutions Enabler is installed. See the following procedure:

  To ensure that the remote SYMAPI server daemon is running on the server where the EMC Solutions Enabler is installed

- Ensure that the EMC Solutions Enabler on the discovery host can communicate with the remote SYMAPI server. See the following procedure:

  To ensure that the EMC Solutions Enabler can communicate with the remote SYMAPI server

To ensure that the remote SYMAPI server daemon is running on the server where the EMC Solutions Enabler is installed

1 Log on with the administrative credentials to the host that you want to use as the remote SYMAPI server and which has in-band connectivity to the EMC Symmetrix array.

2 Type the following command on the host:

```
stordaemon list
```

An example of the daemon list appears.

```
Available Daemons   ('[*]': Currently Running):
[*]  storapid        EMC Solutions Enabler Base Daemon
     storgnsd        EMC Solutions Enabler GNS Daemeon
     storrdfd        EMC Solutions Enabler RDF Daemon
     storevntd       EMC Solutions Enabler Event Daemon
[*]  storwatchd      EMC Solutions Enabler Watchdog Daemon
     storsrmd        EMC Solutions Enabler SRM Daemon
     storstpd        EMC Solutions Enabler STP Daemon
     storsrvd        EMC Solutions Enabler SYMAPI Server Daemon
[*]  storsrvdInst    >>> Running Instance of storsrvd <<<
```

The name for the remote SYMAPI server daemon is `storsrvd`. If you see a [*] for `storsrvd`, that means the remote SYMAPI server daemon is already running on the host. If the daemon is running, proceed to the next procedure.

To ensure that the EMC Solutions Enabler can communicate with the remote SYMAPI server

3 Type the following commands to start the `storsrvd` daemon:

```
stordaemon start storsrvd

   Waiting for daemon to start. This may take several seconds.

stordaemon list
```

An example of the daemon list appears.

```
Available Daemons   ('[*]': Currently Running):
[*]  storapid         EMC Solutions Enabler Base Daemon
     storgnsd         EMC Solutions Enabler GNS Daemeon
     storrdfd         EMC Solutions Enabler RDF Daemon
     storevntd        EMC Solutions Enabler Event Daemon
[*]  storwatchd       EMC Solutions Enabler Watchdog Daemon
     storsrmd         EMC Solutions Enabler SRM Daemon
     storstpd         EMC Solutions Enabler STP Daemon
[*]  storsrvd         EMC Solutions Enabler SYMAPI Server Daemon
```

4 Perform steps 1 and 2 on each host in which you want to configure the remote SYMAPI server.

To ensure that the EMC Solutions Enabler can communicate with the remote SYMAPI server

1 Install EMC Solutions Enabler on the Discovery Host.

2 Change to the SYMAPI configuration directory. By default, the directory is:

- Linux — /var/symapi/config
- Windows — %PROGRAMFILES%\EMC\SYMAPI\config

3 Modify the file "netcnfg" in the SYMAPI configuration directory of the host where the EMC Solutions Enabler is installed. Append the entry for the configured SYMAPI server(s) to the end of the file. The following is an example of adding entries for two SYMAPI servers:

```
#SYMAPI_SERVER - TCPIP node001 WWW.XXX.YYY.ZZZ 2707 -

DC1_SERVER - TCPIP ctrlhost_1 10.200.15.155 2707 -

DC2_SERVER - TCPIP ctrlhost_2 10.249.100.155 2707 -
```

See “EMC Symmetrix configuration prerequisites” on page 147.

Verifying the configuration of a remote SYMAPI server

Verify the remote SYMAPI server configuration before you perform the device setup requirements. Set environment variables to test if the SYMAPI server is configured correctly.

To verify the configuration of a remote SYMAPI server

1. Open an operating system console and log on to the host as root (Linux) or as a user with administrator-level privileges (Windows).
2. Ensure that the SYMCLI commands are in your `PATH` environment.
3. Do one of the following:
   - On Linux, run the following SYMCLI commands to set the server's environment variables:

     ```
     SYMCLI_CONNECT_TYPE=REMOTE; export SYMCLI_CONNECT_TYPE
     SYMCLI_CONNECT=DC1_SERVER; export SYMCLI_CONNECT
     symcfg list
     ```

   - On Windows, run the following SYMCLI commands to set the server's environment variables:

     ```
     set SYMCLI_CONNECT_TYPE=REMOTE
     set SYMCLI_CONNECT=DC1_SERVER
     symcfg list
     ```

4. Ensure that the arrays on different remote SYMAPI server hosts are discovered correctly.

   If you get an error in the output (instead of a list of the Symmetrix arrays), verify that your EMC Solutions Enabler is configured properly. If it is not configured properly, consult the EMC Solutions Enabler install guide for the commands. The install guide provides the detailed instructions on configuring the SYMAPI server and related commands.

5. To unset the environment variables, type the following commands:

   ```
   unset SYMCLI_CONNECT_TYPE
   unset  SYMCLI_CONNECT
   ```

See “EMC Symmetrix configuration prerequisites” on page 147.

EMC CLARiiON configuration prerequisites

The Infrastructure Management Server (IMS) communicates to the EMC CLARiiON enclosures through the NaviSphere Secure CLI (NaviSecCLI) utility, which is the secure CLI for communicating to the NaviSphere package on the enclosure. The NaviSphere Secure CLI utility must be installed on a discovery host. A discovery host can be a Windows or Linux host that you add as an asset to the IMS.

The IMS supports only the EMC CLARiiON CX series enclosures with a firmware version 6.26 and later.

To discover EMC CLARiiON enclosures, ensure that network physical connections and Navisphere CLI are properly configured.

For more information, refer to *Hardware and Software Compatibility List (HSCL).*

Physical connection requirements

The physical connection requirements for EMC CLARiiON enclosure are as follows:

- Network connectivity between the enclosure and discovery host.

Device setup requirements

The device setup requirements for EMC CLARiiON enclosure include using password security file and verifying NaviSecCLI communication with CLARiiON enclosure.

See “Using Password Security file for CLARiiON enclosure” on page 152.

See “Verifying NaviSecCLI communication with CLARiiON enclosure” on page 153.

Using Password Security file for CLARiiON enclosure

To use Secure Sockets Layer for the discovery of CLARiiON array, you need to use Password Security file on the discovery host.

To put the password in a security file

1. Log on as root to the host that manages CLARiiON array.
2. Run the following NaviSecCLI command to create the security file:

   ```
   naviseccli -user userName -password passWord -scope
   userScope AddUserSecurity [-secfilepath
   secFileFolderName]
   ```

   Where:

   userName is a valid user name for the account in the array.

   passWord is the password for the *userName*.

   userScope is the scope of the user that you specified when you created the account in the array. It can be either local (userScope=1) or global (userScope=0).

   secFileFolderName is the directory where you want the security file to reside. This directory is the path that you specify when you configure an array for the discovery.
3. If some CLARiiON arrays have a different user name, repeat step 2 for each user name, specifying a different `secFileFolderName`.

   If the `secFileFolderName` was specified in step 2, use the same name when configuring the corresponding arrays.

Verifying NaviSecCLI communication with CLARiiON enclosure

Use the following procedure to verify the communication of NaviSecCLI with the CLARiiON enclosure.

To verify that the NaviSecCLI can communicate with CLARiiON arrays

1. Enter the following command at the command prompt of CLARiiON management host:

   ```
   naviseccli -h arrayIPAddress -user userName -password
   passWord -scope userScope networkadmin -get
   ```

   Where:

   arrayIPAddress is the IP address, the fully qualified domain name, or the name of one of the storage processors in the array.

   userName is a valid user name for the account in the array.

   passWord is the password for *userName*.

   userScope is the scope of the user that you specified when you created the account in the array. It can be Local, Global, or LDAP.

2. Review the information that displays:

   - If NaviSecCLI can communicate with the arrays, the following information is displayed:

     ```
     Storage Processor:                  SP B
     Storage Processor Network Name:     cx500-2spb
     Storage Processor IP Address:       10.100.18.18
     Storage Processor Subnet Mask:      255.255.248.0
     Storage Processor Gateway Address: 10.100.16.1
     ```

     If this command succeeds, the Management Server uses the credentials to communicate with the array.

   - If NaviSecCLI cannot communicate with the array, the following information is displayed:

     ```
     Broken Pipe

     Valid IP Address with Feature not installed:
     naviseccli -h cx500a -user admin -password password -scope 0
     networkadmin -get
     Management Server - Feature software is not installed or the
     command may have been typed incorrectly
     usage:
       metalun
       migrate
       connection
       mirror
     ```

```
snapview
analyzer
```

EMC VNX configuration prerequisites

Infrastructure Management Server (IMS) discovery can be performed on EMC VNX (block) storage enclosures. You can specify the scope for VNX block as global (value 0) or local (value 1) when you add the enclosure in the IMS.

Configuration steps for VNX block are similar to EMC CLARiiON.

See "EMC CLARiiON configuration prerequisites" on page 152.

NetApp configuration prerequisites

For the Infrastructure Management Server (IMS) to discover a NetApp enclosure, ensure that the storage network physical connections and NetApp server are properly configured.

The NetApp storage objects work on the Data ONTAP operating system, which provides various interfaces to administer the NetApp storage objects. The IMS communicates to the enclosures using the ONTAP SDK interface to get the NetApp enclosure information. This communication occurs through the HTTP protocol (using the port number 80) or through the HTTPS protocol (using the port number 443).

The IMS supports NetApp enclosures that have Data ONTAP 1.4 or later.

Resiliency Platform supports Netapp storage with volumes or qtrees mounted via NFS or with LUNs provisioned via Fibre Channel (FC).

Physical connection requirements

The physical connection requirements for NetApp enclosure discovery are as follows:

- Network connectivity between the discovery host and NetApp enclosure.
- You should be able to connect from the discovery host to NetApp enclosure or server using HTTP and HTTPS connections. Use the following URLs to access the enclosure:
  - https://*netapp_address/na_admin*
    Port 443 is used for HTTPS connection.
  - http://*netapp_address/na_admin*
    Port 80 is used for HTTP connection.
  - *netapp_address* is the IP address or NetApp array name, registered with the Domain Name System (DNS).

Device setup requirements

Setting up the device includes NetApp server configuration and enabling support for MultiStore Virtual Systems on NetApp enclosure.

Configure the array with an IP address or name, and an administrator-level account with valid user name and password. These credentials are used by IMS to access the enclosure for discovery.

Ensure you turn on the following options in the NetApp enclosure. These are required for NetApp SnapMirror operations.

- httpd.admin.enable
- httpd.enable

Ensure that the following licenses are installed and enabled:

- licensed_feature.multistore.enable (required for discovering IP addresses)
- licensed_feature.flex_clone.enable (required for rehearsal operation)

Hitachi configuration prerequisites

For the Infrastructure Management Server (IMS) to discover Hitachi TrueCopy storage arrays, ensure that your storage network's physical connections and device settings are properly configured.

Physical connection requirements

The physical connection requirements are as follows:

- Fiber channel or Small Computer System Interface (SCSI) should directly attach the hosts to the Hitachi TrueCopy array that contains the TrueCopy P-VOL / S-VOL devices.

Device setup requirements

The device setup requirements include the following:

- The host must be configured with Command Control Interface (CCI) provided by Hitachi TrueCopy.
- The IMS discovers all Hitachi storage arrays with a Fiber Channel or SCSI connection to a discovery host where CCI is installed. It discovers TrueCopy as well as Shadow Instances of Hitachi.
- The IMS Supports TrueCopy on all microcode levels on all the arrays, only if the host, HBA, and array combination is supported by Hitachi.
- Veritas Resiliency Platform supports All Levels of CCI for IMS discovery of the Hitachi TrueCopy storage enclosures.

Replication requirements

To discover and manage Hitachi TrueCopy/HUR replication technology, you need to create appropriate HTC/HUR device groups on hosts which have Hitachi Command Devices mapped. Device groups need to be defined on such hosts in both the production and recovery data centers. You must add these hosts to the local IMSs so that Resiliency Platform can discover the device groups.

See “Adding application hosts ” on page 114.

For more information on requirements for array-based replication, see the Solutions guides.

HPE 3PAR configuration prerequisites

To discover HPE 3PAR enclosure, ensure that the storage network connection between the enclosure and the discovery host is properly configured. The discovery host communicates directly with 3PAR enclosure using the IP address that is provided on the Device Configuration panel. The discovery uses SSH communication.

For the Infrastructure Management Server (IMS) to discover HPE 3PAR enclosures, ensure that your storage network’s physical connections and device settings are properly configured.

Physical connection requirements

The hosts should connect to the 3PAR array that contains the RemoteCopy Primary and Secondary devices using Fiber channel or Small Computer System Interface (SCSI).

Replication requirements

The discovery host should be able to communicate with HPE 3PAR array using Secure Shell (SSH). Periodic modes of replication with the mirror_config policy is supported.

REVISED CONTENT

To discover HPE 3PAR enclosure, ensure that the storage network connection between the enclosure and the discovery host or Infrastructure Management Server (IMS) is properly configured.

By default the IMS is the discovery host. You can designate any Windows or Linux host as a discovery host.

The discovery host or the IMS communicates directly with 3PAR enclosure using the IP address that is provided on the Device Configuration panel. The discovery uses SSH communication.

For IMS to discover HPE 3PAR enclosures, ensure that the physical connections of the storage network and the device settings are properly configured.

Physical connection requirements

The hosts should connect to the 3PAR array that contains the RemoteCopy Primary and Secondary devices using Fiber channel or Small Computer System Interface (SCSI). Here the hosts are the virtualization servers that consume storage from the array.

Replication requirements

The discovery host should be able to communicate with HPE 3PAR array using Secure Shell (SSH). Periodic modes of replication with the mirror_config policy is supported.

IBM SVC Global Mirror configuration prerequisites

For the Infrastructure Management Server (IMS) to discover IBM SVC Global Mirror storage arrays, ensure that the physical connections of the storage network and the device settings are properly configured.

Physical connection requirements

The hosts should connect to the SVC array that contains the Global Mirror master or auxiliary vdisks using Fiber channel or Small Computer System Interface (SCSI). Here the hosts are the virtualization servers that consume storage from the array.

Device setup requirements.

- Ensure that the IMS is able to communicate with IBM SVC array using SSH.
- Ensure that that IBM SVC array user has the role of either ‘Administrator’ or ‘SecurityAdmin’ to perform IBM SVC Global Mirror replication operations.
- Ensure that the Global Mirror, FlashCopy consistency group, relationships, and vdisks do not have names ending with `_vrp`.

IBM XIV configuration prerequisites

For the Infrastructure Management Server (IMS) to discover IBM XIV storage arrays, ensure that the physical connections of the storage network and the device settings are properly configured.

The Infrastructure Management Server (IMS) communicates with IBM XIV enclosure through the IBM XIV Command Line Interface (XCLI) utility. The IBM XCLI utility must be installed on a discovery host. A discovery host can be a Windows or Linux host that you add as an asset to the IMS.

Physical connection requirements

The physical connection requirements for IBM XIV enclosure are as follows:

- Network connectivity between the enclosure and the discovery host.
- Network connectivity between the IMS and the discovery host.

Device setup requirements.

The device setup requirement for IBM XIV enclosure includes verifying IBM XCLI communication with the IBM XIV enclosure. Use the below procedure to verify the communication.

Enter the following command at the command prompt of the IBM XCLI management host.

```
xcli  -x -y -u userName -p passWord -m  arrayIPAddress config_get
```

Where,*arrayIPAddress* is the IP address, or the fully qualified domain name, or the name storage array.*userName* is a valid user name for the account in the array.*passWord* is the password for userName.

Below information is displayed if the communication is successful.

```
<XCLIRETURN STATUS="SUCCESS" COMMAND_LINE="-x -y -u vrpuser -p ******
-m 172.31.255.255 config_get">
<OUTPUT>
<config id="0">
                <name value="dns_primary"/>
                <value value="172.31.255.255"/>
                <level value="User"/>
            </config>
            <config id="1">
                <name value="dns_secondary"/>
                <value value="172.31.255.255"/>
                <level value="User"/>
            </config>
            <config id="2">
 ......
            </config>

  </OUTPUT>
</XCLIRETURN>
```

Adding storage enclosures

You can add storage enclosures (arrays) to Resiliency Platform for discovery by an Infrastructure Management Server (IMS). This is required only to monitor array-based replication. For environments that are using other types of replication, you do not need to add storage enclosures.

To add storage enclosures

1 Prerequisites

Ensure that you have configured the storage array for discovery.

See “Configuration prerequisites for adding storage enclosures to an IMS” on page 145.

If you have configured a discovery host, ensure that you have the name.

See “About the discovery host for enclosures” on page 141.

2 Navigate

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

Expand the data center > **Storage**

Click the required enclosure tab.

Launch the **+ Enclosure** wizard.

Tip: You can also access **Manage Asset Infrastructure** from the **Quick Actions** menu.

3 In the **Configure Server** panel, specify the details and click **Next**.

See “Configure Server panel options” on page 160.

4 In the **Select Enclosure** panel, select the enclosures for deep discovery and click **Next**.

See “Select Enclosure panel options to enable discovery for selected enclosures” on page 165.

5 Review the information and click **Finish**.

See “Managing enclosure assets” on page 109.

Configure Server panel options

Use this wizard panel to select the array model and specify the server details for deep discovery of enclosures.

Table 17-3 Configure Server panel options

Field	Description
Select Technology	Select the array model for which you want to enable the deep array discovery. Tip: Information required to complete the configuration depends on your selection of the array model. See Table 17-4 on page 161. for EMC enclosures. See Table 17-5 on page 163. for Hitachi enclosures. See Table 17-6 on page 164. for NetApp enclosures. See Table 17-7 on page 164. for IBM enclosures. See Table 17-8 on page 165. for HPE enclosures.
IMS	If multiple Infrastructure Management Servers (IMS) are configured then you can select one, else the selection option is not displayed.
Discovery Host	Displays the discovery host. If there are multiple discovery hosts, then you can select from the list. Else the selection option is not displayed.

Table 17-4 Add Enclosure panel options for EMC enclosures

Field	Description
For Clariion and VNX:	
User Name	Specify the user name for the enclosure.
Password	Specify the password for the enclosure.
For Clariion:	
Storage Array Name/IP	Name of the IP address of the storage enclosure.

Table 17-4 Add Enclosure panel options for EMC enclosures *(continued)*

Field	Description
Port	The port for communicating with the enclosure for getting the information. The default port number is 443. Ensure the port that you specify here is enabled on the enclosure.
Scope	Specifies the type of the user account on the storage system that you want to log on. The available options are: ■ Global: Choose this option if your account is effective throughout the domain. When the administrator creates a global account, the software copies the definition of this account to the domain directory, which makes it accessible on all storage systems in the domain. ■ Local: Choose this option if your account is effective only on the storage systems for which the administrator creates the account. Using the local account, you can log on to only those storage systems on which you have a local account. ■ LDAP: LDAP maps the user name and the password entries to an external LDAP or Active Directory server for authentication. The user name and the password pairs whose roles are not mapped to the external directory are denied access. For authentication within the local security directory, specify global or local scope.
NAVISEC CLI Location	The location of the NaviSecCLI binary on the discovery host.
Use Secure Sockets Layer	Select this check box to use the secure socket layer for communicating to the enclosure. If you select this check box, you do not have to enter the credentials again when you perform deep discovery for the EMC CLARiiON enclosures.

Table 17-4 Add Enclosure panel options for EMC enclosures *(continued)*

Field	Description
Certificate Location	Specify the security file location on the discovery host.
For Symmetrix:	
SYMAPI Server (Optional)	Specify the SYMAPI Server name that is configured on the discovery host to discover the enclosures. Use this option if the discovery host does not have visibility to gatekeeper devices for Symmetrix enclosures.
SymCLI Location	Specify the location of the SymCLI binaries. SymCLI must be functional to discover the array details. Refer to the enclosure configurations prerequisites section for more details. See "EMC Symmetrix configuration prerequisites" on page 147.
For VNX:	
Block IP/Name:Scope	IP address or name and scope for VNX Block. The IP address or name, and the scope must be separated using a colon. < are we supporting VNX file? if yes, is this field required for VNX file?> For global scope, enter the value as 0. For local scope, enter the value as 1.
CLI Location	Enter the Navisphere CLI location.

Table 17-5 Add Enclosure panel options for Hitachi enclosures

Field	Description
HiCommand server	Enter the name of the HiCommand server.
Protocol	Select the security protocol.

Table 17-5 Add Enclosure panel options for Hitachi enclosures *(continued)*

Field	Description
Port	The port for the HiCommand server. The default port number is 2001. Ensure the port that you specify here is enabled on the enclosure.
User Name	Specify the user name for the enclosure.
Password	Specify the password for the enclosure.

Table 17-6 Add Enclosure panel options for NetApp enclosures

Field	Description
NetApp Server	Enter the name or the IP address for the NetApp server.
Port	Enter the port number. Enter 80 for communicating over HTTP. For communicating over HTTPS, enter 443. Ensure the port that you specify here is enabled on the enclosure.
User Name	Specify the user name for the enclosure.
Password	Specify the password for the enclosure.

Table 17-7 Add Enclosure panel options for IBM enclosures

Field	Description
Enclosure	Enter the name or the IP address of the IBM enclosure.
User Name	Specify the user name for the enclosure.
Password	Specify the password for the enclosure.
CLI Location	The location of the CLI binary on the discovery host.

Table 17-8 Add Enclosure panel options for HPE enclosures

Field	Description
3PAR Enclosure IP/Name	Enter the name or the IP address for the 3PAR enclosure.
User Name	Specify the user name for the enclosure.
Password	Specify the password for the enclosure.

See “Adding storage enclosures ” on page 160.

Select Enclosure panel options to enable discovery for selected enclosures

Use this wizard panel to enable deep discovery for selected enclosures.

Select the check box in the top row to select all the enclosures in the list.

Table 17-9 Add Enclosure panel to enable deep discovery information for enclosures

Field	Description
Display Name	Displays the name of the enclosure.
Enclosure Vendor ID	Displays the ID that is generated for the enclosure.
Serial	Displays the serial number of the enclosure.
Vendor	Displays manufacturer of the enclosure.
Model	Displays the enclosure model information.
Product	Displays the type of the enclosure.

See “Adding storage enclosures ” on page 160.

Configuration prerequisites for adding replication appliances

For appliance based replication, you need to add the appliance to Resiliency Platform for discovery by an Infrastructure Management Server (IMS).

EMC RecoverPoint configuration prerequisites

For the Infrastructure Management Server (IMS) to discover EMC RecoverPoint appliance, ensure that following prerequisites are met:

- The IMS supports EMC RecoverPoint 4.1.
- The discovery host should be able to communicate with RecoverPoint using Secure Shell (SSH).

Adding RecoverPoint appliance for replication

In Veritas Resiliency Platform console, you can add a RecoverPoint appliance to provide continuous data protection with multiple recovery points to restore the applications instantly to a specific point in time.

To add RecoverPoint appliance for replication

1. Prerequisites

   Ensure that you have the name of the discovery host.

   See “About the discovery host for enclosures” on page 141.

2. Navigate

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

   Expand the data center > **Data Mover**

   Click the **RecoverPoint** tab.

   Launch the **+ RecoverPoint** wizard

   Tip: You can also access **Manage Asset Infrastructure** from the **Quick Actions** menu.

3. In the **Add RecoverPoint** panel, specify the details and click **Next**.

   See “Add RecoverPoint panel options for configuration details for RecoverPoint” on page 167.

4. In the **Configure RecoverPoint** panel, review the configuration and click **Finish**.

Add RecoverPoint panel options for configuration details for RecoverPoint

Use this wizard panel to specify the details of the devices and the server information for adding the RecoverPoint appliance for replicating a virtual machine.

Table 17-10 Add RecoverPoint panel options

Field	Description
Infrastructure Management Server	If multiple Infrastructure Management Servers (IMS) are configured then you can select one, else the selection option is not displayed.
Discovery Host	Select the name of the discovery host.
RecoverPoint Server	Enter the name of the RecoverPoint appliance.
Username	Enter the user name for the RecoverPoint appliance.
Password	Enter the password for the RecoverPoint appliance.

See "Adding RecoverPoint appliance for replication" on page 166.

Editing the discovery configuration for an enclosure

You can edit details for enclosure configurations that were added previously to Resiliency Platform.

To edit the discovery configuration for an enclosure

1. Navigate

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

   Expand the data center > **Storage**

   Click the required enclosure tab.

   Tip: You can also access **Manage Asset Infrastructure** from the **Quick Actions** menu.

2. On the row for the enclosure that you want to edit, select the vertical ellipsis > **Edit**.

3. In the **Configure Server** panel, edit the configuration details to change the device discovery. Select **Next**.

   Note: The configuration details are the same as when adding the enclosure.

   See “Configure Server panel options” on page 160.

4. In the next panel, select the check box for the enclosures for which you want to perform the deep discovery configuration. Select **Next**.

5. In the result panel review the information and select **Finish**.

See “Managing enclosure assets” on page 109.

Removing the discovery configuration for an enclosure

You can remove enclosure configurations that were added previously to Resiliency Platform.

To remove the discovery configuration for an enclosure

1. Navigate

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

   Expand the data center > **Storage**

   Click the required enclosure tab.

   Tip: You can also access **Manage Asset Infrastructure** from the **Quick Actions** menu.

2. On the row for the enclosure that you want to remove, select the vertical ellipsis > **Remove**.
3. In the **Verify & Remove** panel, review the information and click **Next**.
4. Click **Finish** to exit the wizard.

See "Managing enclosure assets" on page 109.

Refreshing enclosure discovery information

You can submit a refresh request to update the information displayed on the table of enclosures assets that have been added to Resiliency Platform.

To refresh an enclosure configuration discovery

1. Navigate

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

   Expand the data center > **Storage**

   Click the required enclosure tab.

   Tip: You can also access **Manage Asset Infrastructure** from the **Quick Actions** menu.

2. On the row for the enclosure that you want to refresh, select the vertical ellipsis > **Refresh**.
3. In the **Confirm and Refresh** panel, review the information and click **Next**.
4. Click **Finish** to exit the wizard.

Section 5

Managing users and global product settings

Managing licenses

This chapter includes the following topics:

Managing licenses

Using the Veritas Resiliency Platform console, you can install, view, and manage the licenses. You can also view the report that provides details about the licenses that are deployed for various Veritas Resiliency Platform solutions.

See "About licenses" on page 171.

See "Viewing and managing licenses" on page 172.

See "Viewing the License Entitlement report" on page 173.

About licenses

To create resiliency groups using virtual machines or applications, you need to install a subscription license for Veritas Resiliency Platform. The license is provided for a predefined number of virtual machines for a set duration of time. The license for physical servers is provided for a predefined number of CPU cores for a set duration of time. The extension of the license file is .slf. You can install the file using the Resiliency Platform console.

During the initial setup, a demo license is made available. This demo license is valid for 60 days, letting you evaluate the Resiliency Platform. Before the expiry date, daily notifications are sent based on the warning period that is specified in the license file. You need to purchase a subscription if you intend to use the Resiliency Platform beyond the expiry date of the demo license.

After a subscription has expired, you can continue to perform operations on the resiliency groups that are already created. However create new resiliency group operation is disabled. Note that to be in compliance you are required to repurchase the subscription to continue using Veritas Resiliency Platform.

See “Viewing and managing licenses” on page 172.

See “Viewing the License Entitlement report” on page 173.

Viewing and managing licenses

You can install and view the licenses using the Veritas Resiliency Platform console. The extension of a license file is .slf.

You can view the following information about the installed licenses in a table:

- Name: Name of the license.
- Meter Type: Licenses for applications on physical hosts are categorized under **Per Core** meter type whereas licenses for virtual machines and applications on virtual machines are categorized under **Per Virtual Machine** meter type. If resiliency groups are created using applications that are protected using the Data Mover, then **Per Giga Byte** meters are consumed.
- Type: Type of the license, demo or permanent.
- Version: License version number.
- Purchased Quantity: Number of meters purchased.
- Start Date: The date on which the license is installed.
- Expiry Date: Expiry date of the license.
- Valid For (Days): Indicates the number of days the license is valid for.

To install a license

1. Navigate

   Settings (menu bar) > **Settings** > **Licenses**

2. Click **Browse** to select the .slf file and click **Install License**.

See “About licenses” on page 171.

See “Viewing the License Entitlement report” on page 173.

Viewing the License Entitlement report

This report provides details about the licenses that are deployed for various Veritas Resiliency Platform solutions.

You can view the following information in the table for licenses deployed on physical hosts and virtual machines:

- Total number of subscriptions and expired subscriptions
- Purchased, used, and available quantity
- Number of unmanaged assets

In the **Details** table, you can view the additional information about all the licenses deployed. Information such as the license type (demo or permanent), version, purchased quantity, start and expiry date, and the status of the license.

To view the License Entitlement report

1. Navigation

   Reports (menu bar) > **Inventory Reports**.

2. Click **Run** on the **License** report to view the report in the HTML format or save as a comma-separated (.csv) file.

   Click **Schedule** on the **License Entitlement** report to receive the report on the specified email address.

See “About licenses” on page 171.

See “Viewing and managing licenses” on page 172.

Chapter 19

Managing user authentication and permissions

This chapter includes the following topics:

Managing user authentication and permissions

Veritas Resiliency Platform provides a console for viewing information and performing operations. Managing user authentication and permissions for the console involves the following tasks.

Table 19-1 Process for setting up user authentication and permissions

Task	Details
Configure authentication domains	You can add multiple authentication domains. See "About user authentication in the web console" on page 176. See "Configuring authentication domains " on page 183. See "Unconfiguring authentication domains" on page 187.
Configure user groups and users	Once you configure an authentication domain, you can configure user groups or users for Resiliency Platform from that authentication domain. See "Configuring user groups and users" on page 188.
Assign permissions to groups and users	When you configure user groups or users for Resiliency Platform, they are by default assigned the Guest persona, which gives permission to view information in the web console. Permission to perform operations in the console requires assigning additional personas. For some personas, you can also limit the scope of the operation to selected objects, for example, resiliency groups. See "About user permissions in the web console" on page 176. See "Predefined personas" on page 178. See "About limiting object scope for personas" on page 182. See "Assigning permissions to user groups and users" on page 189. You can also create custom personas. See "Adding custom personas" on page 190. See "Predefined jobs that can be used for custom personas" on page 191.

Table 19-1 Process for setting up user authentication and permissions *(continued)*

Task	Details
Configure Windows global user	To customize the static IP of Windows guest virtual machines in the VMware environment, you need to provide the administrator user name and password to log on to the Windows virtual machines. See "Configuring Windows global user" on page 194.

About user authentication in the web console

By default, the Admin user of the Veritas Resiliency Platform virtual appliance can log in to the web console with access to all views and operations.

The Admin user can configure authentication domains from external identity providers such as Active Directory (AD) and LDAP.

Once an authentication domain is configured, the Admin user can configure user groups and users for Resiliency Platform from that domain. These users can log in to the console with their domain login credentials.

All users and groups that are configured for Resiliency Platform have permission by default to view everything in the web console but not to perform any operations. Permissions for operations must be assigned separately by assigning the appropriate personas to users and groups.

It is recommended not to remove the default Resiliency Platform users or reduce the permissions of the default Resiliency Platform users.

If you change the password of a user who was configured to logon to the domain, you need to edit the configured domain and enter the new password for the user.

See "Editing authentication domains" on page 187.

See "Managing user authentication and permissions" on page 175.

About user permissions in the web console

Veritas Resiliency Platform uses the concepts of personas, job, and objects to define permissions for users in the web console.

Persona

A role that has access to a predefined set of jobs (operations).

The product comes with a set of predefined personas.

See “Predefined personas” on page 178.

You can also add custom personas.

See “Adding custom personas” on page 190.

See “Predefined jobs that can be used for custom personas” on page 191.

All users and groups that are added to Resiliency Platform have the Guest persona by default. The Guest persona allows users to view everything in the web console but not to perform any operations.

Job

A type of task (operation) that a user can perform.

Examples:

Manage resiliency groups

Manage assets

Perform disaster recovery of resiliency groups

Object types and scope

Each job can be performed on certain types of Resiliency Platform objects. Types of objects include data centers, resiliency groups, and virtual business services.

See “About Resiliency Platform features and components” on page 15.

When you assign a persona to a user or group, you define the scope of some jobs by selecting from available objects. For some jobs, the scope is the resiliency domain, which would be the entire scope of the product deployment.

If you want a user to have permissions that are different from the user group to which they belong, you must add the user individually to Resiliency Platform. Permissions assigned at the individual user level override the permissions that the user has as a user group member.

If a user tries to perform an operation for which they do not have authorization, a message is displayed to notify them of the fact; in addition an entry for "authorization check failed" is available in the audit logs.

See “Managing user authentication and permissions” on page 175.

Predefined personas

The following table lists the predefined personas for Veritas Resiliency Platform and their associated jobs and objects. You can assign one or more of these personas to a user or user group to define permissions. Some jobs let you limit the scope by specifying the assets (resiliency groups) on which permissions are assigned.

You can also create custom versions of these personas, except for the Guest and Super admin persona.

Table 19-2 Predefined personas and jobs

Persona	Description and scope	Jobs
Super admin	Can perform all operations on all objects in resiliency domain.	All jobs All objects in resiliency domain
Resiliency Platform admin	Manage Resiliency Managers and Infrastructure Management Servers (IMSs) and data centers. Manage assets. Manage user security settings and other product settings. Manage product updates. Scope: Resiliency domain.	**Manage assets (jobs separated by type):** ■ Manage host assets ■ Manage virtualization assets ■ Manage data mover assets ■ Manage application cluster assets ■ Manage cloud assets ■ Manage copy manager assets ■ Manage enclosure assets ■ Manage access profiles Manage user security settings Manage product settings Manage product updates Manage server deployments

Table 19-2 Predefined personas and jobs *(continued)*

Persona	Description and scope	Jobs
Resiliency Platform Deployment admin	Manage Resiliency Managers and Infrastructure Management Servers (IMSs). Can add an IMS to an existing data center. Manage product updates. Scope: Resiliency domain.	Manage product updates Manage server deployments
Data Center admin	Manage disaster recovery settings and manage assets of specified types. Scope: Specified data center.	Manage DR settings **Manage assets (jobs separated by type):** ■ Manage host assets ■ Manage virtualization assets ■ Manage data mover assets ■ Manage application cluster assets ■ Manage cloud assets ■ Manage copy manager assets ■ Manage enclosure assets ■ Manage access profiles
Resiliency Domain admin	Create, update, and delete resiliency groups, virtual business services (VBSs), and resiliency plans and templates. Start/stop all resiliency groups and VBSs. Scope: Resiliency domain.	Manage resiliency groups Start/stop resiliency groups Manage virtual business services Manage resiliency plan templates Manage resiliency plans Execute custom scripts

Table 19-2 Predefined personas and jobs *(continued)*

Persona	Description and scope	Jobs
Resiliency Group admin	Update and delete specified resiliency groups. Start/stop specified resiliency groups. Start/stop VBSs as long as the VBS contains only the specified resiliency groups. Scope: Specified resiliency groups.	Manage resiliency groups Start/stop resiliency groups
Resiliency Group operator	Start/stop specified resiliency groups. Start/stop VBSs as long as the VBS contains only the specified resiliency groups. Scope: Specified resiliency groups.	Start/stop resiliency groups
VBS admin	Create, update, and delete all virtual business services (VBSs). Start/stop all resiliency groups and VBSs. Scope: Resiliency domain.	Manage virtual business services Start/stop resiliency groups
Resiliency Domain Recovery admin	Configure all resiliency groups for disaster recovery (DR). ?? Perform rehearsal and DR operations: migrate, takeover. Create, update, and delete resiliency plans and templates. Manage disaster recovery network settings. Start/stop all resiliency groups. Scope: Resiliency domain.	Manage resiliency groups Rehearse resiliency groups Recover resiliency groups Manage resiliency plans Manage resiliency plan templates Manage DR settings Start/stop resiliency groups

Table 19-2 Predefined personas and jobs *(continued)*

Persona	Description and scope	Jobs
Resiliency Group Recovery admin	Manage and perform disaster recovery of resiliency groups Start/stop specified resiliency groups. Start/stop or perform DR operations on VBSs as long as the VBS contains only the specified resiliency groups. Scope: Specified resiliency groups.	Manage resiliency groups Start/stop resiliency groups Rehearse resiliency groups Recover resiliency groups
Resiliency Group Recovery operator	Start/stop specified resiliency groups. Perform disaster recovery on specified resiliency groups. Start/stop or perform DR operations on VBSs as long as the VBS contains only the specified resiliency groups. Scope: Specified resiliency groups.	Start/stop resiliency groups Perform disaster recovery of resiliency groups Rehearse resiliency groups
Guest	View all information in console. Assigned by default when user or group is configured for Resiliency Platform.	No operations, only view permission

Table 19-2 Predefined personas and jobs *(continued)*

Persona	Description and scope	Jobs
Resiliency Platform Assets admin	Manage all assets such as enclosure, application, application cluster assets, virtualization, data mover, and cloud.	**Manage assets (jobs separated by type):** ▪ Manage enclosure assets ▪ Manage application assets ▪ Manage virtualization assets ▪ Manage access profiles ▪ Manage cloud assets ▪ Manage application cluster assets ▪ Manage data mover assets

See "Managing user authentication and permissions" on page 175.

About limiting object scope for personas

For some personas, Veritas Resiliency Platform lets you select a subset of objects such as resiliency groups to limit the scope of operations.

See "Predefined personas" on page 178.

For example, you can assign one user the permission to perform operations on resiliency group RG1 and assign another user the permission to perform operations on RG2.

When planning persona assignments in which you select objects to limit the scope, take the following into account:

- Before you can select the objects such as resiliency groups to limit the scope of operations for a persona, the objects must first be created in Resiliency Platform.
- You need to plan for future maintenance on such limited scope personas. If more objects of that type are added later, you may need to edit existing personas for users or user groups in order to add permissions for the new objects.
- Keep in mind that operations on virtual business services (VBSs) that include multiple resiliency groups will fail unless the user performing the operation has permission for operations on all the resiliency groups in the VBS.

The same limitation applies for workflow or resiliency plan operations that include multiple resiliency groups.

For example: a VBS is composed of RG1 and RG2. The operator has permission to perform operations on RG1 but not RG2. If they try to perform operations on the VBS, the operation will fail.

Configuring authentication domains

By default, the Admin user on the Veritas Resiliency Platform virtual appliance can log in to the Resiliency Platform web console with access to all views and operations. The Admin user can configure authentication domains for Resiliency Platform from external identity providers so that other users can be authenticated for access to the console.

To configure authentication domains

To edit authentication domains

To configure authentication domains

1 Prerequisites

The fully qualified domain name (FQDN) or IP address and credentials for the LDAP/AD servers in the authentication domain

2 Navigate

⚙ **Settings** (menu bar)

Under **Product Settings**, click **User Management > Domains**

Note: You can also configure an authentication domain from the Getting Started wizard after setting up the Resiliency Manager and resiliency domain.

3 Click **Configure Domain**.

4 Select a data center and under **Specify server information for each data center**, enter the information for the server at that data center.

Repeat this step for other data centers in the authentication domain. When you select a different data center, the server information fields are cleared so that you can enter information for a different server, but the entries for the previous data center are remembered.

Note: If the same server is used for more than one data center, enter the same server information for each data center.

The remaining fields on the page apply to all data centers; fill these in as required.

See "Options for authentication domain configuration " on page 185.

Once you have entered information for all data centers, click **Next**.

5 Verify and complete the configuration:

In the **Domain Name** field, enter a friendly name for the authentication domain. If you configure the login screen to list domains, this name is listed.

Verify that the applicable data centers are listed. To make any changes, click **Back** to return to the previous screen. Once all is complete, click **Submit**.

6 Verify that the new domain is listed under **Domains**.

You can now configure user groups and users from that domain and assign permissions.

To edit authentication domains

1 Navigate to the domain list as described in the procedure to configure authentication domains.

2 Select the authentication domain you want to edit and select the Edit option.

Note the following guidelines when editing:

- To add server information for a new data center, select the applicable data center and fill in the server information.
- To edit existing server information, select the applicable data center.
- To edit other information, you do not need to select each data center; the same information applies to all.
- If a data center no longer uses a separate server, replace the server information for that data center with the information for the server that is being used.

- To remove a data center from the authentication domain, use the Unconfigure option instead of the Edit option.

  See "Unconfiguring authentication domains" on page 187.

See "Managing user authentication and permissions" on page 175.

Options for authentication domain configuration

The first page of the authentication domain configuration wizard is divided into 2 areas.

Server information by data center

Configuration options applicable to all data centers

Server information by data center

You must specify the server information separately for each data center. When you select a different data center the server information fields clear so you can enter the new information. If the same server is used for multiple data centers, enter the same information for both data centers.

Table 19-3 Server information by data center

Option	Description
Server (Mandatory)	Enter the fully-qualified host name or IP address of the LDAP server. If a secure session is configured with the LDAP server using SSL certificates, you must enter the fully-qualified host name that matches with the fully-qualified host name in the LDAP server certificate.
Port (Mandatory)	Displays the number of the port on which the LDAP server is configured to run. By default, this field displays the port number as 389. You can edit this port number, if required.
Connect using SSL/TLS	Select this check box to use the Secure Sockets Layer (SSL) or Transport Layer Security (TLS) certificates to establish a secure channel between the authentication broker and the LDAP server.
Certificate	Browse to the location of the trusted root CA certificate of the vendor that issued the LDAP server certificate.

Configuration options applicable to all data centers

The remaining fields apply to all data centers; fill these in as required.

Table 19-4 Configuration options applicable to all data centers

Option	Description
The authentication servers require me to log on.	Select this check box if the anonymous operations are disabled on the LDAP server and a bind user ID is required to proceed with configuring the LDAP-based authentication
Bind User Name/DN	Enter the complete Distinguished Name (DN) of the user that is used to bind to the LDAP server. If the LDAP server being used is Active Directory (AD), you can provide the DN in the following formats: username@domainname.com or domainname\username For example, you can provide the DN as Administrator@enterprise.domainname.com ENTERPRISE\Administrator For RFC 2307 compliant LDAP servers, specify complete bind DN. For example, cn=Manager,dc=vss,dc=veritas,dc=com The LDAP or the AD administrator can provide you the bind user name that you can use.
Password	Enter the password that is assigned to the bind user name that you use.
Query Information:	
User (Mandatory)	Under Query Information, enter the user name based on which the system detects the LDAP server-related settings. Ensure that the user name does not contain any special characters. The system determines the search base based on the user name that you specify in this field.
Group	Enter the name of the user group based on which the system detects the LDAP server-related settings. Ensure that the group name does not contain any special characters. The system determines the search base based on the group name along with the user name that you have specified.

See “Configuring authentication domains ” on page 183.

See “Getting started with a new Resiliency Platform configuration” on page 62.

Editing authentication domains

Using Resiliency Platform console, you can edit the configuration of an authentication domain.

To edit an authentication domain

1. Navigate

   ⚙ **Settings** (menu bar)

   Under **Product Settings**, click **User Management > Domains**

2. Right-click the domain and select **Edit**.
3. Edit the values that you want to update and click **Next**.
4. Verify the domain configuration details and click **Submit**.

See “Managing user authentication and permissions” on page 175.

Unconfiguring authentication domains

If an authentication domain is no longer applicable for a data center you can unconfigure it (remove it from Resiliency Platform).

Warning: Any users or user groups that you added from that domain are also removed from Resiliency Platform when you unconfigure an authentication domain.

To unconfigure an authentication domain

1. Navigate

   ⚙ **Settings** (menu bar)

   Under **Product Settings**, click **User Management > Domains**

2. Right-click the domain and select **Unconfigure**.
3. Select the data center. If you select all data centers, any users or user groups that you added from that domain are removed from Resiliency Platform. Click **Submit**.
4. Verify that the domain is removed under **Domains**.

See “Managing user authentication and permissions” on page 175.

Showing domains on login screen

You can set up the login screen to list the available authentication domains. By default, the domains list is not shown and the user must enter a fully qualified username, for example, username@domain or domain\username.

To show domains on login

1 Navigate

⚙ **Settings** (menu bar)

Under **Product Settings**, select **Miscellaneous**

2 Under **Login Settings**, select **Show domains list** and save the setting.

Configuring user groups and users

After you configure an authentication domain for Veritas Resiliency Platform, you can configure user groups and users for Resiliency Platform from that domain.

If you want to assign permissions to a user that are different from the user group as a whole, you must configure the user separately from the group.

To configure user groups and users

1 Prerequisites

The names of the user groups or users that you want to configure, as configured in the authentication domain.

2 Navigate

⚙ **Settings** (menu bar)

Under **Product Settings**, click **User Management > Users & Groups**

Note: To edit or remove an existing user or group, right-click the name in the list and select the appropriate option.

3 Click **Configure User or Group**.

4 Select the authentication domain.

5. Type the name of the user group or user. Click **Verify** so that the wizard can verify the name in the domain.

6. Click **Submit** and verify that the group or user is listed under **Users & Groups**.

   All groups and users that are added have the default persona of Guest. You can add other permissions.

   See “Assigning permissions to user groups and users” on page 189.

See “Managing user authentication and permissions” on page 175.

Assigning permissions to user groups and users

In Veritas Resiliency Platform, permissions use the concept of personas and jobs. When you first add user groups and users to Resiliency Platform, they are assigned the Guest persona, which allows views but no operations. You can assign other permissions. For each persona, there is a set of jobs (operations) and for some jobs, you select objects.

See “About user permissions in the web console” on page 176.

To assign permissions to user groups and users

1. Prerequisites

   The users and groups must be added to Resiliency Platform before you can assign personas.

   See “Configuring user groups and users” on page 188.'

2. Navigate

   ⚙ **Settings** (menu bar)

   Under **Product Settings**, click **User Management > Users & Groups**

3. Double-click the user group or user.

4. Click **Assign Persona**.

5. In the **Assign Persona** page, you can assign one persona at a time. Complete the following steps:

   - Select a persona that you want to assign to that user group or user.
   - Verify that you want to assign the jobs that are listed for that persona.
   - Under **Objects**, view the available objects on which jobs can be performed. To assign permission to selected objects, drag them from the left grid to

the left grid. If there are multiple object types, they are listed on separate tabs. Click any remaining tab and select the objects.

- Click **Submit**.

6 Verify that the correct persona name and associated objects are listed on the user details page.

To edit permissions or unassign personas

1 Navigate

⚙ **Settings** (menu bar)

Under **Product Settings**, click **User Management > Users & Groups**

2 Double-click the user or group.

3 On the details page for the user or group, right-click the persona that you want to unassign or edit, and select the appropriate option.

See “Managing user authentication and permissions” on page 175.

Adding custom personas

Veritas Resiliency Platform provides a set of predefined personas with access to predefined jobs.

You can add custom personas by selecting from the predefined jobs.

For example, the predefined persona Resiliency Platform Admin includes the jobs for managing assets, managing security settings, and managing product settings. You could create an "Asset Manager" persona that includes only the managing assets job.

You cannot customize the Super admin persona, which has access to all jobs and all objects in the resiliency domain. You also cannot customize the Guest persona, which can view all information in the console.

To add custom personas

1 Navigate

⚙ **Settings** (menu bar)

Under **Product Settings**, click **User Management > Persona & Jobs** > **New Persona**

2 In the **New Persona** page, complete the following steps and submit:

- Assign a name and description to the custom persona.
- Select one or more jobs that you want to assign to the persona. The jobs are shown in categories depending on whether the scope is the entire resiliency domain or whether the scope can be customized to specific data centers or assets. Select the job from the appropriate category.

  For example, if you want to assign a permission related to managing any resiliency group in the resiliency domain, select **Manage Resiliency Group** under the category of **For entire Resiliency Domain**. But if you want to limit permissions to specific resiliency groups, select **Manage Resiliency Group** under the category **For specific resiliency groups**.

  See "Predefined jobs that can be used for custom personas" on page 191.

3 Verify that the correct persona name and associated jobs are listed.

You can now assign this persona to users or user groups.

See "Managing user authentication and permissions" on page 175.

Predefined jobs that can be used for custom personas

The following table lists the predefined jobs that you can use to create custom personas for Veritas Resiliency Platform. The jobs are categorized as to whether they provide permissions for the entire resiliency domain or can be customized to specific data centers or assets.

Table 19-5 Jobs for custom personas

Jobs	Description	Scope
View all information	View all information in console.	Resiliency domain

Table 19-5 Jobs for custom personas *(continued)*

Jobs	Description	Scope
Manage user security settings	Manage authentication domains, users and user groups, personas.	Resiliency domain
Manage product settings	Manage general product settings such as alerts and notifications.	Resiliency domain
Manage server deployments	Edit Resiliency Manager information. Join a Resiliency Manager to a domain or leave a domain. Manage IMSs, including add, remove, edit, reconnect operations.	Resiliency domain
Manage product updates	Perform the operations available from the Product Updates page of the console.	Resiliency domain
Manage service objectives	Activate service objectives from templates; manage activated service objectives.	Resiliency domain
Manage assets, by type: ▪ Manage host assets ▪ Manage virtualization assets ▪ Manage data mover assets ▪ Manage application cluster assets ▪ Manage cloud assets ▪ Manage copy manager assets ▪ Manage enclosure assets ▪ Manage access profiles	Add, edit, or remove specific types of asset infrastructure	Resiliency domain or specific data centers
Manage resiliency groups	Create, update, and delete resiliency groups.	Resiliency domain or specific resiliency groups
Start/stop resiliency groups	Start and stop resiliency groups.	Resiliency domain or specific resiliency groups

Table 19-5 Jobs for custom personas *(continued)*

Jobs	Description	Scope
Manage virtual business services	Create, update, and delete virtual business services (VBSs).	Resiliency domain or specific VBSs
Manage resiliency plans	Create, update, and delete resiliency plans. **Note:** The permission to execute a resiliency plan depends on a cumulative check on permissions for individual resiliency groups and VBSs in the plan. See "About limiting object scope for personas" on page 182.	Resiliency domain
Manage resiliency plan templates	Create, update, and delete resiliency plan templates.	Resiliency domain
Execute custom scripts	Execute custom scripts as part of resiliency plans.	Resiliency domain or specific data centers
Rehearse resiliency groups	Perform rehearsal and rehearsal cleanup. **Note:** There is no separate job to perform rehearsal of VBSs. If the assigned scope of this job includes all the resiliency groups in a VBS, Rehearsal operations can be performed on that VBS. See "About limiting object scope for personas" on page 182.	Resiliency domain or specific resiliency groups

Table 19-5 Jobs for custom personas *(continued)*

Jobs	Description	Scope
Recover resiliency groups	Perform Recovery operations such as migrate, takeover, resync. **Note:** There is no separate job to perform disaster recovery of VBSs. If the assigned scope of this job includes all the resiliency groups in a VBS, DR operations can be performed on that VBS. See "About limiting object scope for personas" on page 182.	Resiliency domain or specific resiliency groups
Manage DR settings	Configure disaster recovery network settings, for example, mapping network settings for disaster recovery or replication gateway pairing.	Resiliency domain or specific data centers

See "Predefined personas" on page 178.

See "Adding custom personas" on page 190.

Configuring Windows global user

To customize the static IP of Windows guest virtual machines in the VMware environment, Resiliency Platform requires the administrator user name and password to log on to the Windows virtual machines. The user credentials can be Windows Active Directory user or Workgroup user.

For Windows Active Directory user, the Active Directory should be common for both, the primary and the recovery data center.

If a Windows virtual machine is part of a Windows Active Directory, ensure that you log on to the virtual machine at-least once using the Active Directory credentials.

For more information on customizing network, refer to the *Solutions Guide*.

To configure Windows global user

1. Navigate

   ⚙ **Settings** (menu bar)

   Under **Product Settings**, click **User Management > Windows Global User**

2. Click **+ Configure User** to configure the user.
3. Select between Active Directory and Workgroup.
4. Enter the administrator user name and password. Click **Verify**.

   For Workgroup user, enter user name as workgroupname\username. If the workgroup name is not customized then you can enter only the user name.
5. On successful verification, click **Next** and then **Finish** to submit the information.

Managing service objectives

This chapter includes the following topics:

Managing service objectives

As the name implies, Service Objective indicates the objective or intent with which the assets are being managed in a resiliency group.

See the following topics for information on viewing the predefined and pre-activated service objectives, customizing, and deleting service objectives.

Information on applying service objectives to assets is covered in the Solutions guides.

See "About service objectives" on page 196.

See "Viewing or customizing service objectives" on page 198.

See "Deleting activated service objectives" on page 199.

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.

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.

See “Viewing or customizing service objectives” on page 198.

See “Deleting activated service objectives” on page 199.

Viewing or customizing service objectives

Resiliency Platform provides a default set of pre-activated service objectives. It also provides a set of templates. You can view the activated and template service objectives in the console. You can customize a service objectives by activating a template and providing custom settings.

See “About service objectives” on page 196.

To view service objectives

1. Navigate

   ⚙ **Settings** (menu bar) > **Service Objectives**

2. View previously activated service objectives on the **Activated** tab or view the templates on the **Templates** tab. Expand a row to view details on settings.

To customize service objectives

1. Navigate

   ⚙ **Settings** (menu bar) > **Service Objectives**

2. On the **Templates** tab, locate the service objective and select **Activate**.
3. In the wizard, select from the options available to customize:
   - For remote recovery, you can customize the RPO setting to monitor an RPO. Select **Monitor as** and specify an RPO in minutes. Then select **Verify applicable technologies**.
     Choose between Copy and Replication technologies. Copy is available only if the acceptable RPO is 240 minutes (4 hours) and above.

If a technology that you select requires additional information, the wizard displays a page for you to provide those inputs.

- On the Summary page, provide a name for the customized service objective. You can also select a **Rank**.

The customized service objectives is listed on the **Activated** tab.

Deleting activated service objectives

If you do not want to use a pre-activated service objectives or one that was customized and activated from a template, you can delete it- provided that it is not in use by a resiliency group. Deleting the service objective ensures that it is not listed for selection when creating a new resiliency group.

To delete an activated service objective

1 Navigate

⚙ **Settings** (menu bar) > **Service Objectives**

2 ⋮ On the **Activated** tab, select the vertical ellipsis next to the service objective that you want to delete and select **Delete**.

See “About service objectives” on page 196.

Managing reports

This chapter includes the following topics:

Managing reports

Using the Veritas Resiliency Platform console you can view and generate various reports. You can schedule periodic email updates.

See “About reports” on page 200.

See “Managing report preferences” on page 201.

See “Scheduling a report” on page 203.

See “Running a report” on page 205.

See “Viewing and managing report schedules” on page 206.

About reports

Using the Veritas Resiliency Platform console, you can generate a variety of reports. The following are the broad categories under which the reports are grouped:

- **Inventory**: Reports in this category provide information about the data centers and applications, and the virtual machines that are deployed in the data centers.
- **Recovery Assessment**: This category lists the reports that are related to the disaster recovery operations such as the migrate and take over report, and the rehearsal report.
- **Risk**: This category has two reports; Current Risk and Risk History. These reports show the summary and details of all the current and historical risks that occurred in the environment.

Reports can be scoped on the data center or global. You can subscribe for a report on a daily, weekly, monthly, quarterly, or yearly basis, or on predefined days of the week, or run on demand. Reports are available in the HTML and PDF format, or as a comma-separated file (CSV) file.

You can send a report to multiple recipients by entering the email addresses separated by a comma (,) or a semicolon (;).

See "Managing report preferences" on page 201.

See "Scheduling a report" on page 203.

See "Running a report" on page 205.

Managing report preferences

Using the Veritas Resiliency Platform console, you can create, update, and view your preferences for generating and receiving reports.

To create report preferences

1. Navigate

   Reports (menu bar) > **Settings**.

2. In the **Report preferences** wizard panel, specify the following information and click **Save**.

Scope	Select the scope of the report such as Global or specific data center.
Duration	Select the duration for which you want to receive the report.
Format	Select the delivery format as HTML or CSV.

Email	Enter an email address at which you want to send the report. You can enter multiple email addresses that are separated by a comma (,) or a semicolon (;).
Frequency	Select the start and the end date and the time at which you want to generate and receive the report. Select **Daily** to generate the report on a daily basis. Select **Weekly** to avail the following options: ■ Select **Every Weekday** to receive the report on all week days. ■ Select **Recur every week on** and select one or more week days on which you want to receive the report. Select **Monthly** to avail the following options: ■ Set the monthly recurrence. For example every one month, or every 3 months. ■ Select the day of the month on which you want to receive the report. ■ Or select every weekday of the month on which you want to receive the report. For example every first Monday of the month or every fourth Saturday of the month. Select **Yearly** to avail the following options: ■ Set the yearly recurrence. For example every one year, or every 3 years. ■ Select the day of the month on which you want to receive the report. ■ Or select every weekday of a month on which you want to receive the report. For example every first Monday of January or every fourth Saturday of April. Select **Once** to generate the report only one time.

See "Scheduling a report" on page 203.

See "Running a report" on page 205.

Scheduling a report

Using the Veritas Resiliency Platform console, you can update the report generation schedule for a selected report. The schedule that is defined in the template is then overwritten. You can also enable or disable the report schedule.

To schedule a report

1 Navigate

Reports (menu bar), and expand the report category.

2 In the report row, click on **Schedule**.

3 In the **Schedule Report** wizard panel, specify the following information and click **Schedule**.

4

Name	Enter a name for the report schedule. Only special character under score (_) is allowed.
Description	Enter a description for the report schedule.

Frequency	Select the start and the end date and the time at which you want to generate and receive the report. Select **Daily** to generate the report on a daily basis. Select **Weekly** to avail the following options: ■ Select **Every Weekday** to receive the report on all week days. ■ Select **Recur every week on** and select one or more week days on which you want to receive the report. Select **Monthly** to avail the following options: ■ Set the monthly recurrence. For example every one month, or every 3 months. ■ Select the day of the month on which you want to receive the report. ■ Or select every weekday of the month on which you want to receive the report. For example every first Monday of the month or every fourth Saturday of the month. Select **Yearly** to avail the following options: ■ Set the yearly recurrence. For example every one year, or every 3 years. ■ Select the day of the month on which you want to receive the report. ■ Or select every weekday of a month on which you want to receive the report. For example every first Monday of January or every fourth Saturday of April. Select **Once** to generate the report only one time.
Scope	Select the scope of the report such as Global or specific data center.
From and To	Select the duration for which you want to generate the report.

Format	Select the delivery format as HTML or CSV.
Email	Enter an email address at which you want to send the report. You can enter multiple email addresses that are separated by a comma (,) or semicolon (;).

See “Managing report preferences” on page 201.

See “Running a report” on page 205.

Running a report

On the Veritas Resiliency Platform console, you can run a report on demand. The report is generated and sent to the specified email address. To view the generated report in the browser, do one of the following:

- Click on the report notification.
- Click **Saved** to expand the table, and then double-click on the saved report row.
- Click **Saved** to expand the table, click on the **Action** menu, and then click **View**.

To run a report

1 Navigate

Reports (menu bar).

Click **Inventory Reports** or **Risk Assessment Reports**.

2 Click **Run** on the desired report, specify the following information in the wizard panel, and click **Run**.

Scope	Select the scope of the report such as Global or specific data center.
From and To	Select the duration for which you want to generate the report.
Format	Select the delivery format as HTML or CSV.
Email	Enter an email address at which you want to send the report. You can enter multiple email addresses that are separated by a comma (,) or semicolon (;).

See “Scheduling a report” on page 203.

Viewing and managing report schedules

You can use the Resiliency Platform console to view the details of all the reports and manage the report schedules. You can view a brief description about the report along with the following information:

- Number of times the report is saved.
- Number of times the report is scheduled to run.
- Number of currently running instances of the report.

When a currently running instance of a report is complete, the number of saved report count increases by one and the number of currently running instances count decreases by one.

In each report row you can do the following:

Saved	Click to view additional details such as the generation time, format, status, scope, and user information of all the saved instances of the report. Double-click on a saved report row to view the report. Click on the vertical ellipses to view or remove the report. Saved reports are purged depending on the number of days defined in the **Reports Retention Policy Settings**.
Schedules	Click to view the report generation schedules such as the format, recipient email address, recurrence, and whether the report is enabled or disabled. Click on the **Actions** column to enable, disable, update, or delete the report schedule.
Running	Click to view the format, scope, and user information. You can abort the generation process.
Run	Click to run the report on demand.
Schedule	Click to update the report generation schedule.
Last Run	Displays the last run date and time of the report.

To view reports

- Navigate

  Reports (menu bar)

  Expand **Inventory Reports** or **Risk Assessment Reports** to view details of all the reports.

See “Managing report preferences” on page 201.

See “Scheduling a report” on page 203.

See “Running a report” on page 205.

Chapter 22

Managing settings

This chapter includes the following topics:

Managing settings for alerts and notifications and miscellaneous product settings

See the following topics for information on configuring email and SNMP settings for notifications and reports, setting up rules for event notifications, configuring purge intervals, and changing telemetry settings.

See “Adding, modifying, or deleting email settings” on page 209.

See “Adding, modifying, or deleting SNMP settings” on page 211.

See “Setting up rules for event notifications” on page 211.

See “Modifying the purge setting for logs and SNMP traps” on page 212.

See “Modifying the purge setting for reports” on page 213.

See “Modifying the purge setting for activities” on page 213.

See “Enabling or disabling telemetry collection ” on page 214.

Adding, modifying, or deleting email settings

You can configure email settings to be used for different features, such as sending reports or receiving automatic email notifications of events. Veritas Resiliency Platform manages email notifications via Resiliency Managers. When Resiliency Managers are located in different geographical locations, the required email settings are likely different for each location. In that case, you add a separate email configuration for each location. You can send a test email to verify the settings. You can also modify or delete existing email configurations.

To add, modify, or delete email settings

1 Navigate

⚙ **Settings** (menu bar)

Under **Product Settings**, select **Alerts & Notifications** > **Email**

To add a new email configuration, select **Add Email Configuration**.

To modify or delete an existing one, right-click it and select **Modify** or **Delete.**

2 To add or modify an email configuration, go through the wizard pages and specify the options.

In **Server Information**, specify the following:

Name	Assign a unique name for the email configuration.
Email Server	Valid formats include: Fully Qualified Domain Name (FQDN), IP address, or, if the network handles DNS resolution for host names, a shortened host name. Examples: Host123, Host123.example.com, xxx.yyy.zzz.aaa.
SMTP Port	Enter the SMTP mail server port number. The default is 25.
From Email Address	Enter the email address to be shown as the sender of all the emails that are sent.
Friendly Email Name	Optionally, enter a name to be shown for the From address.
Send To	Enter the email address to which you want to send the email.

3 In **Security**, if you want to implement secure SMTP, select the checkbox and enter the user name and password.

4 In **Select Resiliency Managers**, select a Resiliency Manager in the data center location where these email settings apply.

5 In **Test Email Settings**, enter a valid email address, and enter a subject and message for the test email. Select **Send Test Email** to test your settings.

6 Review the information in the summary and submit

See “Managing settings for alerts and notifications and miscellaneous product settings” on page 208.

Adding, modifying, or deleting SNMP settings

When an event takes place, you can configure SNMP traps to be sent. You can configure the SNMP settings in the web console.

To add, modify, or delete SNMP settings

1 Navigate

⚙ **Settings** (menu bar)

Under **Product Settings**, select **Alerts & Notifications** > **SNMP**

To add a new SNMP configuration, select **Add SNMP Configuration**.

To modify or delete an existing one, right-click it and select **Modify** or **Delete.**

2 To add or modify SNMP settings, specify the following:

Name	Assign a friendly name.
SNMP Server	Enter the IP Address or name of the host where the SNMP trap console is located. Example: Host123.example.com
SNMP Port	Enter the SNMP port number. The default port for the trap is 162.

See “Managing settings for alerts and notifications and miscellaneous product settings” on page 208.

Setting up rules for event notifications

Logs of the type information, warning, or error generate an event. You can view Veritas Resiliency Platform event logs in the web console and set up rules for receiving notifications of events. You can also modify or delete existing rules.

To set up rules for event notifications

1 Prerequisite

Configure the email server for sending notifications. Optionally you can also configure SNMP.

See “Adding, modifying, or deleting email settings” on page 209.

See “Adding, modifying, or deleting SNMP settings” on page 211.

2 Navigate

⚙ **Settings** (menu bar)

Under **Product Settings**, select **Alerts & Notifications**

To add a new rule: Select the **Definition** tab > **New Rule**.

To modify or delete an existing rule: Select the **Rules** tab, right-click the rule, and select **Modify** or **Delete.**

3 In **Configure Rule**, enter or modify the following:

Name	Enter a unique name for this rule.
Send emails to	Enter one or more email addresses separated by a comma
Send SNMP traps to	Optional
Select Notifications	Select one or more events that you want to be notified about

4 Select **Submit**.

The rule is listed on the **Rules** tab.

Modifying the purge setting for logs and SNMP traps

By default, logs and SNMP traps are retained for two years. You can modify this purge setting.

To modify the purge setting for logs and SNMP traps

1 Navigate

⚙ **Settings** (menu bar)

Under **Product Settings**, click **Miscellaneous**

2 Under **Log Settings**, enter the new value for the purge setting, in months, and save the setting.

See “Managing settings for alerts and notifications and miscellaneous product settings” on page 208.

Modifying the purge setting for reports

By default, reports are saved for 7 days. You can modify this purge interval.

To modify the purge setting for reports

1 Navigate

⚙ **Settings** (menu bar)

Under **Product Settings**, click **Miscellaneous**

2 Under **Reports Retention Policy Settings**, enter the new value for the purge setting and save the setting.

See “Managing settings for alerts and notifications and miscellaneous product settings” on page 208.

Modifying the purge setting for activities

By default, the information on tasks performed in activities is saved for 6 months. You can modify this purge setting.

To modify the purge setting for activities

1 Navigate

⚙ **Settings** (menu bar)

Under **Product Settings**, click **Miscellaneous**

2 Under **Activities Settings**, enter the new value for the purge setting and save the setting.

See "Managing settings for alerts and notifications and miscellaneous product settings" on page 208.

Enabling or disabling telemetry collection

Veritas Resiliency Platform can collect usage information via telemetry for the purpose of future product enhancements. You can enable or disable the collection.

The types of telemetry information collected include configuration information, mainly inventory counts, and license information.

For example, information can include number of configured authentication domains, resiliency plans and templates, virtual business services, virtual machines by platform and virtualization technology, virtualization servers by type, resiliency groups by replication type, distribution of hosts over physical and virtual, enclosures by type, virtual machines and applications enabled or not enabled for disaster recovery.

You can view a file showing the collected information.

Telemetry collection requires that the Resiliency Manager have internet connectivity.

To enable or disable telemetry collection

1 Navigate

⚙ **Settings** (menu bar)

Under **Product Settings**, select **Miscellaneous**

2 Under **Telemetry Settings**, select the setting to turn it on or off and save the setting. To download a file showing the information that is collected, select **Show what is collected**.

Section 6

Updating or uninstalling the product

Chapter 23

Updating Resiliency Platform

This chapter includes the following topics:

About updating Resiliency Platform

This chapter covers common aspects of updating a Resiliency Platform deployment.

The topics in this chapter cover the process of applying updates (patches and maintenance release) to the virtual appliance, add-ons, and host packages.

Note: Upgrade from Resiliency Platform 2.0 to Resiliency Platform 2.1 using the Resiliency Manager console is not supported. This upgrade can be done only through klish menu.

About applying updates to Resiliency Platform

Updates to Veritas Resiliency Platform provide significant benefits, such as improved functionality, performance, security, and reliability.

In Veritas Resiliency Platform, you can apply updates to the following:

- Veritas Resiliency Platform virtual appliance
- Veritas Resiliency Platform add-ons
- Host packages on the assets that are added to the Infrastructure Management Server (IMS) as a host

---

Note: It is recommended to apply the update to all the Resiliency Platform components to take the complete advantage of the changes available in the updates.

---

you can apply the update either across all the appliances in your environment through Resiliency Manager console, or through individual appliance's klish menu. Refer to the readme file shipped with the update to check if the update should be applied using the console or using klish menu.

For applying updates to Resiliency Platform, you need to set up a repository server and download the updates to the repository server. Then, you assign the repository server to the Resiliency Platform virtual appliance, where you want to apply the update.

The following figure shows how a repository server is used to apply the updates to Resiliency Platform:

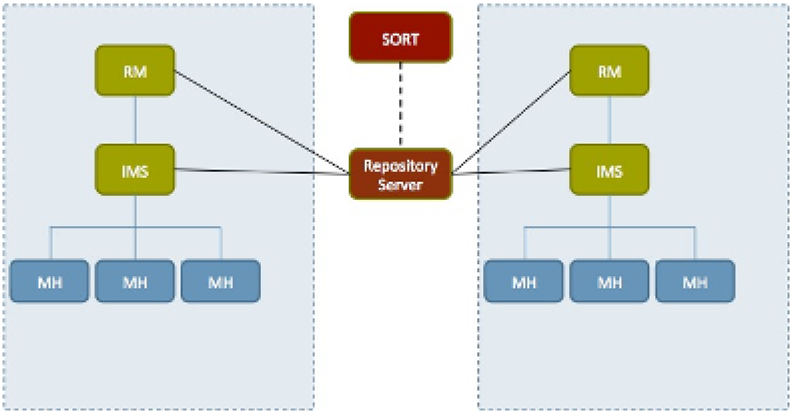

Note: While applying updates, ensure that the virtual appliance remains powered on. Restarting the appliance during the process of applying updates may adversely affect the functionality and the virtual appliance may go into an irrecoverable condition.

The following is an overview of the process of applying updates in Veritas Resiliency Platform:

Table 23-1 Applying updates to Resiliency Platform

Step	Task	Description
1	Make sure that the prerequisites for the repository server are met.	See "Prerequisites for a repository server" on page 220.
2	Download the update	See "Downloading the Resiliency Platform update" on page 221.
3	Set up a repository server You can also use the YUM virtual appliance as the YUM server	See "Setting up the repository server " on page 224. See "Using YUM virtual appliance as YUM server" on page 221.

Table 23-1 Applying updates to Resiliency Platform *(continued)*

Step	Task	Description
4	Apply update to the virtual appliances in the given sequence: Resiliency Manager. In case of two Resiliency Managers in the domain, both the Resiliency Managers in synchronization IMS at production data center IMS at recovery data center Replication Gateway at recovery data center. After applying the update on the Replication Gateway, the add-ons also need to be updated. Replication Gateway at production data center. After applying the update on the Replication Gateway, the add-ons also need to be updated. Reboot the appliances after upgrade.	Refer to the readme file shipped with the update to check if the update should be applied using the console or using klish menu See "Applying updates to virtual appliance using klish menu" on page 227. See "Applying updates to virtual appliances using the console" on page 226. For applying update on Resiliency Manager, see: See "Applying update on Resiliency Managers" on page 229. For updating the add-ons on Replication Gateway, see: See "Applying updates to the hosts" on page 231.
5	Apply update on the host packages	See "Applying updates to the hosts" on page 231.
6	In case of infoScale environment, apply update to the add-on	See "Applying update for InfoScale environment" on page 230.
7	If you had configured a resiliency group using Resiliency Platform Data Mover before applying update to IMS, apply update to the Data Mover bundle	See "Applying updates to Resiliency Platform Data Mover bundle" on page 231.
8	Remove an update from the repository server	See "Removing an update from the repository server" on page 232.

You also have an option of applying a private hotfix, if veritas support provides you one.

Note: All the updates should be downloaded and stored in a common location.

Prerequisites for a repository server

To set up a repository server, make sure that the following prerequisites are met:

- Repository server should be RHEL server version 6.5 with minimum YUM version 3.2.29. Base server installation is recommended for the repository server.
- Web server (HTTP/HTTPS) should be configured on the server. Two-way SSL configuration is recommended for HTTPS.
  Default ports are 80 for HTTP and 443 for HTTPS.
- Repository server should have minimum 50 GB disk space available for repository data.
- Repository server should have connectivity with SORT as well as with the virtual appliances.
- `createrepo` should be installed on the server.
- Perl and Python should be installed on the server. The following modules need to be installed:
  - `Archive::Extract`
  - `Archive::Tar`
  - `Config::Simple`
  - `Cwd`
  - `File::Basename`
  - `File::Copy`
  - `File::Fetch`
  - `File::Path`
  - `Getopt::Long`
  - `JSON`
  - `LWP::Simple`
  - `Time::Local`
  - `XML::Twig`

See “About applying updates to Resiliency Platform” on page 217.

Downloading the Resiliency Platform update

Updates to Veritas Resiliency Platform are made available to the customers on SORT.

To download the update

1. Go to the following location

   https://sort.veritas.com/patch/finder?prod=vrp

2. You can see a list of all the applicable updates for a particular version. Select the required version and update, and then click the **Download** link.

3. Unzip the file that you have downloaded. Following files are extracted to the directory for updates:

   - `Veritas_Resiliency_Platform_Upgrade_Bundle_`*`version_of_the_update`*`.tar.gz`
   - `Veritas_Resiliency_Platform_DataMover_Upgrade_Bundle_`*`version_of_the_update`*`.tar.gz`

   Where, *version_of_the_update* is the version of the update.

See “About applying updates to Resiliency Platform” on page 217.

Using YUM virtual appliance as YUM server

Veritas Resiliency Platform 2.1 introduces the use of a virtual appliance as a YUM server. This feature of Resiliency Platform makes it easier to setup a YUM server. If you do not use the YUM virtual appliance, you will be required to manually install all the required perl modules to setup the repository server.

The Resiliency Platform YUM virtual appliance is available in the OVA and VHD formats.

To use the YUM appliance as YUM server

1. Download the YUM appliance, deploy, and configure it.

   See “Downloading the Resiliency Platform virtual appliances” on page 25.

   See “Deploying the virtual appliance through VMware vSphere Client” on page 43.

   See “Deploying the virtual appliance through Hyper-V Manager” on page 44.

2. After successful deployment of the Resiliency Platform YUM appliance, log into the appliance using admin user. Go to `manage > option`.

3 Add the patch using the `add-patch` klish option. This option requires you to enter the upgrade bundle location and the setup repository bundle location. These locations could be any location which is accessible by the `wget` command.

4 You can use the following commands to list the updates and remove them:

Command	Purpose
`list-patch vrp`	To list the Veritas Resiliency Platform patches. The command also provides the information about the patch that is to be used for configuring the YUM repository on Resiliency Manager, IMS, or Replication Gateway appliances such as the port, repository path, and the protocol.
`remove-patch vrp all`	To remove the Veritas Resiliency Platform patches.
`list-patch datamover`	To list the Resiliency Platform Data Mover patches.
`remove-patch datamover all`	To remove the Resiliency Platform Data Mover patches.

Note: Upgrade is not supported for a YUM virtual appliance.

See “About applying updates to Resiliency Platform” on page 217.

Configuring the YUM repository server

After the virtual appliance deployment, when you log into the virtual appliance console for the first time using the admin user credentials, the bootstrap process is automatically invoked. This bootstrap process is used to set up or configure the Resiliency Platform component for the first time.

The default network protocol for virtual appliance is Dynamic Host Configuration Protocol (DHCP). If the appliance detects DHCP during the first boot or before the completion of bootstrap process, the appliance network automatically gets configured. After the network configuration, you can either use the virtual appliance console or Secure Shell (SSH) to log in and complete the bootstrap process.

If DHCP is not configured in your environment, you have an option to use a static IP for the appliance. Since the appliance network is not automatically configured in this case, you can only use the console to log into the virtual appliance.

To configure a YUM repository server

1. Log in to the virtual appliance console or SSH using the following credentials:

   - **Username:** admin
   - **Password:** P@ssw0rd

   After a successful login, you need to change the password of the admin user.

   See "Password policies for Resiliency Platform virtual appliance" on page 57.

   If you are logged in to SSH, you will be logged off the SSH session after the password change and you need to again log in to complete the rest of the steps of the bootstrap process. If you are logged in to the virtual appliance console, you can continue and complete the rest of the steps of the bootstrap process.

2. Accept the End User License agreement (EULA) to proceed with the configuration.

3. In the **Host Network Settings** section, you can configure the appliance network by using DHCP or static IP.

   If the DHCP configuration is working in the environment, the details are printed on the screen. Confirm if you want to proceed with these settings.

   In case of DHCP, you need to ensure that a Dynamic Host Configuration Protocol (DHCP) server is working in the subnet where the virtual appliance is deployed. In case of static IP, you need to respond to the following additional prompts:

   - **Enter the fully qualified hostname:**
   - **Enter the IP address:**
   - **Enter the Subnet mask:**
   - **Enter the Default Gateway:**
   - **Enter the DNS server (space separated if more than one DNS, maximum 2 DNS entries):**

4. Confirm if you are in Network Address Translation (NAT) environment and want to verify the NAT configuration using an external IP and hostname. You need to respond to the following additional prompts:

   - **Enter the fully qualified external hostname:**
   - **Enter the external IP address:**

Note: The external hostname or IP that you provide are set as the appliance hostname or IP.

5 In the **Appliance Settings** section, do the following:

- Press the Enter key to confirm the use of an NTP server for configuring the date and time.
- You are required to select the time zone. Follow the instructions as displayed on the virtual appliance console or SSH session to select the correct time zone.
- Enter the FQDN or IP address of the NTP server. The appliance verifies the NTP server details. If there are any issues, details are displayed on the screen and you are prompted to enter the details again.

You can reset the timezone and NTP server at a later point of time using klish menu. Changing the system settings can affect the product functionality if incorrect values are set.

6 **Product Settings**: This section configures the appliance as a repository server.

7 After a successful product configuration, a confirmation message will be displayed and you will be logged out of the virtual appliance console.

See “About applying updates to Resiliency Platform” on page 217.

Setting up the repository server

You need to set up a repository server in your environment, download the updates from SORT, and make them available on your repository server.

To set up a repository server

1 Go to the following location and select the product and version:

https://sort.veritas.com/patch

2 You can see a list of all the applicable updates for a particular version. Select the required version and click the required update. On the next page, download the file by clicking the **Setup Repository Bundle Download** link.

3 Copy the file that you have downloaded to a temporary location and extract this tar file.

4 Create a repository path under root directory of the web server.

```
mkdir path_to_repository
```

5 The `setup_conf_repo.pl` file is one of the files that are extracted from the update that is downloaded from SORT. This file is used to configure the repository.

6 To update the repository server with the updates that you have saved on your local system:

```
./setup_conf_repo.pl --add-local-updates --repo-location
path_to_repository --update-location path_to_tar
--metadata-location path_to_master.xml
```

See “About applying updates to Resiliency Platform” on page 217.

Adding a repository server in Resiliency Platform

After configuring a repository server, you need to add the repository server in Veritas Resiliency Platform. There can be multiple repository servers added to Veritas Resiliency Platform at a time.

To add a repository server in Veritas Resiliency Platform

1 Navigate

⚙ **Settings** (menu bar) > **Updates** > **Repository Servers**

2 Click **Add**.

3 In the **Add Repository** Wizard panel, do the following:

- Select the protocol for adding the repository server.
- Enter the fully qualified hostname (FQDN) or IP address of the server that you want to configure as the repository server.
- If you want to modify the default port, enter the port number.
- Enter the repository path that is created under root directory of web server.
- Click **Submit**.

See “About applying updates to Resiliency Platform” on page 217.

Assigning a repository server in Resiliency Platform

You need to assign a repository server to every virtual appliance where you want to apply the updates. You can store all the available updates on this server and apply it on the virtual appliance whenever required.

A single repository server can be assigned to multiple virtual appliances but one virtual appliance can be assigned only one repository server at a time.

To assign a repository server to a virtual appliance

1 Navigate

⚙ **Settings** (menu bar) > **Updates**

2 Select the server names (virtual appliances) to which you want to assign a repository server.

3 Click **Assign Repository**. Select the repository server that you want to assign to the virtual appliances.

Click **Submit**.

See “About applying updates to Resiliency Platform” on page 217.

Applying updates to virtual appliances using the console

You can apply updates to the virtual appliances using the console.

Before applying the update using the console, you need to first add a repository server to Resiliency Platform and then assign a repository server to the virtual appliance.

See “Adding a repository server in Resiliency Platform” on page 225.

See “Assigning a repository server in Resiliency Platform” on page 226.

Replication Gateway updates must be applied on the recovery site Replication Gateway first and then on the production site Gateway.

To apply updates to the virtual appliances using the console

1 Prerequisites:

All the Resiliency Managers in the domain should have same version of update installed on them.

Ensure that following services are running on the local as well as remote Resiliency Manager:

- User Interface service
- Database service
- Messaging service
- Core service
- Task service
- Event service

2. Navigate

   ⚙ **Settings** (menu bar) > **Updates**

3. Select the server name or virtual appliance on which you want to apply the update.
4. Select the update that you want to apply from **New Updates**.
5. Click **Upgrade**.
6. Verify the details of the update and click **Submit**.

Note: If the process of applying updates on the appliance takes more than 30 minutes, the session times out and you need to confirm if you want to continue the session and refresh the page. The progress of the task of applying updates can be tracked from **Recent Activities**.

See “About applying updates to Resiliency Platform” on page 217.

Applying updates to virtual appliance using klish menu

You can use the klish menu to perform the upgrade related tasks in Resiliency Platform.

Replication gateway updates must be applied on the cloud replication gateway first and then on the on-premises gateway.

To apply updates to virtual appliance using klish menu

1. Prerequisites

   Before applying the update, you must stop the environmental services in the given sequence using the following klish commands:

   - Stop all the IMS services at production data center through klish:

     ```
     manage>services ims stop ALL
     ```

   - Stop all the Resiliency Manager services at production data center through klish:

     ```
     manage>services rm stop ALL
     ```

   - Stop all the IMS services at recovery data center through klish:

     ```
     manage>services ims stop ALL
     ```

   - Stop all the Resiliency Manager services at recovery data center through klish:

     ```
     manage>services rm stop ALL
     ```

2. It is recommended to power off the appliance and take a snapshot of the appliance before applying the updates.

3. You need to log into the virtual appliance as admin and go to the updates sub-menu.

4. Following is a list of commands that you can run to perform the operations that are related to the updates:

   - To configure the repository:

     `config-repository` *`FQDN_or_IP_of_the _repository_server protocol port_number Repository _path_on_repository_server`*

     If you enter HTTPS as protocol, you are required to copy the content from the SSL certificate, paste it on prompt, and press enter key.

   - To view the current configuration of the repository:

     `show-repository`

   - To view the current version of the appliance or the version of the update installed on the appliance:

     `list-updates`

   - To show the readme file for the specified update:

```
show-readme version_of_the_update
```

- To apply the specified update:

  ```
  apply-update version_of_the_update
  ```

- To remove the current repository configuration:

  ```
  remove_repository
  ```

5. After applying updates, you may want to refresh the information about the applicable updates on each of the virtual appliances or servers. If you apply the updates using klish, you need to refresh the information to reflect the current status of the updates in the Resiliency Manager web console.
6. Navigate

   ⚙ **Settings** (menu bar) > **Updates** > **Available Updates**

   Click **Refresh**.

See “Using klish” on page 256. for a complete list of options available with `Updates` command.

See “About applying updates to Resiliency Platform” on page 217.

Applying update on Resiliency Managers

If you have two Resiliency Managers in your resiliency domain then you need to apply the update on both the Resiliency Managers in synchronization.

For the purpose of upgrade, one of the Resiliency Managers is designated as import node and the other one as non-import node:

- The first Resiliency Manager configured in the resiliency domain is called the import node. If the first Resiliency Manager is no longer in the resiliency domain, then the Resiliency Manager configured next in the domain is designated as the import node.
- The other Resiliency Manager in the domain is called the non-import node.

If you do not remember the sequence in which the Resiliency Managers are configured in the domain, you can start applying the update on any one of the Resiliency Managers and the process will guide you about the import node and non-import node.

To apply update on Resiliency Managers

1. Start the process of applying update on the non-import node. Wait until upgrade on the non-import node reaches to a stage where you are prompted to switch to the import node.

2. Start the process of applying update on the import node. Once the upgrade process completes successfully, verify that all the Resiliency Manager services are running on the import node by running following klish command on the import node:

   ```
   manage> services rm show
   ```

   If the upgrade on the import node fails, you need to stop all the Resiliency Manager services to avoid the data loss by running following klish command on the import node:

   ```
   manage> services rm stop ALL
   ```

3. Now switch to the non-import node and resume the upgrade there.

   In case the upgrade on import node had failed, you need to take the non-import node offline as soon as the upgrade on the non-import node completes, irrespective of success or failure of the upgrade.

Applying update for InfoScale environment

If a Veritas InfoScale Operations Manager server is added to Resiliency Manager, you need to apply update on Veritas Resiliency Platform Enablement add-on (VRTSsfmitrp) for Veritas InfoScale Operations Manager after you apply the update on the resiliency Manager.

To apply update for InfoScale environment

1. Go to the following location

   https://sort.veritas.com/vom

2. You can see a list of all the applicable add-ons for a particular Veritas InfoScale Operations Manager version. Select the required version and download the `Veritas Resiliency Platform Enablement` add-on.

3. Log into Veritas InfoScale Operations Manager server console. Go to **Settings > Deployment > Upload Solutions** and upload the .sfa file that you have downloaded and extracted within the patches folder.

4. Right click and install the .sfa file on Veritas InfoScale Operations Manager server.

5. Restart the Web server by clicking **Restart Web Server** on the task bar.

6. Go to **Settings > Deployment** in Veritas InfoScale Operations Manager server console and install the .sfa file on the required hosts.
7. Log into the Resiliency Manager console. Remove the Veritas InfoScale Operations Manager servers from production as well as recovery data center once and then add them again.
8. Edit the resiliency groups that were created before applying the update. You are prompted to include all the service groups belonging to a single group dependency.

See "About applying updates to Resiliency Platform" on page 217.

Applying updates to Resiliency Platform Data Mover bundle

If any resiliency group has already been configured using Resiliency Platform Data Mover before applying update to IMS, you need to apply update to Resiliency Platform Data Mover bundle.

To apply updates to Resiliency Platform Data Mover bundle

1. Use the following command of klish:

   ```
   manage>datamover vmware-iofilter upgrade
   ```

2. The command will display a list of applicable hosts. Enter the name of the host where you want to upgrade the Resiliency Platform Data Mover bundle.

Note: If automatic DRS is not enabled, you need to put the ESX hosts into maintenance mode to go ahead with the bundle update.

See "About applying updates to Resiliency Platform" on page 217.

Applying updates to the hosts

Updates for the add-ons and for the host packages installed on the assets that are added as a host, becomes available under the **Managed Hosts** section in the Resiliency Platform console. These components can be upgraded from the console.

To apply updates to the hosts

1. Navigate

   ⚙ **Settings** (menu bar) > **Updates**

2. Under **Available updates**, go to **Managed Hosts** section. Select the hosts on which you want to apply the update, and click **Upgrade**.

In the list of hosts, some of the hosts may be listed for both the production as well as recovery data centers. You need to apply the update only on a host which is listed under production data center.

Removing an update from the repository server

You can remove a particular update from the repository server.

To remove an update from the repository server

1. Go to the `ITRP/RM` directory on the repository server. This directory is created under the repository path that you had provided while setting up the repository.
2. Run the following commands:
   - To remove the directory created for a particular update:

     ```
     rm -rf patch_version_dir
     ```

   - To clear the older data, and then refresh and build the repository with the existing patches in the `RM` directory:

     ```
     createrepo --update RM
     ```

See “About applying updates to Resiliency Platform” on page 217.

Chapter 24

Uninstalling Resiliency Platform

This chapter includes the following topics:

- About uninstalling Resiliency Platform

About uninstalling Resiliency Platform

In the current version, there is no provision for uninstalling Resiliency Platform. If you do not want to use the Resiliency Platform product any longer, you can remove the Resiliency Platform virtual appliance node using the appropriate hypervisor manager in your environment.

If you want to decommission a Resiliency Platform virtual appliance node while continuing to use the product on other nodes in the resiliency domain, you should first use the web console to remove the node from the Resiliency Manager database. For example, you can remove a Resiliency Manager node from the domain if another Resiliency Manager node is active.

See "Removing a Resiliency Manager from a resiliency domain" on page 66.

If you want to remove an Infrastructure Management Server (IMS), you first need to remove the association of the IMS with the resiliency Manager before removing the virtual appliance node:

See "Removing an IMS" on page 72.

Section 7

Troubleshooting and maintenance

Chapter 25

Troubleshooting and maintenance

This chapter includes the following topics:

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.

See "Setting up rules for event notifications" on page 211.

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

See "Modifying the purge setting for logs and SNMP traps" on page 212.

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.

Accessing Resiliency Platform log files

You can use `logs-gather` option available with `support` command of klish menu to access the Resiliency Platform log files.

To access Resiliency Platform log files

1 Log in to the Resiliency Platform virtual appliance console or SSH session as an admin user.

2 Go to the **support** under **main menu**.

3. Run the logs-gather command with any of the log collection options that are available.

   See “Using klish” on page 256.

   The command collects the logs according to the option that you use with the command.

4. Once the logs are collected, a URL for downloading the log zip file is provided to you. You can enter the URL in a browser. You will be prompted to enter the admin user credentials and download the zip file.

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 25-1 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 25-2 lists the predefined risks available in Resiliency Platform. These risks are reflected in the current risk report and the historical risk report.

Table 25-2 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

Table 25-2 Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
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
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.

Table 25-2 Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
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
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

Table 25-2 Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
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.
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.

Table 25-2 Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
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.
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.

Table 25-2 Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
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.
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

About the Resiliency Manager services

The Resiliency Manager is a server that includes a set of loosely coupled services, a data repository, and a management console. The following is a list of services that can be started or stopped via klish on the Resiliency Platform virtual appliance.

See “Using klish” on page 256.

Table 25-3 Resiliency Manager services

Service or component name	Description
Database service (DB)	Supports the main data repository.
Core service	Provides the default platform functionality. Also includes critical capabilities such as security management, data repository access and external systems communication.
Licensing service	Provides the licensing capability.
Workflow Service (WF)	Provides the platform-level capability to deploy and execute workflows for other services in the platform.
Reporting Service	Provides the platform-level capability to deploy and run reports for other services in the platform.
Messaging Service (MQ)	The Messaging Service is the backbone of internal communication between all services in a Resiliency Manager.
Authentication Service (AT)	Provides consistent tokens and certificates across identity providers that can be used by Resiliency Platform authorization and rule-based access control (RBAC).
Scheduler Service	Provides the platform-level capability to schedule execution of a job (report, workflow, API, etc.) for other services in the platform. Though schedule settings are maintained at the main data repository and available consistently to all Resiliency Managers, the schedule runs at only one Resiliency Manager instance.
User Interface Service (UI)	Provides the web-based user interface for the product.
Recovery Automation Service (RA)	Provides disaster recovery capability for virtual machines and applications.

Components of Resiliency Platform virtual appliance

Following components are deployed while deploying the Resiliency Platform virtual appliance:

Table 25-4

Components	Description
Operating System	Hardened CentOS 6.7 Minimal operating system. The operating system is hardened or customized to contain only those packages that are required to run the application.
Veritas Resiliency Platform	Veritas Resiliency Platform provides core and standard services framework for the solution.
Resiliency Manager	Serves as the management console for Resiliency Platform. It also includes the database and the Resiliency Platform services.
Infrastructure Management Server (IMS)	Serves as the infrastructure manager or asset manager for Resiliency Platform.
Command Line Interface Shell (klish)	Command Line Interface Shell (klish) is used to provide the user a limited menu-based access to the operating system and the application.

See “About deploying the Resiliency Platform virtual appliance” on page 24.

Troubleshooting discovery of assets

When asset infrastructure is added to the Infrastructure Management Server (IMS), or when changes are made to the infrastructure, the IMS discovers and correlates the asset information and displays the information on the Assets page of the Resiliency Platform console. The discovery can take some time before the information is updated on the console. Until discovery is complete, some information needed to configure resiliency groups may be missing from the Assets page on the console.

See “About the frequency of asset information discovery” on page 247.

If changes have been made to the asset infrastructure, you can use the Refresh operation on the assets to speed up discovery so that updated asset information is displayed more quickly in the console. To use the Refresh operation, right-click the asset and select Refresh.

Note: Occasionally, the data discovered from the Infrastructure Management server (IMS) may not be updated properly in the Resiliency Manager database. This situation may result in displaying incorrect information about the resiliency group state, replication state, and replication type. In such a case, refresh the appropriate assets in both the data centers.

If you are configuring replication using storage arrays in a VMware vCenter Server environment, you can use the following guidelines to speed up discovery or to troubleshoot information that is not being updated:

Table 25-5 Configuring asset infrastructure in IMS for storage arrays in VMware environment

Situation	Troubleshooting/best practices
Adding storage arrays as enclosures to IMS	Ensure that the storage arrays that are added to the IMS are the ones that provide storage to the ESX servers managed by the vCenter Server that is added to the IMS.
More than one IMS in data center	Ensure that the vCenter Server that is managing the ESX servers, and the enclosure providing storage to those servers, are added to the same IMS.
Refreshing the IMS after a change in infrastructure	Ensure that you use the Refresh operation on the correct vCenter Servers and enclosures where the change was made.
Refreshing the IMS after a change in infrastructure, where there is more than one IMS	Ensure that you use the Refresh operation in the correct IMS.

In the VMware and EMC SRDF environment, the general guideline is to add/refresh the enclosure before adding/refreshing the VMware vCenter Server.

Table 25-6 Configuring or refreshing asset infrastructure in IMS for VMware and EMC SRDF environment

Situation	Recommended sequence
You have not yet added the asset infrastructure.	Add the enclosure information in the IMS and let the discovery complete before adding the vCenter Server to the IMS.
You later provision new storage from an enclosure that is already configured in the IMS and mount datastores from the new storage.	Refresh the enclosure in the IMS, let the refresh task on the enclosure complete, and then refresh the vCenter Server in the IMS.
You provision storage from a new enclosure.	Add the new enclosure in the IMS and then refresh the vCenter Server after the enclosure discovery completes.

Table 25-6 Configuring or refreshing asset infrastructure in IMS for VMware and EMC SRDF environment *(continued)*

Situation	Recommended sequence
You are provisioning storage from an enclosure that is already configured in the IMS to new ESX servers managed by a vCenter Server.	Refresh the enclosure first, then add the vCenter Server to the IMS or refresh it if it is already added to the IMS.

In the VMware and NetApp SnapMirror environment, the general guideline is add/refresh the vCenter Server first, then add/refresh the NetApp enclosure.

Table 25-7 Configuring or refreshing asset infrastructure in IMS for storage arrays in VMware and NetApp SnapMirror environment

Situation	Recommended sequence
You have not yet added the asset infrastructure.	Add the vCenter Server to the IMS first and let the discovery complete before you add the NetApp enclosure.
You later provision storage from an existing NetApp enclosure and mount NFS datastores on ESX servers.	Refresh the vCenter Server first in the IMS, let the discovery complete and then refresh the NetApp enclosure.
You later provision storage from a new NetApp enclosure and mount NFS datastores on that ESX servers.	Refresh the vCenter Server first in the IMS, wait for the vCenter Server discovery to complete, and then add the new NetApp enclosure.

The recommended sequence for adding or modifying asset infrastructure for application discovery in the NetApp SnapMirror replication environment is as follows: Ensure that discovery of the hosts is complete before you add or refresh the NetApp enclosures.

See “Adding the asset infrastructure” on page 106.

About the frequency of asset information discovery

After you add the asset infrastructure, for example virtualization servers, to Veritas Resiliency Platform, the Infrastructure Management Server (IMS) discovers information about the assets and the information is displayed in the console. Thereafter, the IMS continues to discover and update the information. The following table describes how often the IMS performs discovery.

If you make changes to the asset infrastructure, such as adding or removing virtual machines, you can use the Refresh operation on assets to manually initiate the IMS discovery.

Note: The discovery is triggered when configuration changes occur on the hosts. If configuration changes are not detected on the managed hosts, the communication between the host and IMS is restricted to the heartbeat communication that occurs every five minutes.

Asset type	Discovery interval in minutes	Discovered information
Host	1440	The operating system and networking for the host. Typically, this information does not change frequently.
Applications	360	Supported applications and their storage dependencies.
Hyper-V	120	Virtual machines and storage discovery.
VMware	360	ESX servers, virtual machines, and their storage dependencies.
Enclosures	360	Logical devices, physical devices, host associations, replications, and other storage array-specific properties.
Replication Appliance	360	Replication and storage dependencies
Replication Gateway	120	Information about the replication gateway, including gateway pair information and Veritas Replication Set information.

Troubleshooting the connection between Resiliency Managers in a resiliency domain

Multiple Resiliency Managers that are part of the same domain synchronize their databases using built-in replication. Each Resiliency Manager has its own web console but because the data is synchronized, all consoles show the same data. Operations can be performed from any console and the results show in all the consoles in the resiliency domain.

In some cases the connection is lost between Resiliency Managers. In such a case, if you login to the console, a message is displayed to warn you about this and request that you confirm whether the other Resiliency Manager is down (outage).

If the Resiliency Manager administrator confirms that the other Resiliency Manager is down, you can click the confirmation on the message box and continue working on the console. When the other Resiliency Manager is brought up, the changes are synchronized.

However, if you check and the other Resiliency Manager is not down, the problem is in the network connection. In this case, you should not attempt to work in the Resiliency Manager console until the network connection is restored.

Troubleshooting removing a Resiliency Manager from a resiliency domain

In some cases you may want to remove a Resiliency Manager node from a resiliency domain.

Before removing a Resiliency Manager node, you should first remove the Resiliency Manager from the resiliency domain using the Leave Domain operation in the Veritas Resiliency Platform web console. Completing this operation ensures that the Resiliency Manager is cleanly decommissioned and that all references to it are removed from the Resiliency Manager database and no longer appear in the web console user interface. The Leave Domain operation has prerequisites that are documented in the procedure topic.

See “Removing a Resiliency Manager from a resiliency domain” on page 66.

The following gives more details for trouble scenarios, for example, if the Resiliency Manager you want to remove is not online or if the operation does not complete successfully.

Unable to bring the Resiliency Manager online	The Leave Domain operation requires that both Resiliency Managers be online. However, if you are unable to bring the Resiliency Manager you want to remove online, there is no problem with leaving it in a down state. The resiliency domain and other infrastructure components continue to function. If an Infrastructure Management Server (IMS) was connected to the Resiliency Manager that was down, the IMS will automatically reconnect itself to another Resiliency Manager in the same domain. In addition, you can add another Resiliency Manager and join it to the domain.
Unable to complete Leave Domain operation	The Leave Domain operation is a multistep process. First the Resiliency Manager decommissions itself. Then all references to it are removed from the Resiliency Manager database. Finally any IMS connected to the decommissioned Resiliency Manager is rerouted to another Resiliency Manager. You can use the Activities pane to verify that the Leave Domain operation completes. If the process fails before all steps are complete, the partially removed Resiliency Manager no longer operates. However, the resiliency domain continues to function.

Using Veritas Services and Operations Readiness Tools to find a Unique Message Identifier description and solution

You can use Veritas Services and Operations Readiness Tools (SORT) to find a Unique Message Identifier (UMI) description and solution.

To find a Unique Message Identifier description and solution

1. Point your Web browser to the following URL:

   http://sort.veritas.com

2. In the search field on the top right of any SORT page, enter the UMI code, and then click the search icon.

3. On the **Search Result** page, in the **Error codes** pane, click the link to your message code. If you have a large number of search results, use the check boxes at the top of the page to display only error codes to find your code more easily.

   The **Error Code details** page for the UMI code displays, which provides the description and any possible solutions.

4. If the information on the page does not provide an adequate solution to your issue, you can click one of the links on the page to do one of the following things:

   - Comment on the UMI or its solution.
   - Request a solution.
   - Add a solution of your own.

Section 8

Reference

Virtual appliance security features

This appendix includes the following topics:

Operating system security

Veritas Resiliency Platform appliance operating system is hardened against potential security exploitation by removing the operating system packages that are not used by the Resiliency Platform. All the default yum repository files that are shipped with the operating system are removed.

The Control + Alt + Delete key combination has been disabled to avoid any accidental reboot of the virtual appliance. Exec-shied is enabled to protect the virtual appliance from stack, heap, and integer overflows.

Management Security

Only two users are available on the appliance: admin user and support user. These two user accounts are used to access the appliance based on the requirement.

Only admin login is available for the appliance. The password policy of admin login is modified to prompt the user to change the password on the first login.

See “Password policies for Resiliency Platform virtual appliance” on page 57.

If the admin user password is lost, you need to contact Veritas support for resetting the admin user password.

On successful completion of the product bootstrap, admin user can only access a limited menu of commands through klish. Besides admin user, support user is also supported in the appliance but remote login of support user is disabled. To access the support user, one need to login as an admin and go through **klish**. An option `support > shell` is provided in the **klish** menu to switch the user to support and access the bash shell of support. After selecting this option, the support user is given superuser privileges. Using this option is not recommended and it should be used only with the assistance of technical support.

Timeout of the bash shells of all users is set to 900 seconds.

Network security

The TCP timestamp responses are disabled in Resiliency Platform virtual appliance. Another network security feature of the appliance is that during the product bootstrap process, only those ports that are used by the product for communication and data transfer, are opened through the firewall and all the other communications are blocked.

Uncommon network protocols such as DCCP, SCTP, RDC, TIPC have been disabled so that any process cannot load them dynamically.

Access control security

Resiliency Platform virtual appliance implements certain access control measures. The umask is set to 0700 across the appliance. The access permissions of some of the files such as home folder of root, the log directory etc. is restricted. All the security and the authorization messages are logged into the appliance.

Physical security

In the Resiliency Platform virtual appliance, the USB storage access is disabled.

Using klish menu in Resiliency Platform

This appendix includes the following topics:

- About klish
- Using klish

About klish

Once the Veritas Resiliency Platform virtual appliance is deployed and configured, you are given limited, menu-based access to the operating system and the product. You need to use Command Line Interface Shell (klish) menu to manage the configuration-related changes to the product.

You can use klish menu to do the following:

- Manage the Veritas Resiliency Platform appliance
- Monitor the Veritas Resiliency Platform appliance activities
- Change some of the network configurations
- Change the system settings
- Access the Veritas Resiliency Platform logs
- Manage Veritas Resiliency Platform updates and patches

See “Using klish” on page 256.

Using klish

After the product configuration, whenever you log in to the Resiliency Platform appliance, you get the main menu of klish. This menu is the starting point, from which you can configure, manage, monitor, and support your application using the command line.

You can reconfigure or modify some of the appliance settings that are configured through the product bootstrap. Following are the settings that you can reconfigure using klish:

- **Network settings:** You can reconfigure the subnet mask, default gateway, DNS server, and search domains using the klish menu.
  You cannot reconfigure the hostname that you had configured through the bootstrap process. In case of static DHCP, you cannot change the network settings using the klish menu. You cannot change the network settings for any component that is configured in the cloud environment.
- **System settings:** You can reset the timezone and NTP server using klish menu. Changing the system settings can affect the product functionality if incorrect values are set.

You can also perform logical volume management (LVM) operations such as adding a disk or removing a disk using the klish menu.

You can press the **tab** or **space** key to display the menu options. Press **?** key to display detailed help.

If a klish command is expected to perform any operation on an entity such as start or stop services, it goes into lock mode and does not allow same operation from any other session till the first operation gets completed. You can see a message similar to the following:

Warning: Try to get lock. Please wait...

Error: Can't get lock.

Table B-1 Options available in the **main** menu

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
manage	Manage appliance Table B-2

Table B-1 Options available in the **main** menu *(continued)*

Menu option	Description
monitor	Monitor appliance activities Table B-7
network	Network configuration Table B-9
settings	Appliance settings Table B-15
support	Access logs Table B-19
updates	Manage updates and patches Table B-21

Table B-2 Options available with **manage** command

Menu option	Description
back	Return to the previous menu
configure	Configure Resiliency Platform component or show the configured component Table B-3
datamover	Manage Resiliency Platform Data Mover activities and objects This option is available only on a Replication Gateway or Storage Proxy appliance Table B-4
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
infra-appliance	List or remove Replication Gateway appliance Table B-5

Table B-2 Options available with **manage** command *(continued)*

Menu option	Description
services	Manage the appliance services ■ If the appliance has been configured as a Resiliency Manager or IMS, use **rm** or **ims** as first parameter and options available in the services menu as second parameter. Table B-6 See "About the Resiliency Manager services " on page 243. ■ If the appliance has been configured as a Replication Gateway, use the options available in the services menu as first parameter.

Table B-3 Options available with **configure** command

Menu option	Description
ims_register	Register the IMS using the registration URL obtained after initiating the Add IMS operation See "Adding an IMS " on page 70.
ims	Configure Infrastructure Management Server
rm	Configure Resiliency Manager
show	Show the configured component

Table B-4 Options available with **datamover** command

Menu option	Description
start	Start a Veritas Replication Set
abort	Stop a Veritas Replication Set
delete	Delete a Veritas Replication Set
clear-admin-wait	Clear the admin Wait status for the Veritas Replication Set

Table B-5 Options available with **infra-appliance** command

Menu option	Description
list	List the Replication gateway appliance.

Table B-5 Options available with **infra-appliance** command *(continued)*

Menu option	Description
remove	Remove the Replication Gateway appliance. You need to remove the Gateway pair before you remove the Gateway. See "Removing a Veritas Replication Gateway pair" on page 103.

Table B-6 Options available with **services** command

Menu option	Description
show	Show Resiliency Platform services. the short service names displayed here are used while exercising other options with services command such as restart, start, status.
restart	Restart Resiliency Platform services Two options available are: `restart` *`all`* where, *all* means all the services. `restart` *`service name`* where, *service name* is the name of a particular service. You can provide multiple service names (comma separated).
start	Start Resiliency Platform services Two options available are: `start` *`all`* where, *all* means all the services. `start` *`service name`* where, *service name* is the name of a particular service. You can provide multiple service names (comma separated).
status	Check the status of Resiliency Platform services Two options available are: `status` *`all`* where, *all* means all the services. `status` *`service name`* where, *service name* is the name of a particular service. You can provide multiple service names (comma separated).

Table B-6 Options available with **services** command *(continued)*

Menu option	Description
stop	Stop Resiliency Platform services Two options available are: `stop` *`all`* where, *all* means all the services. `stop` *`service name`* where, *service name* is the name of a particular service. You can provide multiple service names (comma separated).

Table B-7 Options available with **monitor** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
datamover	Display VRP Datamover activities and objects This option is available only on a Replication Gateway or Storage Proxy appliance Table B-8
FSusage	Display filesystem usage
help	Display an overview of the CLI syntax
top	Display the top process information
uptime	Display the uptime statistics for the appliance
who	Display who is currently logged into the appliance

Table B-8 Options available with **datamover** command

Menu option	Description
repl-sets	Display the details about Veritas Replication Sets including RPO, connection state, replication state
update-sets	Display the list of current update sets which are in transit
ingress-data	Display the IO statistics for the data transfer from protected virtual or physical machine to Gateway (IOReceiver statistics)

Table B-8 Options available with **datamover** command *(continued)*

Menu option	Description
network-data	Display the network related statistics for data transfer between production site Gateway and recovery site Gateway (Transceiver statistics)
disk-data	Display the IO statistics for the data write on recovery site disks (Applier statistics)

Table B-9 Options available with **network** command

Menu option	Description
back	Return to the previous menu
dns	Show or change the DNS Table B-10
exit	Log out from the current CLI session
gateway	Show or change the Gateway Table B-11
help	Display an overview of the CLI syntax
hostname	Show the hostname
ip	Show or change the IP address Table B-12
netmask	Show or change the netmask Table B-13
search-domain	Show or change the domain Table B-14

Table B-10 Options available with **dns** command

Menu option	Description
set	Configure Domain Name Server
show	Show the current Domain Name Server

Table B-11 Options available with **gateway** command

Menu option	Description
set	Configure Gateway
show	Show the current Gateway

Table B-12 Options available with **ip** command

Menu option	Description
set	Configure the IP address for additional NIC
show	Show the current IP address

Table B-13 Options available with **netmask** command

Menu option	Description
set	Configure the netmask
show	Show the current netmask

Table B-14 Options available with **search-domain** command

Menu option	Description
add	Add search-domain
remove	Remove the search domain name
show	Show the search domain settings

Table B-15 Options available with **settings** command

Menu option	Description
back	Return to the previous menu
change-password	Change the admin user password for the appliance
date	Display the current date and time for the appliance Table B-16
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax

Table B-15 Options available with **settings** command *(continued)*

Menu option	Description
lvm	Perform operations related to logical volume manager on the appliance Table B-17
ntp	Perform operations related to NTP server
poweroff	Shut down the appliance
reboot	Restart the appliance
timezone	Show or change the timezone for the appliance Table B-18

Table B-16 Options available with **date** command

Menu option	Description
show	Show the time and date

Table B-17 Options available with **lvm** command

Menu option	Description
add-disk	Add disk to the data volume. You need to attach a disk before adding it.
list-free-disk	List the free disks
initialize-free-disk	Initialize the newly attached free disk
list-used-disk	List the disks used by the data volume
remove-disk	Remove disk from the data volume. Make sure that you have an extra disk to migrate the data before removing a disk.

Note: In case you initialize the newly-added disk during add-disk or remove-disk operation, the existing data on the new disk is deleted.

Table B-18 Options available with **timezone** command

Menu option	Description
set	Set the timezone for the appliance

Table B-18 Options available with **timezone** command *(continued)*

Menu option	Description
show	Show the current timezone for the appliance

Table B-19 Options available with **support** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
loggather	■ If the appliance has been configured as a Resiliency Manager or an IMS, then various options will be available for collecting the Resiliency Manager and IMS logs. Table B-20 ■ If the appliance has been configured as a Replication Gateway, then `loggather` command will collect the logs of the Replication gateway.
shell	Open the bash shell prompt for support user

Table B-20 Options available with **loggather** command

Menu option	Description
basic	Gather logs of Resiliency Manager and IMS without database
full	Gather logs of Resiliency Manager and IMS with database
fullims	Gather logs of IMS with database
fullrm	Gather logs of Resiliency Manager with database
ims	Gather logs of IMS
rm	Gather logs of Resiliency Manager

Table B-21 Options available with **updates** command

Menu option	Description
config-local-iso-repository	Configure the repository from locally mounted ISO image on CD-ROM

Table B-21 Options available with **updates** command *(continued)*

Menu option	Description
apply-update	Apply the specified update
back	Return to the previous menu
config-repository	Configure the repository Table B-22
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
list-updates	List the applicable updates
remove-repository	Remove current repository configuration
show-readme	Show readme for the specified update
show-repository	Show current repository configuration
show-version	Show appliance version

Table B-22 Options available with **config-repository** command

Menu option	Description
hostname	hostname of the repository server
protocol	Protocol on which the repository server is configured
port	Port on which the repository server is configured
RepoPath	Path on which the repository server is configured

See “About klish” on page 255.

See “Accessing Resiliency Platform log files” on page 236.

Tips on using the web console

This appendix includes the following topics:

- Tour of the Resiliency Platform web console screen
- Using Quick Actions for shortcuts to common tasks
- Filtering and searching for objects in the web console
- About settings in the web console
- About the Resiliency Platform Dashboard
- Web console icons

Tour of the Resiliency Platform web console screen

The numbered screen areas are illustrated below the table.

Table C-1 Overview of the web console screen areas

Screen areas	Description
1 - Menu bar	Menu options for reports, resiliency plans, views, settings, notifications, inbox, and online help. See "Menu bar options" on page 267.
2 - Navigation pane	Icons to open pages for configuring and implementing start/stop and disaster recovery operations. See "Navigation pane options" on page 269.

Table C-1 Overview of the web console screen areas *(continued)*

Screen areas	Description
3 - Dashboard	The console home page - clicking the Home icon in the navigation pane returns to the Dashboard. View an overview of assets in the resiliency domain and their current status. Drill down for details. See "About the Resiliency Platform Dashboard" on page 272.

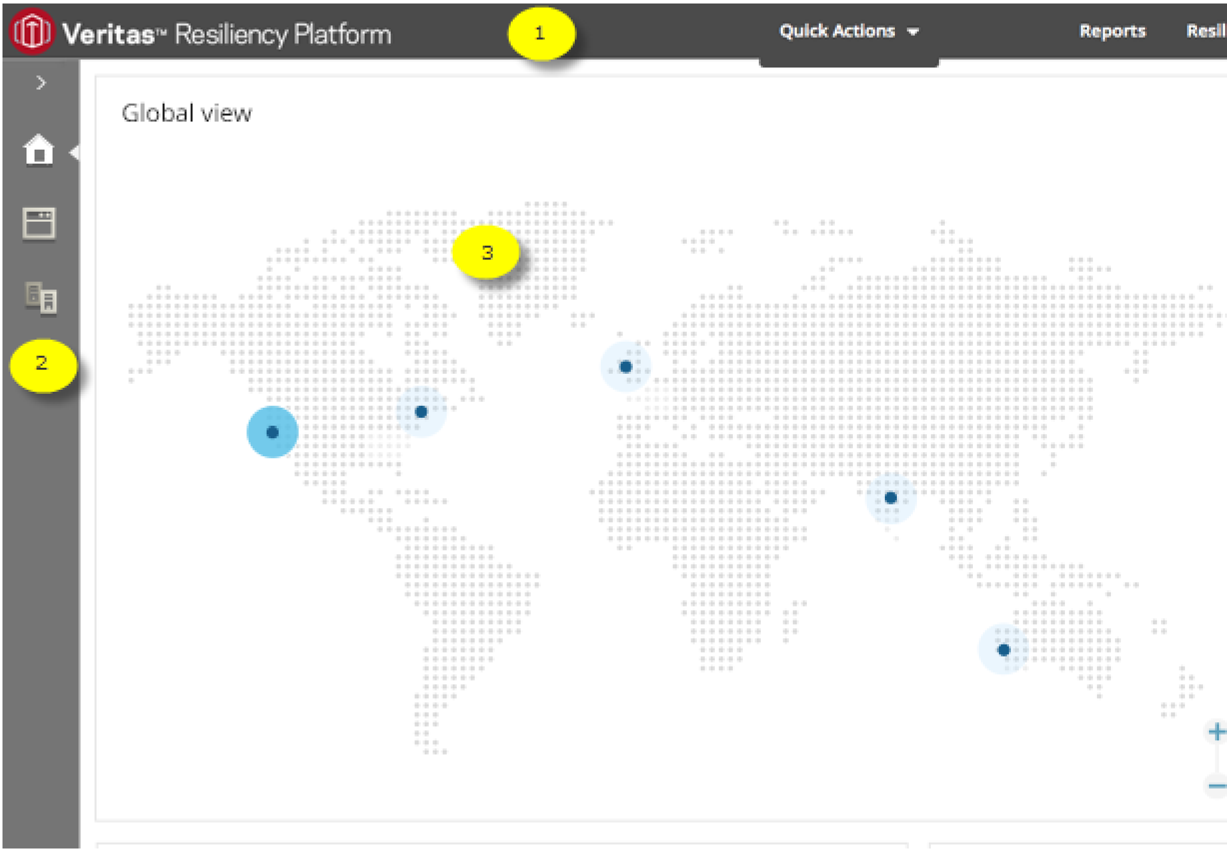

Menu bar options

The menu bar is located at the top of the console window.

Table C-2 Menu bar options for the Veritas Resiliency Platform web console

Options	Description
Quick Actions ▾	Open drop-down selection of shortcuts to common tasks. See "Using Quick Actions for shortcuts to common tasks" on page 269.
Reports	Schedule and run reports. View reports showing data center and asset status. See "About reports" on page 200.
Resiliency Plans	Create and run custom resiliency plans for starting, stopping, and migrating resiliency groups. See the Solutions guide for details on resiliency plans.
⊞	More views View activities, risks, and logs. See "Viewing events and logs in the console" on page 235.
⚙	Settings Open Settings page for configuring and maintaining product infrastructure and other settings. See "About settings in the web console" on page 270.
🔔	Notifications Display most recent notifications. Requires alerts and notifications to be enabled using Settings page. See "Managing settings for alerts and notifications and miscellaneous product settings" on page 208.
✉	Inbox View actions to be completed.
?	Help Open Help window where you can search all help or filter by category.

Table C-2 Menu bar options for the Veritas Resiliency Platform web console *(continued)*

Options	Description
👤	Log out of console. Shows Resiliency Manager, resiliency domain, and data center.

Navigation pane options

The navigation pane is located on the left side of the console window.

Click the arrow on the top of the navigation pane to expand or contract the pane and view labels for icons.

Table C-3 Left navigation pane options for the Veritas Resiliency Platform web console

Options	Description
🏠	Returns to Home page Dashboard
▤	Opens the Assets page for configuring resiliency groups, viewing details of assets, and performing start and stop or disaster recovery operations
▥	Opens page for configuring disaster recovery settings such as network mapping and replication gateway pairs

Using Quick Actions for shortcuts to common tasks

In the Veritas Resiliency Platform web console, you can use the Quick Actions pull-down menu for shortcuts to go to common tasks.

To use Quick Actions for shortcuts to common tasks

1 Navigate

On the top menu bar, click **Quick Actions** to display available shortcuts.

2 The menu of available shortcuts is displayed. Click the desired shortcut.

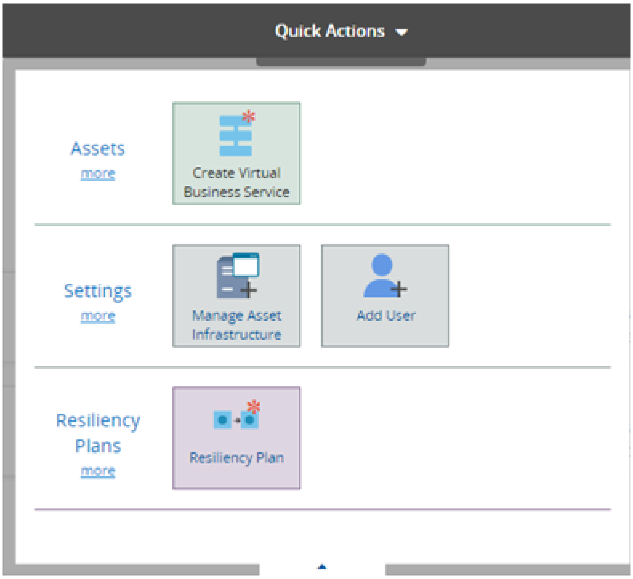

Filtering and searching for objects in the web console

On pages that list multiple objects, for example, virtual machines listed on the Assets page, the web console lets you select object types as a filter or search by first letters of a name. To see the full list again, clear the filter or search field.

You can also double-click to drill down to a more detailed view. For example, you can drill down from a row of a table that lists virtual machines, or from a Dashboard graphic showing information on virtual machine status.

About settings in the web console

With appropriate permissions you can modify Veritas Resiliency Platform infrastructure and other product settings under the **Details view** of the **Settings** page. In the **Map view** , you can see the map view of the data centers. You access the Settings page from the menu bar.

⚙ Settings

Table C-4 Settings page options

Type of setting	Description	More information
Infrastructure	Manage Resiliency Managers and resiliency domains Add and manage Infrastructure Management Servers (IMS) Add the asset infrastructure Manage data centers and their network settings	See "Adding a Resiliency Manager to an existing resiliency domain" on page 65. See "Adding the asset infrastructure" on page 106. For information on managing data centers, refer to the `Solutions Guide`. See "Managing data centers" on page 75.
Updates	View and deploy product updates	See "About applying updates to Resiliency Platform" on page 217.
Application Support	View and manage application types	For information on application support, refer to the `Solutions for Applications` guide.
Service Objectives	View and manage service objectives	See "Viewing or customizing service objectives" on page 198.
User Management	Configure authentication domains, users and user groups, and user personas (roles)	See "Managing user authentication and permissions" on page 175.
Licenses	View and manage product licenses	See "Managing licenses" on page 171.
Alerts and Notifications	Configure email and SNMP for alerts and notifications, configure rules for notification	See "Managing settings for alerts and notifications and miscellaneous product settings" on page 208.
Miscellaneous	Configure general product settings such as purge intervals and telemetry	See "Managing settings for alerts and notifications and miscellaneous product settings" on page 208.

Note: Additional settings that relate to disaster recovery configuration are available from the navigation pane.

See “Tour of the Resiliency Platform web console screen” on page 266.

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.

Web console icons

The following is a summary of icons that appear on the Veritas Resiliency Platform web console.

Table C-5 Web console icons

Icon	Description	Location
⊞	More views Menu options for Activities, Logs, Risks	Menu bar
⚙	Settings Opens Settings page	Menu bar

Table C-5 Web console icons *(continued)*

Icon	Description	Location
🔔	Notifications Displays notifications Requires alerts and notifications to be enabled using Settings page	Menu bar
✉	Inbox View actions to be completed.	Menu bar
?	Help Opens Help window where you can search all help or filter by category	Menu bar
👤	Log out of console Shows user login and information about Resiliency Manager, resiliency domain, and data center	Menu bar
🏠	Home Returns to the Home page Dashboard	Navigation pane
🗔	Assets Opens the Assets page for configuring resiliency groups, viewing details of assets, and performing start and stop or disaster recovery operations	Navigation pane
🖧	Disaster Recovery Settings Opens page for configuring disaster recovery settings such as network mapping and replication gateway pairs	Navigation pane
⋮	Vertical ellipsis Displays list of actions for selected object	To the right of a selected object in a list

Glossary

Index