Veritas™ Resiliency Platform 1.0: Deployment Guide

July 2015

VERITAS™

Veritas Resiliency Platform: Deployment Guide

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product version: 1.0

Document version: 1.0

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. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our website at the following URL:

www.symantec.com/business/support/index.jsp

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/contact_techsupp_static.jsp

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information

- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
  - Error messages and log files
  - Troubleshooting that was performed before contacting Symantec
  - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan	customercare_apj@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

http://www.symantec.com/connect/storage-management

Documentation

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

doc_feedback@symantec.com

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

https://www-secure.symantec.com/connect/storage-management/forums/storage-and-clustering-documentation

Contents

Chapter 1

Overview of Resiliency Platform deployment

This chapter includes the following topics:

About Veritas Resiliency Platform

Veritas Resiliency Platform offers a unified approach for visibility and control of IT service continuity for applications, virtual machines, and complex, multi-tier business services across a global landscape.

Resiliency Platform has the following core capabilities:

Recovery

Resiliency Platform provides a disaster recovery (DR) solution using data centers on premises in different geographical locations. The management console simplifies recovery, with single-click rehearsal and recovery operations.

Visibility

The console Dashboard provides visibility into the health of applications, virtual machines, and multi-tier business services.

Orchestration

Resiliency Platform can assist in data center day-to-day workload automation activities. For instance, virtual machines or IT services can be started and stopped for maintenance.

About Resiliency Platform features and components

The following is a brief introduction to Veritas Resiliency Platform key components and features. Administrators responsible for deploying and configuring Resiliency Platform 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 16.

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 17.

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 18.

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 the IMS for discovery and monitoring.

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.

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 migrate/takeover the entire VBS.

See "Resiliency Platform component diagram" on page 14.

Resiliency Platform deployment infrastructure

A typical deployment of Veritas Resiliency Platform 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 Veritas Resiliency Platform deployment infrastructure and how Resiliency Manager, IMS, and various assets associated with the IMS interact with each other.

Figure 1-1 Deployment infrastructure

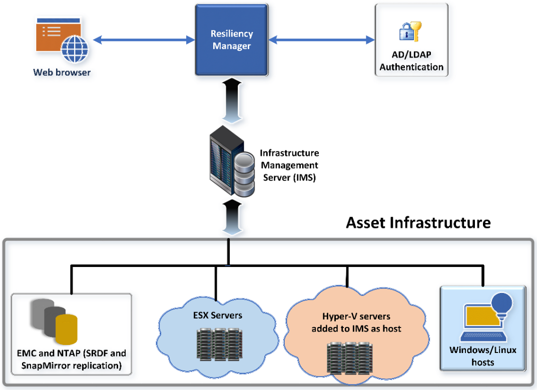
Web browser
Resiliency Manager
AD/LDAP Authentication
Infrastructure Management Server (IMS)
Asset Infrastructure
EMC and NTAP (SRDF and SnapMirror replication)
ESX Servers
Hyper-V servers added to IMS as host
Windows/Linux hosts

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 the IMS for IMS discovery and monitoring. The asset infrastructure can include hosts (Windows or Linux servers), virtualization servers for Hyper-V and VMware, and enclosures (storage arrays).

See “Resiliency Platform deployment infrastructure” on page 14.

This diagram does not show the replication details for the asset infrastructure.

See “Replication in a Resiliency Platform deployment” on page 19.

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

Figure 1-2 Resiliency Platform components

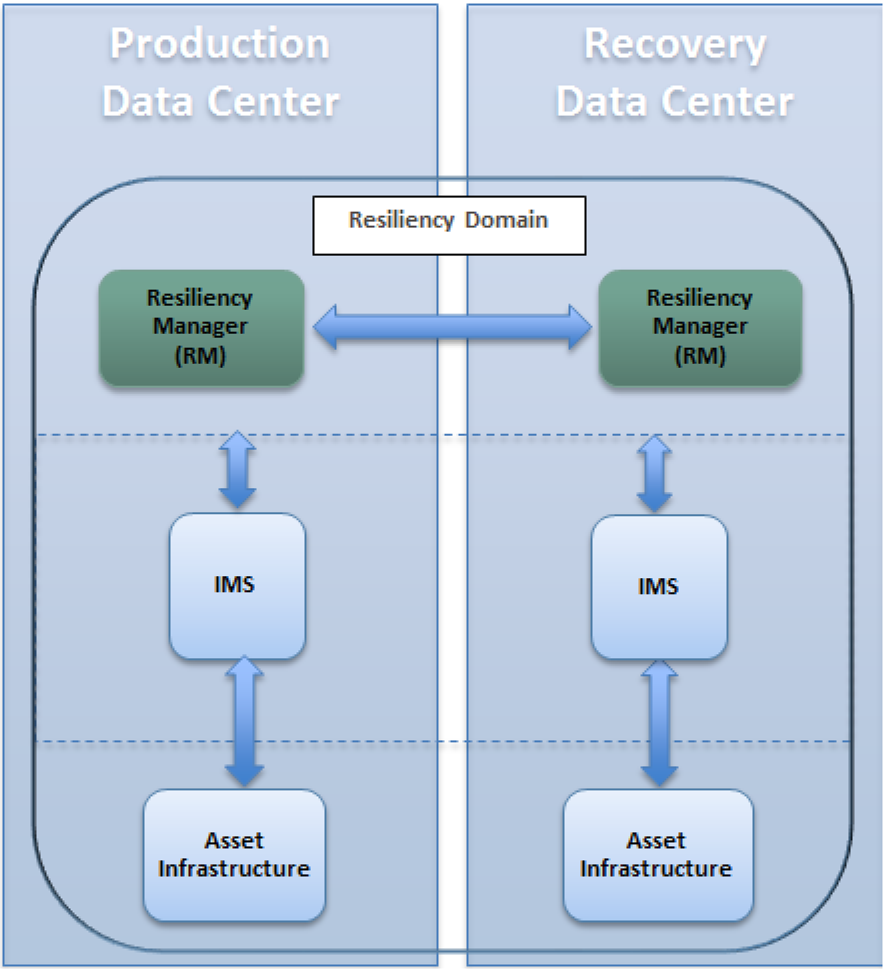

Resiliency domain

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.

See “Resiliency Platform component diagram” on page 14.

See “Resiliency Manager” on page 17.

See “Infrastructure Management Server (IMS)” on page 18.

Resiliency Manager

The Resiliency Manager component of Veritas Resiliency Platform 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 14.

See “Resiliency domain” on page 16.

See “Infrastructure Management Server (IMS)” on page 18.

Infrastructure Management Server (IMS)

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

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

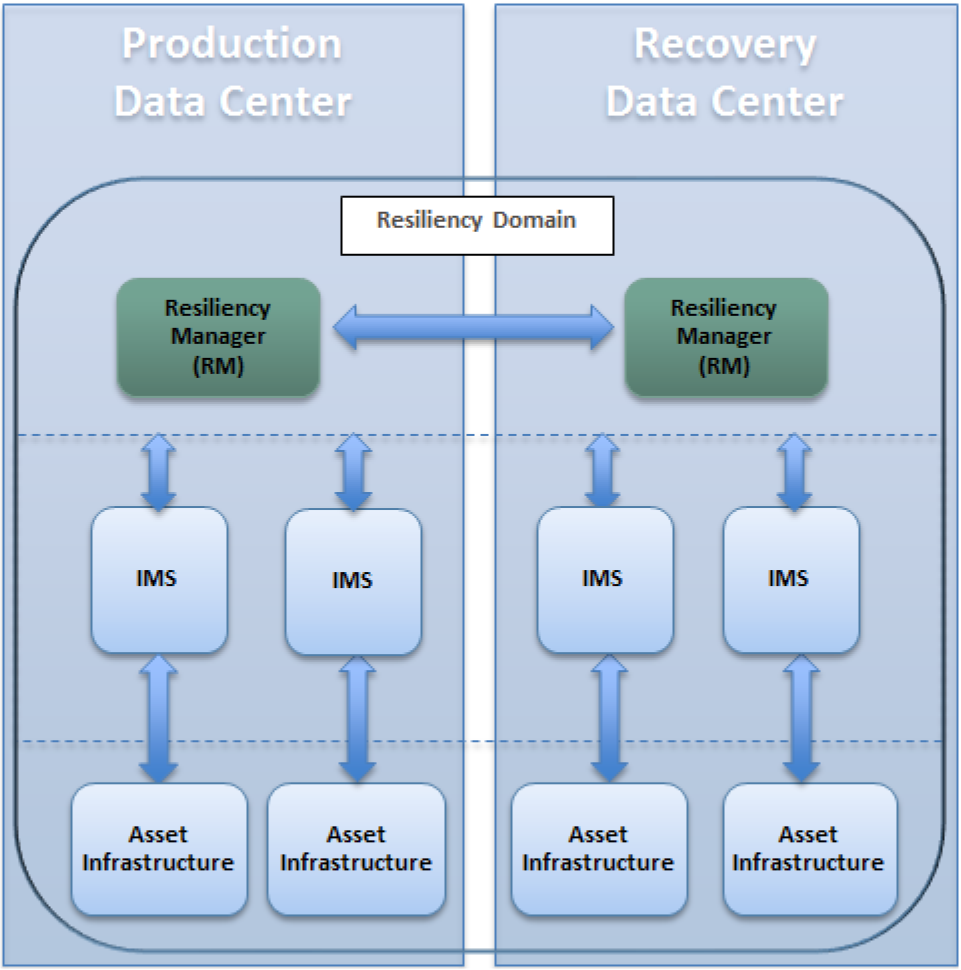

See “Resiliency domain” on page 16.

See “Resiliency Manager” on page 17.

Replication in a Resiliency Platform deployment

Veritas Resiliency Platform provides data recovery from your production data center to your recovery data center using two forms of replication: array-based replication (block-based replication) and hypervisor-based replication. Figure 1-4 depicts how replication is performed in Veritas Resiliency Platform.

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

Figure 1-4 Replication in a Resiliency Platform deployment

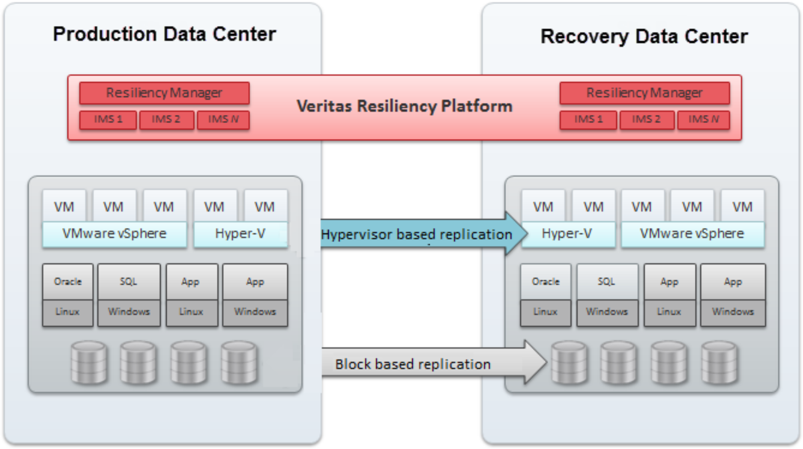

More information on configuring replication for use with Resiliency Platform is available in the solutions guides:

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 16.

See "Resiliency Manager" on page 17.

See "Infrastructure Management Server (IMS)" on page 18.

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 27.
2	Configure the virtual appliances as Resiliency Platform components	Deployment Guide See "About configuring the Resiliency Platform components" on page 31.
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 35.
4	Finish configuring the settings for the resiliency domain	Deployment Guide See "Adding an IMS " on page 44. See "Managing user authentication and permissions" on page 101. See "Managing settings for alerts and notifications and general product settings" on page 124.
5	Add the asset infrastructure to the Infrastructure Management Server (IMS)	Deployment Guide See "Adding the asset infrastructure to an Infrastructure Management Server (IMS)" on page 62.

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

Step	Description	More information
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 for Virtual Business Services

Chapter 2

System requirements

This chapter includes the following topics:

Supported hypervisors for 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:

- ESX 5.5
- vCenter Server 5.5

See “About deploying the Resiliency Platform virtual appliance” on page 27.

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 and Infrastructure Management Server (IMS):

Disk space	60 GB
RAM	16 GB
Virtual CPU	8

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

In addition to the above mentioned resources, you need a 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.

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.

See “About deploying the Resiliency Platform virtual appliance” on page 27.

Virtual appliance security requirements

Veritas Resiliency Platform virtual appliance implements a number of features to ensure the security of the product.

See “About virtual appliance security” on page 150.

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.
- 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 2-1 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
14176	Used for communication between the Resiliency Manager and Infrastructure Management Server (IMS)	Resiliency Manager and IMS Resiliency Managers of the two data centers	Bi-directional	TCP
7000	Used for database replication	Resiliency Managers of the two data centers	Bi-directional	TCP
7001	Used for database replication	Resiliency Managers of the two data centers	Bi-directional	TCP
22	Used for communication between remote host to the appliance CLISH access	Appliance and the hosts	Bi-directional	TCP
123	Used for NTP synchronization	Appliance and the NTP server	Bi-directional	TCP

Table 2-2 Ports used for 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 Resiliency Managers of the two data centers	Bi-directional	TCP
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 CLISH 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

See “About deploying the Resiliency Platform virtual appliance” on page 27.

Deploying the Resiliency Platform virtual appliance

This chapter includes the following topics:

About deploying the Resiliency Platform virtual appliance

Veritas Resiliency Platform is installed 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.

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.

You can deploy a Resiliency Platform virtual appliance using Hyper-V Manager as well as using 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 29.

See "Deploying the virtual appliance through VMware vSphere Client" on page 28.

Downloading the Resiliency Platform virtual appliance

You can download Veritas Resiliency Platform virtual appliance from the following URL:

- Symantec FileConnect:
  https://symantec.flexnetoperations.com/control/symc/registeranonymouslicensetoken
- Trialware:
  http://www.symantec.com/products-solutions/trialware/

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. The names of the .zip file and the VHD file for Hyper-V are as follows:

- Download file name:

  `Veritas_Resiliency_Platform_Hyper-V_Virtual_Appliance_1.0.0.0_IE.zip`
- VHD file names:

  `itrp-ra-disk1`

  `itrp-ra-disk2`

To deploy Resiliency Platform virtual appliance through VMware, you need to download an Open Virtualization Archive (OVA) file. The name of the OVA file for VMware is as follows:

`Veritas_Resiliency_Platform_VMWare_Virtual_Appliance_1.0.0.0_IE.ova`

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. 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. Power on the virtual machine.

You can now configure the Resiliency Platform component.

See "About configuring the Resiliency Platform components" on page 31.

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 appliance” on page 28.

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. 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 31.

Chapter 4

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, the virtual appliance bootstrap process requires you to configure the Resiliency Platform component that you have deployed. The following settings are configured as part of this process to set up the component:

- **Network settings:** Settings such as hostname, IP address, subnet mask, default gateway, and DNS server.
- **System 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: Before using the hostname and the IP address in the **Network settings**, you need to register them with the DNS server. The hostname and the IP address that you use for product configuration, cannot 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. If you want to change these settings later, you need to use the Command Line Interface SHell (CLISH) menu for changing these settings.

See “Configuring the Resiliency Platform component through the product bootstrap” on page 32.

Prerequisites for configuring the Resiliency Platform component

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 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.

Configuring the Resiliency Platform component through the product bootstrap

After Veritas Resiliency Platform (Resiliency Platform) deployment, you need to set up the Resiliency Platform component during the product bootstrap. The bootstrap script is used to set up the Resiliency Platform component only for the first time. Later, if you want to reconfigure the Resiliency Platform component, you need to use the CLISH menu.

To configure the Resiliency Platform node through the product bootstrap

1. Log in to the virtual appliance console for the first time using the following credentials:
   - **Username:** admin
   - **Password:** P@ssw0rd

After a successful login, you need to change the password of the admin user. The new password that you enter must not be a dictionary word, and must be at least six characters long.

2 The bootstrap process is automatically invoked once you change the admin password after deploying the virtual appliance. The first step in the bootstrap process is to display the End User License Agreement (EULA). Accept the EULA to proceed with the configuration.

3 In the **Network Settings** section, you need to enter your choice for the network type. Type **1** for selecting static IP or **2** for selecting static DHCP.

In case of static 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 In the **System 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 timezone. Select the appropriate options to set your timezone and verify the displayed information.
- Enter the FQDN or IP address of the NTP server.

5 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 test or evaluation purposes, you can deploy a Resiliency Manager and Infrastructure Management Server (IMS) on the same virtual appliance.

6 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 31.

Chapter 5

Setting up the resiliency domain

This chapter includes the following topics:

- About the web console
- Connecting to the Resiliency Platform web console
- Getting started with a new Resiliency Platform configuration
- Adding a Resiliency Manager to an existing resiliency domain
- Removing a Resiliency Manager from a resiliency domain
- Viewing the status of a Resiliency Manager in a data center

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 35.

Note: For the best console experience, use a minimum resolution of 1280x960.

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 35.

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 161.

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
  The initial credentials that are required are for the Admin user of the Resiliency Manager.
  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 and click **Login**.

See "Getting started with a new Resiliency Platform configuration" on page 35.

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 38.

Prerequisites:

- The basic configuration includes the Resiliency Manager and Infrastructure Management Server (IMS).
  See “About Resiliency Platform features and components” on page 13.
- 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 5-1 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 128.
3. Enable Solutions Licenses	You can select a license file to apply or enable the trial license. See “About licenses” on page 99.

Table 5-1 Getting Started wizard *(continued)*

Wizard step	Details
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. See "Options for Configure Domain" on page 109. 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 101.

Table 5-1 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>
<p>Optional.</p>
<p>Add an Infrastructure Management Server (IMS). Optionally you can add more than one.</p>
<p>You can also add an IMS later from the <b>Settings</b> page.</p>
<ul>
<li>Choose one of the following
<ul>
<li>To add an IMS co-located with the Resiliency Manager, choose <b>Enable internal IMS</b>.</li>
<li>If you deployed an IMS separate from the Resiliency Manager, choose <b>Connect to IMS</b>.</li>
</ul></li>
<li>Fill in the following information and click <b>Add</b>:</li>
</ul>
<table>
<tr><td>Data Center Location</td><td><p>Select the data center location, for example, the city.</p><p>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.</p></td></tr>
<tr><td>Server Name</td><td><p>If you are adding an IMS separate from the Resiliency Manager, enter the fully qualified host name.</p><p>For the login, use the Admin login credentials for the IMS virtual appliance.</p></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></td><td><p>When you are done adding IMSs, you can exit the wizard. You can complete any steps you skipped from the <b>Settings</b> page.</p><p>See “About settings in the web console” on page 165.</p></td></tr>
</table>

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 44.

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 45.
- 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)

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

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

⋮ Click 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 148.

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 147.

Leaving Domain

The Leave Domain operation has been initiated and is in progress.

See “Removing a Resiliency Manager from a resiliency domain” on page 39.

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)

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

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

Chapter 6

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 18.

A Resiliency Platform domain can extend across data centers in different locations. Typically 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 35.
   - 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 plus the Admin credentials for the IMS virtual appliance.

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 **Infrastructure Management Server +**.

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 43.

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.

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 to an Infrastructure Management Server (IMS)” on page 62.

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.

See “Configuring DNS server settings for a data center” on page 51.

Removing an IMS

In the Veritas Resiliency Platform 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.

Caution: Removing an IMS also removes the Resiliency Manager repository records of the asset infrastructure that was added to that IMS. Assets managed by that IMS that were added to Resiliency Platform resiliency groups are removed from the resiliency groups. You may want to edit the resiliency groups before removing the IMS to remove the assets managed by this IMS. If the resiliency group contains only assets from this IMS, you can remove the resiliency group.

More information on resiliency groups is available in the solutions guides.

To remove an IMS

1 Navigate

⚙ **Settings** (menu bar)

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

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

⋮ Click 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 Veritas Resiliency Platform 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 to an Infrastructure Management Server (IMS)" on page 62.

To modify an IMS

1 Navigate

⚙ **Settings** (menu bar)

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

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

⋮ Click the vertical ellipsis for the IMS > **Modify**.

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 43.

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 Veritas Resiliency Platform web console to reconnect an Infrastructure Management Server (IMS) to a Resiliency Manager. 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)

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

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

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

Managing data centers

In Veritas Resiliency Platform, 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 43.

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.

See “Configuring network settings for data centers” on page 50.

To view data centers

1. Navigate

   ⚙ **Settings** (menu bar)

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

2. You can expand or contract each data center listed. Click the arrow to the right of a data center to expand the data center and view the information about any associated Resiliency Manager and IMS, including status of connections.

To add a data center

1. Navigate

   ⚙ **Settings** (menu bar)

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

2. Click **Data Center +**.
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.

To modify a data center

1. Navigate

   ⚙ **Settings** (menu bar)

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

2. ⋮ Click the vertical ellipsis next to the data center name, then click **Edit Data Center**.
3. Edit the data center information and submit the changes.

To delete a data center

1 Navigate

⚙ **Settings** (menu bar)

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

2 Click the arrow to the right of a data center to expand the data center and verify 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.

3 ⋮ Click the vertical ellipsis next to the data center name, then click **Delete Data Center**.

Configuring network settings for data centers

In the Veritas Resiliency Platform web console, you can configure DNS and subnet settings for data centers. These settings are used for disaster recovery operations between data centers.

To configure network settings for data centers

1 Navigate

⚙ **Settings** (menu bar)

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

⋮ Click the vertical ellipsis next to the data center name, then click **DNS & Network Settings**.

2 (For application disaster recovery only) On the **DNS** tab, add, modify, or remove information about the DNS server.

See “Configuring DNS server settings for a data center” on page 51.

3 On the **Subnets** tab, add, modify, or remove information about subnets.

See “Configuring subnet information for a data center” on page 52.

See “Managing data centers” on page 48.

Configuring DNS server settings for a data center

In the Veritas Resiliency Platform console, you can configure the DNS server settings for the data center. These settings are used for performing DNS updates at the time of application migration between the data centers for application disaster recovery. They must be configured on both data centers.

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

Note: If you add an IMS to an existing data center after you add a DNS server 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 test confirms that the newly added IMS can be used to perform DNS updates.

To configure DNS server settings for a data center

1 Prerequisites

You must have the following information:

- The IP address of the DNS server
- The name of the domain, and associated credentials. For TSIG authentication, you need the TSIG key and TSIG private files. For GSSAPI authentication, you need the user name and keytab file.
- A test host name and IP address for performing a test operation by updating the DNS server for validating the specified DNS configuration.

2 Navigate

⚙ **Settings** (menu bar)

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

⋮ Click the vertical ellipsis next to the data center name, then click **DNS & Network Settings** > **DNS** tab.

Any DNS servers already added for the data center are listed. You can modify or remove them, or add a new DNS server.

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

Or select **Add Existing DNS Server** if the DNS server has already been added to Resiliency Platform for another data center.

4. Specify the IP address for the DNS server.
5. Add one or more domains to the DNS server:
   - Fill in the domain name and the authentication type. For TSIG, browse to the key and private files. For GSSAPI, enter the user name and browse to the keytab file.
   - Enter a test host name and IP address and click **Test**. If the test is successful, the **Add** button is enabled.
   - Click **Add** to add the domain to the DNS server.
6. If you are done adding domains, click **Save**.
7. To modify or remove a DNS server, go to the **DNS** tab and select **Modify** or **Remove** for that DNS server.

See “Configuring subnet information for a data center” on page 52.

See “Managing data centers” on page 48.

Configuring subnet information for a data center

To configure subnet information for a data center

1. Navigate

   ⚙ **Settings** (menu bar)

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

   ⋮ Click the vertical ellipsis next to the data center name, then click **DNS & Network Settings** > **Subnets** tab.

   Any subnets already added for the data center are listed. You can modify or remove them, or add a new subnet.

2. To add a new subnet, click **Add** and specify the IP address for the subnet and gateway. Optionally, select the virtualization servers that are part of the subnet.
3. Click **Add** at the bottom of the form.

See “Configuring DNS server settings for a data center” on page 51.

See “Managing data centers” on page 48.

Chapter 7

Adding the asset infrastructure to an Infrastructure Management Server

This chapter includes the following topics:

- Managing solutions for the hosts

About the asset infrastructure

The data center assets that you add to the Infrastructure Management Server (IMS) for IMS discovery and monitoring are referred to as the asset infrastructure.

The asset infrastructure can include hosts (Windows/Linux servers) and virtualization servers for Hyper-V and VMware. For configuring disaster recovery, if you are using enclosures (storage arrays) for replication, you also add the information about the storage arrays to the IMS. See the following use cases for the recommended sequence in which to add assets. 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.

See the following for more information on asset infrastructure requirements for the different use cases of virtual machines and applications.

Table 7-1 Use cases for adding asset infrastructure

Use case	More information
Virtual machine discovery in Hyper-V environment	See "Configuring IMS asset infrastructure for Hyper-V virtual machines" on page 55.
Virtual machine discovery in VMware environment	See "Configuring IMS asset infrastructure for VMware virtual machines" on page 55.
Application discovery on Hyper-V virtual machines	See "Configuring IMS asset infrastructure for application discovery on Hyper-V virtual machines" on page 57.
Application discovery on VMware virtual machines	See "Configuring IMS asset infrastructure for application discovery on VMware virtual machines" on page 58.
Application discovery on physical systems	See "Configuring IMS asset infrastructure for application discovery on physical systems" on page 60.

See "Adding the asset infrastructure to an Infrastructure Management Server (IMS)" on page 62.

Configuring IMS asset infrastructure for Hyper-V virtual machines

The following table describes the asset infrastructure to be added to the Infrastructure Management Server (IMS) for discovery and management of Hyper-V virtual machines and for discovery of storage arrays used for replication for disaster recovery.

Table 7-2 Configuring IMS asset infrastructure for Hyper-V virtual machines

Task	Details	IMS wizard
Add the Hyper-V virtualization server to the IMS	See “Managing Hyper-V assets for an IMS” on page 69. See “Managing host assets for an IMS” on page 63.	Add Hosts
For EMC SRDF replication: Configure a discovery host for the storage array Not applicable for Hyper-V Replica	Configure a discovery host with the vendor-specific array management tools and connect it to the enclosure See “About the discovery host” on page 80. See “Configuration prerequisites for adding enclosures to an IMS” on page 80.	not applicable
Add the array discovery host to the IMS	See “Managing host assets for an IMS” on page 63.	Add Hosts
Add the storage array to the IMS	See “Managing enclosure assets for an IMS” on page 79.	Add Enclosures

Configuring IMS asset infrastructure for VMware virtual machines

The following tables describe the asset infrastructure to be added to the Infrastructure Management Server (IMS) for discovery and management of VMware virtual machines. They include information on adding the storage arrays used for replication for disaster recovery.

Note: The tables also show the recommended sequence for adding the assets so that the IMS completes discovery most quickly.

Refer to the appropriate table depending on whether you are using EMC SRDF or NetApp SnapMirror replication.

Table 7-3 Configuring IMS asset infrastructure for VMware virtual machines (EMC SRDF replication)

Task	Details	IMS wizard
For EMC SRDF replication:Configure a discovery host for the storage array to be used for replication	Configure a discovery host with the vendor-specific array management tools and connect it to the enclosure See "About the discovery host" on page 80. See "EMC Symmetrix configuration prerequisites" on page 80.	not applicable
Add the storage array discovery host to the IMS	See "Managing host assets for an IMS" on page 63.	Add Hosts
Add the storage array to the IMS	See "Managing enclosure assets for an IMS" on page 79.	Add Enclosures
Add the VMware vCenter Server to the IMS	If you added a storage array for replication, for IMS discovery to proceed most quickly, ensure that the enclosure discovery is complete before you add the vCenter Server to the IMS. See "Managing VMware virtualization assets for an IMS" on page 70.	Add Virtualization Server

Table 7-4 Configuring IMS asset infrastructure for VMware virtual machines (NetApp SnapMirror replication)

Task	Details	IMS wizard
Add the VMware vCenter Server to the IMS	See "Managing VMware virtualization assets for an IMS" on page 70.	Add Virtualization Server

Table 7-4 Configuring IMS asset infrastructure for VMware virtual machines (NetApp SnapMirror replication) *(continued)*

Task	Details	IMS wizard
For NetApp SnapMirror replication: Add the storage array to the IMS	For IMS discovery to proceed most quickly, ensure that the vCenter Server discovery is complete before you add the storage array to the IMS. When you enter the array information in the IMS wizard, you must specify the IMS as the discovery host. See "Managing enclosure assets for an IMS" on page 79.	Add Enclosures

Configuring IMS asset infrastructure for application discovery on Hyper-V virtual machines

The following table describes the asset infrastructure to be added to the Infrastructure Management Server (IMS) for discovery and management of applications on Hyper-V virtual machines. It includes information on adding the storage arrays used for replication for disaster recovery.

Table 7-5 Configuring IMS asset infrastructure for application discovery on Hyper-V virtual machines

Task	Details	IMS wizard
Add the Hyper-V virtualization server to the IMS	See "Managing Hyper-V assets for an IMS" on page 69. See "Managing host assets for an IMS" on page 63.	Add Hosts

Table 7-5 Configuring IMS asset infrastructure for application discovery on Hyper-V virtual machines *(continued)*

Task	Details	IMS wizard
Add the Hyper-V virtual machines to the IMS as hosts	See “Managing host assets for an IMS” on page 63. **Note:** Discovery of custom applications requires an additional step after adding the hosts. You must also add the applications on the Assets page. More information is available on adding custom applications. See *Veritas Resiliency Platform Solutions for Applications*.	Add Hosts
For array-based replication, configure a discovery host for the storage array	Configure a discovery host with the vendor-specific array management tools and connect it to the enclosure See “About the discovery host” on page 80. See “Configuration prerequisites for adding enclosures to an IMS” on page 80.	not applicable
Add the storage array discovery host to the IMS	See “Managing host assets for an IMS” on page 63.	Add Hosts
Add the storage array to the IMS	See “Managing enclosure assets for an IMS” on page 79.	Add Enclosures

Configuring IMS asset infrastructure for application discovery on VMware virtual machines

The following table describes the asset infrastructure to be added to the Infrastructure Management Server (IMS) for discovery and management of applications on VMware virtual machines. It includes information on adding the storage arrays used for replication for disaster recovery.

Table 7-6 Configuring IMS asset infrastructure for application discovery on VMware virtual machines (EMC SRDF replication)

Task	Details	IMS wizard
For EMC SRDF replication: Configure a discovery host for the storage array to be used for replication	Configure a discovery host with the vendor-specific array management tools and connect it to the enclosure See "About the discovery host" on page 80. See "EMC Symmetrix configuration prerequisites" on page 80.	not applicable
Add the storage array discovery host to the IMS	See "Managing host assets for an IMS" on page 63.	Add Hosts
Add the storage array to the IMS	See "Managing enclosure assets for an IMS" on page 79.	Add Enclosures
Add the VMware vCenter Server to the IMS	See "Managing VMware virtualization assets for an IMS" on page 70.	Add Virtualization Server
Add the VMware virtual machines to the IMS as hosts	See "Managing host assets for an IMS" on page 63. **Note:** Discovery of custom applications requires an additional step after adding the hosts. You must also add the applications on the Assets page. More information is available on adding custom applications. See *Veritas Resiliency Platform Solutions for Applications*.	Add Hosts

Table 7-7 Configuring IMS asset infrastructure for VMware virtual machines (NetApp SnapMirror replication)

Task	Details	IMS wizard
Add the VMware vCenter Server to the IMS	See "Managing VMware virtualization assets for an IMS" on page 70.	Add Virtualization Server

Table 7-7 Configuring IMS asset infrastructure for VMware virtual machines (NetApp SnapMirror replication) *(continued)*

Task	Details	IMS wizard
Add the VMware virtual machines to the IMS as hosts	See "Managing host assets for an IMS" on page 63. **Note:** Discovery of custom applications requires an additional step after adding the hosts. You must also add the applications on the Assets page. More information is available on adding custom applications. See *Veritas Resiliency Platform Solutions for Applications*.	Add Hosts
For NetApp SnapMirror replication: Add the storage array to the IMS	For IMS discovery to proceed most quickly, ensure that the vCenter Server discovery and virtual machine discovery is complete before you add the storage array to the IMS. When you enter the array information in the IMS wizard, you must specify the IMS as the discovery host. See "Managing enclosure assets for an IMS" on page 79.	Add Enclosures

Configuring IMS asset infrastructure for application discovery on physical systems

The following table describes the asset infrastructure to be added to the Infrastructure Management Server (IMS) for discovery and management of applications on physical systems. It includes information on adding the storage arrays used for replication for disaster recovery.

Table 7-8 Configuring IMS asset infrastructure for application discovery on physical systems (EMC SRDF replication)

Task	Details	IMS wizard
Add the physical servers on which the applications are running to the IMS	See "Managing host assets for an IMS" on page 63. **Note:** Discovery of custom applications requires an additional step after adding the hosts. You must also add the applications on the Assets page. More information is available on adding custom applications. See *Veritas Resiliency Platform Solutions for Applications*.	Add Hosts
For EMC SRDF replication: Configure a discovery host for the storage array to be used for replication	Configure a discovery host with the vendor-specific array management tools and connect it to the enclosure See "About the discovery host" on page 80. See "Configuration prerequisites for adding enclosures to an IMS" on page 80.	not applicable
Add the array discovery host to the IMS	See "Managing host assets for an IMS" on page 63.	Add Hosts
Add the storage array to the IMS	See "Managing enclosure assets for an IMS" on page 79.	Add Enclosures

Table 7-9 Configuring IMS asset infrastructure for application discovery on physical systems (NetApp SnapMirror replication)

Task	Details	IMS wizard
Add the physical servers on which the applications are running to the IMS	See “Managing host assets for an IMS” on page 63. **Note:** Discovery of custom applications requires an additional step after adding the hosts. You must also add the applications on the Assets page. More information is available on adding custom applications. See *Veritas Resiliency Platform Solutions for Applications*.	Add Hosts
For NetApp SnapMirror replication: Add the storage array to the IMS	When adding or refreshing a configuration, ensure that host discovery is complete before adding/refreshing the NetApp enclosure. When you enter the array information in the IMS wizard, you must specify the IMS as the discovery host. See “Managing enclosure assets for an IMS” on page 79.	Add Enclosures

Adding the asset infrastructure to an Infrastructure Management Server (IMS)

After you add an Infrastructure Management Server (IMS) to the resiliency domain, you add the asset infrastructure to the IMS.

The asset infrastructure can include hosts, virtualization servers, and enclosures (storage arrays used for replication). The assets you must add depend on the use case and environment.

See “About the asset infrastructure” on page 54.

To add the asset infrastructure to an IMS

1. Prerequisites

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

2. Navigate

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

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

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

   The **Settings** page for the IMS is displayed. You can add, edit, or remove assets.

   Tip: You must add Hyper-V servers as hosts rather than as virtualization servers.

   See "Managing host assets for an IMS" on page 63.

   See "Managing Hyper-V assets for an IMS" on page 69.

   See "Managing VMware virtualization assets for an IMS" on page 70.

   See "Managing enclosure assets for an IMS" on page 79.

Managing host assets for an IMS

When adding the asset infrastructure to an Infrastructure Management Server (IMS), several different use cases require that you add hosts.

See "About adding host assets to an IMS" on page 63.

See "Adding hosts to an IMS" on page 66.

See "About using a CSV file for adding hosts to an IMS" on page 66.

See "Removing hosts from an IMS " on page 68.

See "Refreshing host discovery information for an IMS" on page 68.

About adding host assets to an IMS

Host assets that you add to an Infrastructure Management Server (IMS) can include physical systems, virtual machines, and Hyper-V servers, depending on the use case, as described in the table.

Table 7-10 Use cases for adding host assets to an IMS

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*.
Hyper-V virtualization server	Hyper-V servers used for virtualization must be added as host assets, rather than as virtualization servers.
Hardware replication	For storage array-based replication, you may need to install array-specific software on a host and add the host to the IMS to use as a discovery host. See "About the discovery host" on page 80. More information is available on requirements for adding enclosures for array-based replication. See "Managing enclosure assets for an IMS" on page 79.

When you add hosts to an IMS, 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. See the prerequisite in the following list for Windows hosts.

The following are prerequisites for adding host assets to an IMS:

- Ensure that the IMS can communicate with the host.
  See "Network and firewall requirements" on page 24.
- 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.
  See "About using a CSV file for adding hosts to an IMS" on page 66.

- For Linux hosts, in order to install the host package while adding the host, ensure that the PasswordAuthentication field is set to **yes** in the `/etc/ssh/sshd_config file` on the host.
- For Windows systems, you must manually install the VRTSsfmh host package on one Windows host before you can add the Windows host to the IMS using the web console. You can then add any remaining Windows hosts using the console, and the IMS installs the host package on the subsequent Windows hosts.
  See "Installing the host package on a Windows host" on page 65.
- For installing the host package while adding a Windows host, ensure that the User Access Control (UAC) is disabled on the host. The host package installation will not work if UAC is enabled on the host.

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

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

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 appliance" on page 28.
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 for an IMS” on page 63.

About using a CSV file for adding hosts to an IMS

When adding hosts to an Infrastructure Management Server (IMS), you have the option to import the information from a comma-separated (.csv) file. The CSV file must include the “.csv” extension. The following is an example of a CSV file:

```
Host,User,Password
host1.abc.com,username1,password1
host2.abc.com,username2,password2
```

The first line in the CSV file must appear as above. You replace the subsequent lines with your hosts, user names, and passwords.

See “Adding hosts to an IMS” on page 66.

Adding hosts to an IMS

After adding an Infrastructure Management Server (IMS) to the resiliency domain, you can add host assets to the IMS.

To add hosts to an IMS

1. Prerequisites

   For use cases and prerequisites for adding hosts to an IMS, see the following topic.

   See “About adding host assets to an IMS” on page 63.

2. Navigate

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

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

3. Under the data center, locate the IMS and click **Manage Asset Infrastructure**.
4. On the IMS **Settings** page, click **Host** to view information on already added hosts, then click **Add Hosts**.

   Or to go directly to the Add Hosts wizard, click **Add Hosts** in the **Settings** page menu bar.

5 In the wizard, 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 the host package is already present on the host that is being added, select **None**.
- If you select **Install managed host package on Linux/Unix**, the **Use root password** option is enabled. Select this option if you want to install the host package on a Linux/Unix host as a non-root user. Provide the non-root username, non-root password, and root password for the specified host. You can use this option if the Secure Shell (SSH) access is disabled for the root login on the host where you want to install the host package.
- Before you can add the first Windows host to the IMS, you must manually install the host package on the host. Then add the host using the **None** option in this wizard.

  See “Installing the host package on a Windows host” on page 65.

  After the first Windows host is added, to add more Windows hosts, run the wizard again and select **Install managed host package on Windows**. Then, for **Select Windows Managed Host**, select the host added previously. If there are multiple Windows hosts listed, you can select any one.

6 Choose from the following methods of adding a host:

- Type the information on the table row. To add a blank table row, click **Add**.
- Click **Clone** to clone the selected table row, and then edit the clone.
- To import the information from a CSV file, click **Import**.

7 Verify that the host has been added successfully.

8 In the **Recent tasks** pane, verify that the **Install add-on** tasks for the **Veritas Resiliency Platform Enablement** add-on, and **Veritas Resiliency Platform Applications Enablement** add-on are successfully completed on the host.

If the add-ons are not successfully installed, then you need to manually install them on the host.

See “Installing add-ons on the hosts” on page 94.

If installation of the **Veritas Resiliency Platform Applications Enablement** add-on fails and an error message to run the cleanup script is displayed, you need to run the following script on the managed host:

```
C:\ProgramData\Symantec\VRTSsfmh\spool\addons\store
\VRTSitrpappdr-1.0.0.0\CleanUp_Scripts\cleanup.bat
```

See “Managing host assets for an IMS” on page 63.

Removing hosts from an IMS

You can remove one or more hosts that were added to an Infrastructure Management Server (IMS).

Before removing a host, you need to uninstall all the add-ons that were installed on the host.

See “Uninstalling add-ons from the hosts” on page 96.

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.

To remove hosts from an IMS

1. Navigate

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

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 **Remove**.

   To remove more than one host, hold down **Ctrl** as you select hosts from the list.

5. Confirm that you want to remove the host.

See “Managing host assets for an IMS” on page 63.

Refreshing host discovery information for an IMS

You can submit a refresh request to update the information displayed for the hosts that have been added to the Infrastructure Management Server (IMS). Once the refresh operation is complete, the Assets page in the console is also updated.

To refresh a host discovery for the IMS

1 Navigate

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

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 (press CRTL to select multiple hosts) 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 for an IMS” on page 63.

Managing Hyper-V assets for an IMS

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

See “About Microsoft Hyper-V virtualization discovery” on page 69.

See “Prerequisites for Microsoft Hyper-V virtualization discovery by the IMS” on page 70.

See “Adding hosts to an IMS” on page 66.

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 for an IMS" on page 69.

Prerequisites for Microsoft Hyper-V virtualization discovery by the IMS

You can add Microsoft Hyper-V hosts to an Infrastructure Management Server (IMS) for virtualization discovery.

Note: When adding the Hyper-V hosts to the IMS in the console, you choose the option to add hosts rather than to add virtualization servers.

See "Managing Hyper-V assets for an IMS" on page 69.

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

Table 7-11 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 the IMS as a host in the Resiliency Platform console. ■ The Hyper-V role must be enabled. ■ The Windows Management Instrumentation (WMI) service must be running.
Exported storage discovery	■ The Windows Management Instrumentation (WMI) service must be running on the guest.

See "Adding hosts to an IMS" on page 66.

Managing VMware virtualization assets for an IMS

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

The VMware discovery provides the following information:

- Information on vCenter servers
- Information on the ESX servers that the vCenter server manages
- Information on the virtual machines that are configured on the ESX servers

See "Prerequisites for adding VMware servers for discovery by the IMS" on page 71.

See "Adding virtualization servers for VMware discovery by the IMS" on page 76.

See "Editing a VMware virtualization discovery configuration in the IMS" on page 77.

See "Removing a virtualization discovery configuration from the IMS" on page 78.

See "Refreshing a VMware virtualization discovery configuration" on page 78.

See "Refreshing an ESX Server discovery" on page 79.

Prerequisites for adding VMware servers for discovery by the IMS

Ensure that the following requirements are met to add the VMware vCenter or ESX servers to the Infrastructure Management Server (IMS):

- Ensure that the IMS server can ping the vCenter servers or the ESX servers from which it can discover the information on VMware Infrastructure.
- Ensure that you have appropriate privileges to log on to the vCenter server or the ESX server. You must provide login credentials to add the servers to the IMS.
- Ensure that you have the Browse Datastore privileges on the vCenter or the ESX server that you want to discover.
- Ensure that you have configured near real-time discovery of VMware events.
  See "About near real-time discovery of VMware events" on page 71.
  See "Setting up near real-time discovery of VMware events" on page 73.
  See "Configuring the VMware vCenter Server to generate SNMP traps" on page 74.

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).

To set up the near real-time discovery of VMware events by the IMS, you must configure the vCenter Server to generate SNMP traps and send them to the IMS address. The recommended sequence is to do this before adding the vCenter Server to the IMS.

See "Setting up near real-time discovery of VMware events" on page 73.

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 7-12 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 7-12 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 7-13 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 74.
Using the Resiliency Platform console:		

Table 7-13 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.	See "Adding virtualization servers for VMware discovery by the IMS" on page 76. 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. See "Refreshing a VMware virtualization discovery configuration" on page 78.

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.

See "About near real-time discovery of VMware events" on page 71.

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.

Field	Description
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 a new action to send a notification trap. Specify to send the notification trap as in the following example:

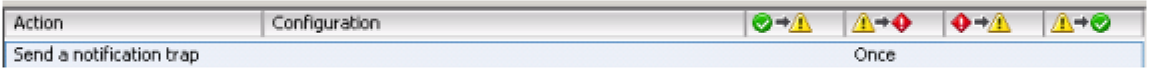

Action	Configuration	✓→⚠	⚠→◆	◆→⚠	⚠→✓
Send a notification trap			Once		

See “About near real-time discovery of VMware events” on page 71.

See “Setting up near real-time discovery of VMware events” on page 73.

Adding virtualization servers for VMware discovery by the IMS

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

The VMware discovery provides the following information:

- Information on vCenter servers
- Information on the ESX servers that the vCenter server manages
- Information on the virtual machines that are configured on the ESX servers

To add a virtualization server for VMware discovery by the IMS

1. Prerequisites:

   See “Prerequisites for adding VMware servers for discovery by the IMS” on page 71.

2. Navigate

   Settings (menu bar) > **Infrastructure**

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

3. Under the data center, locate the IMS and click **Manage Asset Infrastructure**.
4. On the IMS **Settings** page, click **Virtualization** to view information on existing virtualization servers, then click **Add Virtualization Server**.

   Or to go directly to the Add Virtualization Server wizard, click **Add Virtualization Server** in the **Settings** page menu bar.

5. In the wizard, specify the information for the vCenter server and click **Next**.

   Tips:

   - For Configuration Name, specify a name that will help you identify this configuration

- Specify the fully-qualified name of the vCenter that you want to discover along with its port number, separated by a colon. If the vCenter Web service is running on a default port, you do not need to specify the port number.
- When entering the login credentials, you can use a read-only user account if it has the Browse Datastore permissions on the virtualization servers.
- When entering the username, you must enter in the format username@domainname, not domainname\username.

6. Choose to automatically discover all ESX servers or manually specify names of ESX servers to discover. Click **Finish**.
7. In the **Result** panel, view the progress of the configuration. After the configuration is complete, click **OK**.
8. After you add a vCenter server, to view all ESX servers that the vCenter server manages, click **vCenter** under **Data Center**.

If any changes are made on the virtualization server after adding it to the IMS, you need to refresh the server discovery configuration.

See “Refreshing a VMware virtualization discovery configuration” on page 78.

Editing a VMware virtualization discovery configuration in the IMS

You can edit a virtualization discovery configuration in the Infrastructure Management Server (IMS) to modify the following information:

- Name of the configuration.
- Credentials to log on to the vCenter.
  When entering the username, you must enter in the format username@domainname, not domainname\username.

To edit a virtualization discovery configuration in the IMS

1. Navigate

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

2. Under the data center, locate the IMS and click **Manage Asset Infrastructure**.
3. On the IMS **Settings** page, click **Virtualization**.
4. In the **Virtualization Configurations** details list, right-click the configuration that you want to edit.
5. In the **Edit Configuration** wizard panel, modify the required information, click **Next**.

6. In the **Edit Configuration** wizard panel, edit the method for virtualization discovery of the servers, click **Finish**.
7. In the **Result** panel, view the progress of the configuration, click **OK**.

See “Adding virtualization servers for VMware discovery by the IMS” on page 76.

Removing a virtualization discovery configuration from the IMS

To remove a virtualization discovery configuration from the IMS

1. Navigate

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

2. Under the data center, locate the IMS and click **Manage Asset Infrastructure**.
3. On the IMS **Settings** page, click **Virtualization**.
4. Right-click the virtualization server and select **Remove Configuration**.
5. In the **Remove Virtualization Configuration** wizard panel, click **Remove**.
6. In the **Result** panel, click **OK**.

Refreshing a VMware virtualization discovery configuration

You can submit a refresh request to update the information displayed on the table of virtualization servers that have been added to the Infrastructure Management Server (IMS).

To refresh a virtualization discovery configuration

1. Navigate

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

2. Under the data center, locate the IMS and click **Manage Asset Infrastructure**.
3. On the IMS **Settings** page, click **Virtualization**.
4. Right-click the virtualization configuration and select **Refresh Configuration**.
5. In the **Refresh Virtualization Configuration** wizard panel, click **Refresh**.
6. In the **Result** panel, click **OK**.

Refreshing an ESX Server discovery

You can refresh the Infrastructure Management Server (IMS) discovery of one or more ESX servers that are configured under a selected VMware vCenter Server.

To refresh the discovery of an ESX server

1. Navigate

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

2. Under the data center, locate the IMS and click **Manage Asset Infrastructure**.
3. On the IMS **Settings** page, click **Virtualization**.
4. Under the **Virtualization Configurations** tab, you can view the details of virtualization configuration. For example, the name of the virtualization server (vCenter Server) used in the configuration, its type, and other parameters. Select the desired virtualization configuration.
5. The **Configured Virtualization Servers** tab lists the ESX servers managed under the vCenter Server that is part of the selected virtualization configuration.
6. Right-click the required ESX server and click **Refresh.** Press Ctrl or Shift for the selection of multiple ESX servers.
7. In the **Refresh Virtualization Server** wizard panel, click **Refresh**.
8. In the **Result** panel, click **OK**.

Managing enclosure assets for an IMS

You can add storage enclosures (arrays) to an Infrastructure Management Server (IMS) for discovery of storage information to monitor array-based replication.

This does not apply for environments that are using Hyper-V Replica.

See “About the discovery host” on page 80.

See “Configuration prerequisites for adding enclosures to an IMS” on page 80.

See “Adding storage enclosures for discovery by the IMS” on page 85.

See “Editing the discovery configuration for an enclosure” on page 90.

See “Removing the discovery configuration for an enclosure” on page 93.

See “EMC Symmetrix storage enclosure commands” on page 170.

See “NetApp storage enclosure commands” on page 172.

About the discovery host

A discovery host is a Windows or Linux host on which are installed vendor-specific array management tools that the Infrastructure Management Server (IMS) uses for discovery and monitoring the enclosure (storage array). When you add the enclosure to the IMS, you specify a discovery host in the Add Enclosure wizard. The discovery host must also be added to the IMS.

Note: In the case of NetApp, you specify the IMS as a discovery host when you add the enclosure. You do not add a separate discovery host.

See "Adding hosts to an IMS" on page 66.

See "EMC Symmetrix configuration prerequisites" on page 80.

See "NetApp configuration prerequisites" on page 85.

See "Managing enclosure assets for an IMS" on page 79.

Configuration prerequisites for adding 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.

See "About the discovery host" on page 80.

See "EMC Symmetrix configuration prerequisites" on page 80.

See "NetApp configuration prerequisites" on page 85.

See "Adding the asset infrastructure to an Infrastructure Management Server (IMS)" on page 62.

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 81.

See “Verifying the configuration of a remote SYMAPI server” on page 84.

See “Managing enclosure assets for an IMS” on page 79.

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 80.

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 80.

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.

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 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: httpd.admin.enable and httpd.enable. These are required for NetApp SnapMirror operations.

Ensure that following licenses are installed and enabled: licensed_feature.multistore.enable (required for discovering IP addresses) and licensed_feature.flex_clone.enable (required for rehearsal operation).

Adding storage enclosures for discovery by the IMS

After adding Infrastructure Management Server (IMS), you add the asset infrastructure. For array-based replication environments, the asset infrastructure includes the storage enclosures. This does not apply for environments that are using Hyper-V Replica.

To add storage enclosures for discovery

1 Prerequisites

   Ensure that you have configured the storage array for discovery.

   See "Configuration prerequisites for adding enclosures to an IMS" on page 80.

   Ensure that you have the name of the discovery host.

   See "About the discovery host" on page 80.

2 Navigate

   Settings (menu bar) > **Infrastructure**

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

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

4 On the IMS **Settings** page, click **Device** to view information on existing enclosures, then click **Add Enclosure**.

   Or to go directly to the Add Enclosure wizard, click **Add Enclosure** in the **Settings** page menu bar.

5 In the **Add Enclosure** wizard, select the vendor and the enclosure model. Click **Next.**

   See "Add Enclosure panel options for vendor and product selection" on page 87.

6 Specify the discovery details and click **Next**.

   See "Add Enclosure panel options for configuration details" on page 87.

7 Choose the enclosures and enable the deep array discovery. Click **Finish**.

   See "Add Enclosure panel options to enable or disable discovery for selected enclosures" on page 89.

8 When the enclosure is successfully added, you can verify the information on the **Enclosure Configurations** tab.

See "Managing enclosure assets for an IMS" on page 79.

Add Enclosure panel options for vendor and product selection

Table 7-14 Add Enclosure panel for vendor and product selection

Field	Description
Enclosure Vendor	Select the enclosure vendor from the drop-down list.
Select product with appropriate discovery method	
Product	Select the array model for which you want to enable the deep array discovery.
Discovery Method	Displays the discovery method that is used for the discovery of the selected array model.
Additional Information	Provides the information about the discovered objects, prerequisites, CLI version, and other details about the enclosure discovery.

See “Adding storage enclosures for discovery by the IMS” on page 85.

Add Enclosure panel options for configuration details

Use this wizard panel to specify the details of the devices and the server information for adding the following enclosures for deep discovery.

Table 7-15 lists the options for EMC Symmetrix enclosures.

Table 7-16 lists the options for NetApp enclosures.

Table 7-15 Add Enclosure panel options for EMC Symmetrix enclosure

Field	Description
Discovery Host	Enter the name of the discovery host. A host that runs on Linux or Windows can be designated as a discovery host. The discovery host must be added as a host to the IMS.
SYMAPI Server	Specify the SYMAPI Server name that is configured on the discovery host to discover the EMC Symmetrix enclosures. Use this option if discovery host does not have visibility to gatekeeper devices for Symmetrix enclosures.

Table 7-15 Add Enclosure panel options for EMC Symmetrix enclosure *(continued)*

Field	Description
SymCLI Location	
Use Default	Choose this option if you have stored the SymCLI binaries on the default location. 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 80.
Custom	Choose this option if you have stored the SymCLI binaries on any other location. Enter the path to the location in the field. You must ensure that the SymCLI binaries are available on the discovery host.
Enable performance metering	Not applicable for Veritas Resiliency Platform. Clear the check box.

Table 7-16 Add Enclosure panel options for NetApp enclosures

Discovery Host	Specify the name of the IMS.
NetApp Server	Enter the name or the IP address for the NetApp server.
Port	The port for the NetApp server. Enter 80 for communicating to the NetApp server over HTTP. For communicating to the NetApp server over HTTPS, enter 443. Ensure the port that you specify here is enabled on the enclosure.
Username	Enter the user name for the enclosure.
Password	Enter the password for the enclosure.
Enable NAS discovery	Required. Select the check box..

Table 7-16 Add Enclosure panel options for NetApp enclosures *(continued)*

Enable performance metering	Not applicable for Veritas Resiliency Platform. Clear the check box.

See “Adding storage enclosures for discovery by the IMS” on page 85.

Add Enclosure panel options to enable or disable discovery for selected enclosures

Use this wizard panel to enable or disable deep discovery for selected enclosures.

Select the check box in the top row to select all the enclosures in the list. The check box is selected by default.

For **Discovery**, choose **Enable** to perform the discovery; otherwise choose **Disable**.

Table 7-17 Add Enclosure panel to update the deep discovery information for enclosures

Field	Description
Configuration Name	Enter a name for the deep discovery operation that you want to perform.
Enclosures	
Display Name	Displays the name of the 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.

Table 7-17 Add Enclosure panel to update the deep discovery information for enclosures *(continued)*

Field	Description
Connectivity	This field is displayed only for the EMC Symmetrix enclosures. It indicates the following: ▪ **Local**: Whether the enclosure is connected to the host locally. ▪ **Remote**: Whether the enclosure is connected to another enclosure, using the EMC Symmetrix Remote Data Facility (SRDF), which might have been connected to the host locally.
Discovery	Choose **Enable** to perform the deep discovery. Choose **Disable** to disable the deep discovery.
Configured Name	If the enclosure was already add for deep discovery, the configuration name that was entered at that time is displayed in this field.

See “Adding storage enclosures for discovery by the IMS” on page 85.

Editing the discovery configuration for an enclosure

You can edit details for enclosure configurations that were added previously to the Infrastructure Management Server (IMS).

To edit the discovery configuration for an enclosure

1 Navigate

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

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

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

4 Expand **Enclosures** to locate the vendor.

5 In the vendor configurations list, right-click the enclosure and select **Edit Configuration**.

6 In the **Edit Configuration** wizard panel, edit the configuration details to change the device discovery. Click **Next**.

See “Edit Configuration panel options to modify the device discovery” on page 91.

7 In the next panel, select the check box for the enclosures for which you want to perform the deep discovery configuration. Click **Finish**.

8 In the result panel review the information and click **OK**.

See “Managing enclosure assets for an IMS” on page 79.

Edit Configuration panel options to modify the device discovery

Use this wizard panel to edit the configuration for deep discovery for the following enclosures:

Table 7-18 lists the options for EMC Symmetrix enclosures.

Table 7-19 lists the options for NetApp enclosures.

Table 7-18 Edit configuration panel options for EMC Symmetrix enclosure

Field	Description
Enclosure Vendor	Displays the vendor name of the enclosure.
Product	Displays the array model for which the deep array discovery is enabled.
Discovery Method	Displays the discovery method that is used for the discovery of the selected array model.
Discovery Host	Modify the name of the host. This host must be a part of the Management Server domain.
SYMAPI Server	Specify the SYMAPI Server name that is configured on the discovery host to discovery EMC Symmetrix enclosures. Use this option if discovery host does not have visibility to gatekeeper devices for Symmetrix enclosures.
SymCLI Location	

Table 7-18 Edit configuration panel options for EMC Symmetrix enclosure *(continued)*

Field	Description
Use Default	Choose this option if you have stored the SymCLI binaries on the default location. You must ensure that the SymCLI binaries are available on the discovery host.
Custom	Choose this option if you have stored the SymCLI binaries on any other location. Enter the path to the location in the field. You must ensure that the SymCLI binaries are available on the discovery host.
Enable performance metering	Enables the performance metering for the enclosure. Clear the check box to disable performance metering.

Table 7-19 Edit configuration panel options for NetApp enclosures

Field	Description
Enclosure Vendor	Displays the vendor name of the enclosure.
Product	Displays the array model for which the deep array discovery is enabled.
Discovery Method	Displays the discovery method that is used for the discovery of the selected array model.
Discovery Host	Modify the name of the host. This host must be a part of the Management Server domain.
NetApp Server	Modify the name or the IP address for the NetApp server.
Port	Modify the port for the NetApp server. Enter 80 for communicating to the NetApp server over HTTP. For communicating to the NetApp server over HTTPS, enter 443. Ensure the port that you specify here is enabled on the enclosure.
Username	Modify the user name for the enclosure.

Table 7-19 Edit configuration panel options for NetApp enclosures *(continued)*

Field	Description
Password	Modify the password for the enclosure.
Enable NAS discovery	Select the check box to enable the NAS discovery for the NetApp enclosure.
Enable performance metering	Enables the performance metering for the enclosure. Clear the check box to disable performance metering.

Removing the discovery configuration for an enclosure

You can remove enclosure configurations that were added previously to the Infrastructure Management Server (IMS).

To remove the discovery configuration for an enclosure

1. Navigate

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

2. Under the data center, locate the IMS and click **Manage Asset Infrastructure**.
3. On the IMS **Settings** page, click **Device**.
4. Expand **Enclosures** to locate the vendor.
5. In the vendor configurations list, right-click the enclosure and select **Remove Configuration**.
6. In the **Unconfigure Device** panel, click **Yes**.

See “Managing enclosure assets for an IMS” on page 79.

Refreshing enclosure discovery information for an IMS

You can submit a refresh request to update the information displayed on the table of enclosures assets that have been added to the Infrastructure Management Server (IMS).

To refresh an enclosure configuration discovery

1 Navigate

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

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

3 On the IMS **Settings** page, click **Device** and click the **Enclosure Configurations** tab.

4 To refresh the configuration, right click and select **Refresh Configuration**. To refresh an enclosure, right-click a configured enclosure and select **Refresh Enclosure**.

5 Click **OK**.

Managing add-ons for the hosts

See “Installing add-ons on the hosts” on page 94.

Installing add-ons on the hosts

You can install the add-ons on the hosts that are added to the Infrastructure Management Server (IMS).

To install add-ons on the hosts

1 Navigate

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

(Or, click **Quick Actions** (menu bar) > **Add Assets**)

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

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

4 Expand **Add-ons** to select the add-on that you want to install.

5 In the **Add-ons** tab, right-click the add-on, and select **Install**.

6 In the **Install - Selects hosts** wizard panel, select the hosts, and click **Finish**.

See “Install - Select hosts panel options for add-ons” on page 95.

7 In the **Result** panel, click **Close**

8 Those add-ons which require web server restart, click **Restart Web server**.

If installation of the **Veritas Resiliency Platform Applications Enablement** add-on fails and an error message to run the cleanup script is displayed, you need to run the following script on the managed host:

```
C:\ProgramData\Symantec\VRTSsfmh\spool\addons\store
\VRTSitrpappdr-1.0.0.0\CleanUp_Scripts\cleanup.bat
```

Install - Select hosts panel options for add-ons

Use this wizard panel to select the hosts on which you want to install the add-on.

You can do one of the following:

- Select the hosts explicitly and install the add-on on the selected hosts.
- Select the platform.

If you select a specific platform, the add-on is installed on all the hosts using that platform. Also the add-on is installed on all the new hosts that are added to the IMS in the future.

For example, if you select Windows, the add-on is installed on all the hosts that use Windows platform. Also when a new Windows host is added to the IMS, the add-on is installed on the host.

Table 7-20 Select hosts panel options

Field	Description
Hosts	Select to view the list of all the hosts where the add-on is not installed. Select **Show all applicable hosts (Overwrites existing installation)** to list all the hosts on which you can install the add-on. It includes: ■ Hosts on which the add-on is not installed currently. ■ Hosts on which the add-on is installed currently. In this case, the existing add-on installation is overwritten.

Table 7-20 Select hosts panel options *(continued)*

Field	Description
Platform	Select to install the add-on on all the hosts using the specific platform. This option is useful for installing the add-on whenever a new host using the specific platform is added to the IMS. Select **Force install (Overwrites existing installation)** to overwrite existing add-on installation on the hosts.

Uninstalling add-ons from the hosts

You need to manually uninstall all the add-ons before you uninstall the host packages from the Infrastructure Management Server (IMS).

To uninstall add-ons from the hosts

1. Navigate

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

   (Or, click **Quick Actions** (menu bar) > **Add Assets**)

2. Under the data center, locate the IMS and click **Manage Assets**.
3. On the IMS **Settings** page, click **Deployment**.
4. Expand **Add-ons** to select the add-on that you want to uninstall.
5. In the **Add-ons** tab, right-click the add-on, and select **Uninstall**.
6. In the **Uninstall** panel, confirm the action of uninstalling the add-on from all the hosts. Select **Ignore checks (if any) before uninstalling** to ignore the checks before uninstalling.
7. Click Yes to uninstall the add-on from all the hosts.

If you try to uninstall the **Veritas Resiliency Platform Applications Enablement** add-on using Windows Programs and Features, you are asked to reboot the system. You can ignore this message.

If uninstallation of the **Veritas Resiliency Platform Applications Enablement** add-on fails and an error message to run the cleanup script is displayed, you need to run the following script on the managed host:

```
C:\ProgramData\Symantec\VRTSsfmh\spool\addons\store
\VRTSitrpappdr-1.0.0.0\CleanUp_Scripts\cleanup.bat
```

Managing solutions for the hosts

See “Installing a solution on the hosts” on page 97.

See “Uninstalling a solution from the hosts” on page 97.

Installing a solution on the hosts

Using the Resiliency Platform console, you can install a solution on the hosts that are added to the Infrastructure Management Server (IMS).

To install a solution on the hosts

1 Navigate

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

(Or, click **Quick Actions** (menu bar) > **Add Assets**)

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

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

4 Expand and right-click the solution, click **Install**.

5 In the **Install Select hosts** wizard panel, select the hosts on which you want to install the hot fix, package, or patch, and click **Finish**.

6 In the **Result** panel, click **OK**

Uninstalling a solution from the hosts

Using the Resiliency Platform console, you can uninstall a solution from the hosts.

To uninstall a solution from the hosts

1 Navigate

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

(Or, click **Quick Actions** (menu bar) > **Add Assets**)

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

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

4 Locate the solution, and right-click the solution that you want to uninstall, select **Uninstall**.

5 In the **Uninstall** panel, review the information. If you want to uninstall the solution from a specific host, select the host

Select **Ignore checks (if any) before uninstalling** to ignore the checks before uninstalling. Click **Yes**.

See “Installing a solution on the hosts” on page 97.

Chapter 8

Basic Resiliency Platform tasks

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 99.

See "Viewing and managing licenses" on page 100.

See "Viewing the License report" on page 101.

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 (Symantec license file). 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 100.

See “Viewing the License report” on page 101.

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 (Symantec license file).

You can view the following information about the installed licenses in a table:

- Name: Name of the license.
- Meter Type: Licenses for physical hosts are categorized under **Per Core** meter type whereas licenses for virtual machines are categorized under **Per Virtual Machine** meter type.
- 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 99.

See “Viewing the License report” on page 101.

Viewing the License 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 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** report to receive the report on the specified email address.

See “About licenses” on page 99.

See “Viewing and managing licenses” on page 100.

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 8-1 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 102. See "Configuring authentication domains " on page 108. See "Unconfiguring authentication domains" on page 110.
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 111.
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 103. See "Predefined personas" on page 104. See "About limiting object scope for personas" on page 107. See "Assigning permissions to user groups and users" on page 112. You can also create custom personas. See "Adding custom personas" on page 113. See "Predefined jobs that can be used for custom personas" on page 114.

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.

See "Managing user authentication and permissions" on page 101.

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).

Resiliency Platform comes with a set of predefined personas.

See "Predefined personas" on page 104.

You can also add custom personas.

See "Adding custom personas" on page 113.

See "Predefined jobs that can be used for custom personas" on page 114.

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 13. 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 Resiliency Platform 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 101.

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 8-2 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

Table 8-2 Predefined personas and jobs *(continued)*

Persona	Description and scope	Jobs
Resiliency Platform admin	Manage Resiliency Managers and Infrastructure Management Servers (IMSs) and data centers. Manage assets for the IMS. Manage user security settings and other product settings. Manage product updates. Scope: Resiliency domain.	Manage assets Manage user security settings Manage product settings Manage product updates Manage server deployments
Resiliency Platform Deployment admin	Manage Resiliency Managers and Infrastructure Management Servers (IMSs) and data centers. Manage product updates. Scope: Resiliency domain.	Manage product updates Manage server deployments
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 service (VBSs) Manage resiliency plan templates Manage resiliency plans
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 (VBSs) Start/stop resiliency groups

Table 8-2 Predefined personas and jobs *(continued)*

Persona	Description and scope	Jobs
DR admin	Configure all resiliency groups for disaster recovery (DR). Perform DR operations: migrate, takeover, rehearsal. Create, update, and delete resiliency plans and templates. Manage disaster recovery network settings. Start/stop all resiliency groups. Scope: Resiliency domain.	Manage disaster recovery of resiliency groups Perform disaster recovery of resiliency groups Manage resiliency plans Manage resiliency plan templates Manage disaster recovery network settings Start/stop resiliency groups
Resiliency Group DR admin	Manage and perform disaster recovery of 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.	Start/stop resiliency groups Manage disaster recovery of resiliency groups Perform disaster recovery of resiliency groups
Resiliency Group DR operator	Start/stop specified resiliency groups. Start/stop VBSs as long as the VBS contains only the specified resiliency groups. Perform disaster recovery on specified resiliency groups. Scope: Specified resiliency groups.	Start/stop resiliency groups Perform disaster recovery of resiliency groups

Table 8-2 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
Guest	View all information in console. Assigned by default when user or group is configured for Resiliency Platform.	No operations, only view permission

See "Managing user authentication and permissions" on page 101.

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 104.

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 start or stop 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

1 Prerequisites

The fully qualified domain name (FQDN) or IP address and credentials for the LDAP server

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**.

Note: To edit an existing authentication domain, right-click it and select the appropriate option.

4 Enter the information on the first wizard page and click **Next**.

See “Options for Configure Domain” on page 109.

5 Specify a friendly name for the authentication domain and select the applicable data centers. Click **Configure**.

6 Verify that the new domain is listed under **Domains**.

You can now configure user groups and users from that domain and assign permissions.

See “Managing user authentication and permissions” on page 101.

Options for Configure Domain

Table 8-3 Options for Configure Domain

Option	Description
Server Name (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.
This server requires 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=symantec,dc=com The LDAP or the AD administrator can provide you the bind user name that you can use.

Table 8-3 Options for Configure Domain *(continued)*

Option	Description
Password	Enter the password that is assigned to the bind user name that you use.
Use Secure Sockets Layer	Select this check box to use the Secure Sockets Layer (SSL) certificates to establish a secure channel between the authentication broker and the LDAP server.
Certificate location	Browse to the location of the trusted root CA certificate of the vendor that issued the LDAP server certificate.
Query Information:	
User Name (Mandatory)	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 Name	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 108.

See “Getting started with a new Resiliency Platform configuration” on page 35.

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 101.

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**.

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 112.

See "Managing user authentication and permissions" on page 101.

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 103.

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.

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 101.

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** > **Add Persona**

2 In the **Add 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 114.

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 101.

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 8-4 Jobs for custom personas

Job	Description	Scope
View all information	View all information in console.	Resiliency domain
Manage assets	Add assets to the IMS, remove assets that were added previously.	Resiliency domain or specific data centers
Manage user security settings	Manage authentication domains, users and user groups, personas.	Resiliency domain
Manage product settings	Manage product settings such as email/notification.	Resiliency domain

Table 8-4 Jobs for custom personas *(continued)*

Job	Description	Scope
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 resiliency groups	Create, update, and delete resiliency groups.	Resiliency domain or specific resiliency groups
Start/stop resiliency groups	Start/stop resiliency groups.	Resiliency domain or specific resiliency groups
Manage virtual business services (VBSs)	Create, update, and delete virtual business services (VBSs).	Resiliency domain or specific VBSs
Manage resiliency plans	Create, update, and delete resiliency plans	Resiliency domain
Manage resiliency plan templates	Create, update, and delete resiliency plan templates.	Resiliency domain
Manage disaster recovery of resiliency groups	Configure resiliency groups for disaster recovery (DR).	Resiliency domain or specific resiliency groups
Perform disaster recovery of resiliency groups	Perform DR operations: migrate, takeover, rehearsal	Resiliency domain or specific resiliency groups
Manage disaster recovery network settings	Configure disaster recovery network settings, for example, data center settings to be mapped for disaster recovery.	Resiliency domain

See “Predefined personas” on page 104.

See “Adding custom personas” on page 113.

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 116.

See "Managing report preferences" on page 116.

See "Scheduling a report" on page 119.

See "Running a report" on page 121.

See "Viewing and managing report schedules" on page 122.

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.
- **Risk 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.

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 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 116.

See "Scheduling a report" on page 119.

See "Running a report" on page 121.

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**.

3

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 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 119.

See “Running a report” on page 121.

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 116.

See “Running a report” on page 121.

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 119.

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 116.

See “Scheduling a report” on page 119.

See “Running a report” on page 121.

Managing settings for alerts and notifications and general product settings

See the following topics for information on configuring email and SNMP settings for notifications and reports, setting up rules for event notifications, and configuring purge settings for logs and traps, and some general product settings.

See "Adding, modifying, or deleting email settings" on page 124.

See "Adding, modifying, or deleting SNMP settings" on page 126.

See "Setting up rules for event notifications" on page 126.

See "Modifying purge settings for logs and SNMP traps" on page 127.

See "Enabling or disabling telemetry collection " on page 128.

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. The 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.

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. Click **Send Test Email** to test your settings.

6 Review the information in the summary and submit

See “Managing settings for alerts and notifications and general product settings” on page 124.

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 Veritas Resiliency Platform 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 general product settings" on page 124.

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 124.

See “Adding, modifying, or deleting SNMP settings” on page 126.

2 Navigate

⚙ **Settings** (menu bar)

Under **Product Settings**, select **Alerts & Notifications**

To add a new rule: Click the **Definition** tab > **New Rule**.

To modify or delete an existing rule: Click 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 Events	Select one or more events that you want to be notified about

4 Click **Submit**.

The rule is listed on the **Rules** tab.

Modifying purge settings 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 **General Settings**, click **General**

2 Under **Logs**, enter the new value for the purge settings, in months.

3 Click **Save**.

See “Managing settings for alerts and notifications and general product settings” on page 124.

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 **General Settings**, click **General**

2 Under **Telemetry**, select the setting to turn it on or off. To download a file showing the information that is collected, click **Show what is collected**.

Chapter 9

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.

About applying updates to Resiliency Platform

Updates to Veritas Resiliency Platform provide significant benefits, such as improved functionality, performance, security, and reliability.

- 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

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. You can apply the updates using the web console or using the CLISH menu.

The following figure shows how a repository server is used to apply the updates to Resiliency Platform:

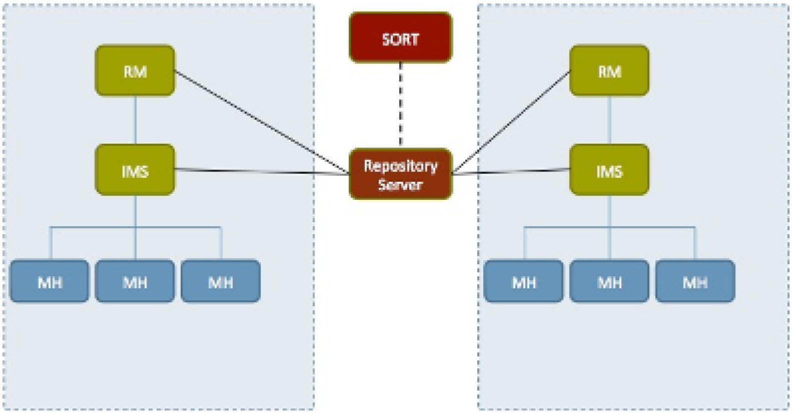

Note: While applying updates, you need to ensure that the virtual appliance remains powered on. Restarting the appliance during the process of applying updates may adversely affect the functionality.

The following is an overview of the process of applying updates in Veritas Resiliency Platform:

Table 9-1 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 131.
2	Set up a repository server and download the updates from SORT	See "Setting up the repository server and downloading updates" on page 132.
3	Add the repository server to Veritas Resiliency Platform. There can be multiple repository servers added to Veritas Resiliency Platform at a time.	See "Adding a repository server in Resiliency Platform" on page 133.
4	Assign a repository server to the virtual appliance where you want to apply the update. 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.	See "Assigning a repository server in Resiliency Platform" on page 134.
5	Apply updates using web console or using CLISH menu	See "Applying updates to virtual appliances using the console" on page 134. See "Applying updates to virtual appliance using CLISH menu" on page 135.
6	Refresh the information about applicable updates	See "Refreshing the information about applicable updates" on page 136.
7	Remove an update from the repository server	See "Removing an update from the repository server" on page 137.

Prerequisites for a repository server

To set up a repository server, make sure that the following prerequisites are met:

- Repository server should be any Linux server with minimum yum version 3.2.29.
- Perl and Python should be installed on the server. Perl modules `JSON.pm` and `Cinfig::Simple.pm` need to be installed on the Linux server.
- Web server (HTTP/HTTPS) should be configured on the server.
- Repository server should have minimum 50 GB disk space.

- Repository server should have connectivity with SORT as well as with the virtual appliances.

See “About applying updates to Resiliency Platform” on page 130.

Setting up the repository server and downloading updates

You need to set up a repository server in your environment, download the updates from SORT, and make them available on your repository server.

You have following three options to set up a repository and download the updates from SORT:

- Download a specified update or download all the updates released after a specified date.
- Download a specified update on your local system and then update the repository system with the downloaded updates. You can use this option if your repository server does not have SORT connectivity. To use this option, you need to download `master.xml` file from SORT.
- Download the metadata of the applicable updates to your repository server. Once you add the repository server using the Resiliency Manager console, you can view the list of applicable updates in the Resiliency Manager console. You can then decide which update you want to download.

Note: In case you plan to update the repository server with the updates or metadata that you have saved on you local system, you need to always use the latest `master.xml` file, irrespective of which update you plan to use.

To set up a repository server and download the updates

1. Create a repository path under root directory of the web server.

   mkdir *path_to_repository*

2. Download the `setup_conf_repo.pl` file from FileConnect.
3. Do one of the following:
   - To download the updates and update the repository server with those updates, do one of the following:
     - To download a particular update to the repository server, and update the repository:

```
./setup_conf_repo.pl --download-updates --repo-location
path_to_repository --product-version base_version --product
product_abbreviation --release-name release_name
```

  - To download multiple updates that are released after a particular date or after a particular update version, and update the repository:

    ```
    ./setup_conf_repo.pl --download-updates --repo-location
    path_to_repository product-version base_version --product
    product_abbreviation --start-date yyyy-mm-dd
    ```

- 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
  ```

- To download the metadata of the applicable updates to your repository server:

  ```
  ./setup_conf_repo.pl --refresh-metadata --repo-location
  path_to_repository --product-version base_version --product
  product_abbreviation
  ```

- To update the repository server with the metadata file `master.xml` that you have saved on your local system:

  ```
  ./setup_conf_repo.pl --refresh-metadata --repo-location
  path_to_repository --metadata-location path_to_master.xml
  ```

See “About applying updates to Resiliency Platform” on page 130.

Adding a repository server in Resiliency Platform

After configuring a repository server, you need to add the repository server in Veritas Resiliency Platform.

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 130.

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.

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 130.

Applying updates to virtual appliances using the console

You can apply updates to the virtual appliances using the console.

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 130.

Applying updates to virtual appliance using CLISH menu

You can use the CLISH menu to perform the upgrade related tasks in Resiliency Platform.

You need to log into the virtual appliance as admin and go to the updates sub-menu. Following is a list of commands that you can run to perform operations related to 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-repo
  ```

- 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_repo
  ```

For a complete list of options available with `Updates` command, seeSee “Using CLISH” on page 153.

See “About applying updates to Resiliency Platform” on page 130.

Refreshing the information about applicable updates

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 CLISH, you need to refresh the information to reflect the current status of the updates in the Resiliency Manager web console.

To refresh the information about applicable updates

1 Navigate

⚙ **Settings** (menu bar) > **Updates** > **Available Updates**

2 Click **Refresh**.

See “About applying updates to Resiliency Platform” on page 130.

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
     ```

Chapter 10

Uninstalling Resiliency Platform

This chapter includes the following topics:

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 39.

Chapter 11

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 126.

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 purge settings for logs and SNMP traps” on page 127.

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 CLISH menu to access the Resiliency Platform log files.

To access Resiliency Platform log files

1. Log in to the Resiliency Platform virtual appliance console 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 CLISH” on page 153.

   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 on any host connected to the virtual appliance. Log in as an admin user and download the zip file.

Use the `vomgather.pl` script to gather logs from the hosts:

- On Unix/Linux:

  ```
  /opt/VRTSsfmh/adm/vomgather.pl --dir mydir
  ```

- On Windows:

  ```
  "C:\Program Files\Veritas\VRTSsfmh\bin\perl.exe" "C:\Program
  Files\Veritas\VRTSsfmh\adm\vomgather.pl" --dir mydir
  ```

where, *`mydir`* is the directory to store gathered data.

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, such as the `xprtld` process being down on the Control Host, require intervention and troubleshooting to resolve.

You can display risks in the following ways:

Table 11-1 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.

Table 11-1 Ways to display risks *(continued)*

To display ...	Do the following:
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.

Table 11-2 describes each Resiliency Platform risk.

Table 11-2 Risks and Descriptions

Risk	Description
CTRL_HOST_DOWN	The `xprtld` process is down on the Control Host, and configured resources are in unknown state. Discovered contents can be stale.
HOST_SFMH_REINSTALLED	The host is disconnected. The probable cause is that the host has been reinstalled. Changes you make after this condition are not reflected on the Resiliency Manager. To correct this issue, remove and re-add this host to the Infrastructure Management Server (IMS).
HOST_DISCONNECTED_MAC_CHANGED	The host is disconnected. The probable cause is that the media access code (MAC) address of host has changed. Changes you make after this condition are not reflected on the Resiliency Manager. To correct this issue, remove and re-add this host to the Infrastructure Management Server (IMS).
VMWARE_DISCOVERY_FAILED	VMware discovery failed.
FS_FILESYSTEM_FULL	The file system is at 100% usage.

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 CLISH on the Resiliency Platform virtual appliance.

See "Using CLISH" on page 153.

Table 11-3 Resiliency Manager services

Service or component name	Description
Database service (DB)	Supports the Resiliency Platform 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 11-4

Components	Description
Operating System	Hardened CentOS 6.6 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 (Resiliency Platform)	Veritas Resiliency Platform (Resiliency Platform) that provides core and standard services framework for the solution.
Infrastructure Management Server (IMS)	Serves as the infrastructure manager or asset manager for Resiliency Platform RA.
Command Line Interface Shell (CLISH)	Command Line Interface Shell (CLISH) 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 27.

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 146.

If changes have been made to the asset infrastructure, you can use the Refresh operation on assets in the IMS to speed up discovery so that updated asset information is displayed more quickly in the console. To use the Refresh operation, display the asset infrastructure page for the IMS, select the asset type, right-click the asset and select Refresh.

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 11-5 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 11-6 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 11-6 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, than 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 11-7 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 to an Infrastructure Management Server (IMS)” on page 62.

About the frequency of asset information discovery

After you add the asset infrastructure, for example virtualization servers, to the Infrastructure Management Server (IMS), the 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 in the IMS to manually initiate the 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.

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 39.

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 Symantec Operations Readiness Tools to find a Unique Message Identifier description and solution

You can use Symantec 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.symantec.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.

Virtual appliance security features

This appendix includes the following topics:

About virtual appliance security

This chapter covers the following:

See “Operating system security” on page 150.

See “Management Security” on page 150.

See “Network security” on page 151.

See “Access control security” on page 151.

See “Physical security” on page 151.

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.

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. The new password must not be a dictionary word and must be at least six characters long. If the admin user password is lost, Symantec may access the root using the grub access, and reset the admin user password.

On successful completion of the Resiliency Platform bootstrap, admin user can only access a limited menu of commands through CLISH. 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 **CLISH**. An option `support > shell` is provided in the **CLISH** 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.

See “Network and firewall requirements” on page 24.

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 authorization messages are logged into the appliance.

Physical security

In the Resiliency Platform virtual appliance, the USB storage access is disabled.

Using CLISH menu in Resiliency Platform

This appendix includes the following topics:

- About CLISH
- Using CLISH

About CLISH

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 (CLISH) menu to manage the configuration-related changes to the product.

You can use CLISH 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 CLISH” on page 153.

Using CLISH

When you enter the Resiliency Platform CLISH menu, you enter in the main menu. 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 CLISH:

- **Network settings:** You can reconfigure the subnet mask, default gateway, DNS server, and search domains using the CLISH 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 CLISH menu. You can not 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 CLISH menu. Changing the system settings can impact the product functionality if incorrect values are set.

You can press the **tab** or **space** key to display the menu options. Press **?** key to display detailed help.

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
monitor	Monitor appliance activities Table B-5
network	Network configuration Table B-6
settings	Appliance settings Table B-12

Table B-1 Options available in the **main** menu *(continued)*

Menu option	Description
support	Access logs Table B-16
updates	Manage updates and patches Table B-18

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
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
services	Manage the appliance services Use **rm** or **ims** as first parameter and options available in the services menu as second parameter. Table B-4 See "About the Resiliency Manager services " on page 143.

Table B-3 Options available with **configure** command

Menu option	Description
ims	Configure Infrastructure Management Server
rm	Configure Resiliency Manager
show	Show the configured component

Table B-4 Options available with **services** command

Menu option	Description
show	Show Resiliency Platform services

Table B-4 Options available with **services** command *(continued)*

Menu option	Description
restart	Restart Resiliency Platform services Two options available are: `restart` *`all`* where, *all* means all the Resiliency Manager 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 Resiliency Manager 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 Resiliency Manager services. `status` *`service name`* where, *service name* is the name of a particular service. You can provide multiple service names (comma separated).
stop	Stop Resiliency Platform services Two options available are: `stop` *`all`* where, *all* means all the Resiliency Manager services. `stop` *`service name`* where, *service name* is the name of a particular service. You can provide multiple service names (comma separated).

Table B-5 Options available with **monitor** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
FSusage	Display filesystem usage
help	Display an overview of the CLI syntax

Table B-5 Options available with **monitor** command *(continued)*

Menu option	Description
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-6 Options available with **network** command

Menu option	Description
back	Return to the previous menu
dns	Show or change the DNS Table B-7
exit	Log out from the current CLI session
gateway	Show or change the Gateway Table B-8
help	Display an overview of the CLI syntax
hostname	Show the hostname
ip	Show or change the IP address Table B-9
netmask	Show or change the netmask Table B-10
search-domain	Show or change the domain Table B-11

Table B-7 Options available with **dns** command

Menu option	Description
set	Configure domain name server
show	Show the current domain name server

Table B-8 Options available with **gateway** command

Menu option	Description
set	Configure Gateway
show	Show the current Gateway

Table B-9 Options available with **ip** command

Menu option	Description
set	Configure the IP address
show	Show the current IP address

Table B-10 Options available with **netmask** command

Menu option	Description
set	Configure the netmask
show	Show the current netmask

Table B-11 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-12 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-13
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax

Table B-12 Options available with **settings** command *(continued)*

Menu option	Description
lvm	Perform operations related to Logical Volume Manager on the appliance See Table B-14 on page 158.
ntp	Perform operations related to NTP server
poweroff	Shutdown the appliance
reboot	Restart the appliance
timezone	Show or change the timezone for the appliance See Table B-15 on page 158.

Table B-13 Options available with **date** command

Menu option	Description
show	Show the time and date

Table B-14 Options available with **lvm** command

Menu option	Description
add-disk	Add disk to the data volume
list-free-disk	Lists free disks
list-used-disk	Lists disks used by the data volume
remove-disk	Remove disk from the data volume

Table B-15 Options available with **timezone** command

Menu option	Description
set	Set the timezone for the appliance
show	Show the current timezone for the appliance

Table B-16 Options available with **support** command

Menu option	Description
back	Return to the previous menu

Table B-16 Options available with **support** command *(continued)*

Menu option	Description
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
loggather	Collect Operating System, Resiliency Manager and IMS logs Table B-17
shell	Open the bash shell prompt for support user

Table B-17 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-18 Options available with **updates** command

Menu option	Description
apply-update	Apply the specified update
back	Return to the previous menu
config-repository	Configure the repository Table B-19
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

Table B-18 Options available with **updates** command *(continued)*

Menu option	Description
show-version	Show appliance version

Table B-19 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 CLISH” on page 152.

See “Accessing Resiliency Platform log files” on page 140.

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 162.
2 - Navigation pane	Icons to open pages for configuring and implementing start/stop and disaster recovery operations. See "Navigation pane options" on page 164.

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 167.

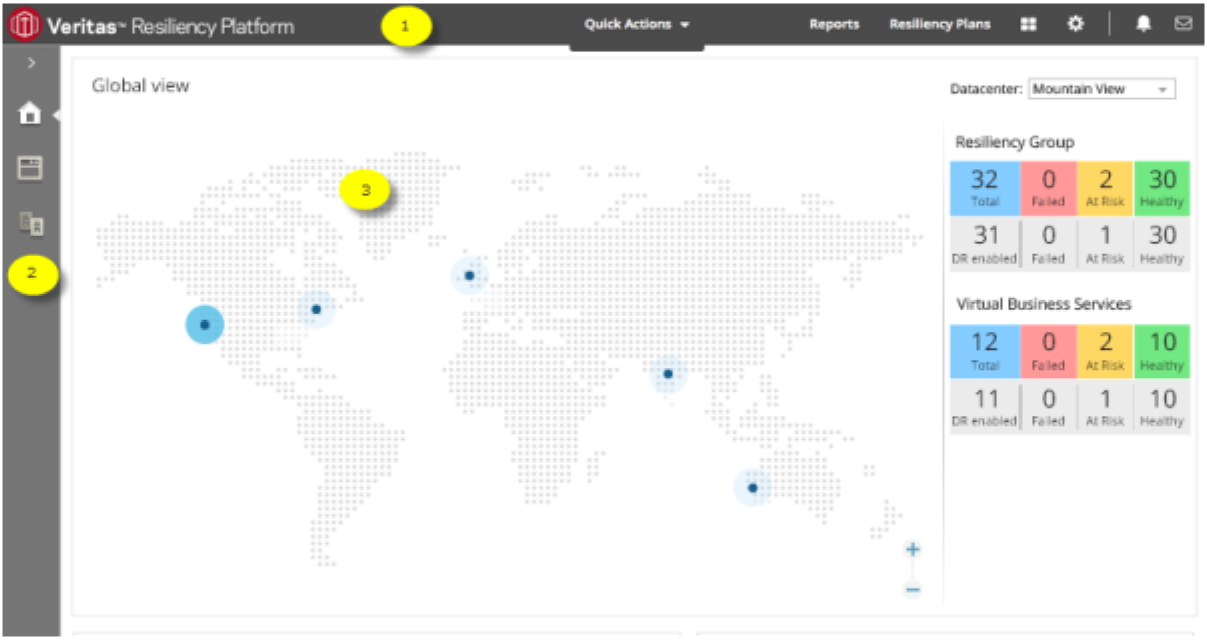

Resiliency Group

32 Total	0 Failed	2 At Risk	30 Healthy
31 DR enabled	0 Failed	1 At Risk	30 Healthy

Virtual Business Services

12 Total	0 Failed	2 At Risk	10 Healthy
11 DR enabled	0 Failed	1 At Risk	10 Healthy

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 164.
Reports	Schedule and run reports. View reports showing data center and asset status. See "About reports" on page 116.
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.
⚙	Settings Open Settings page for configuring and maintaining product infrastructure and other settings. See "About settings in the web console" on page 165.
🔔	Notifications Display most recent notifications. Requires alerts and notifications to be enabled using Settings page. See "Viewing events and logs in the console" on page 139. See "Managing settings for alerts and notifications and general product settings" on page 124.
✉	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.

Note: 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 Assets page for configuring and viewing resiliency groups and performing start/stop operations or disaster recovery
🗄	Opens configuration page for disaster recovery settings

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.

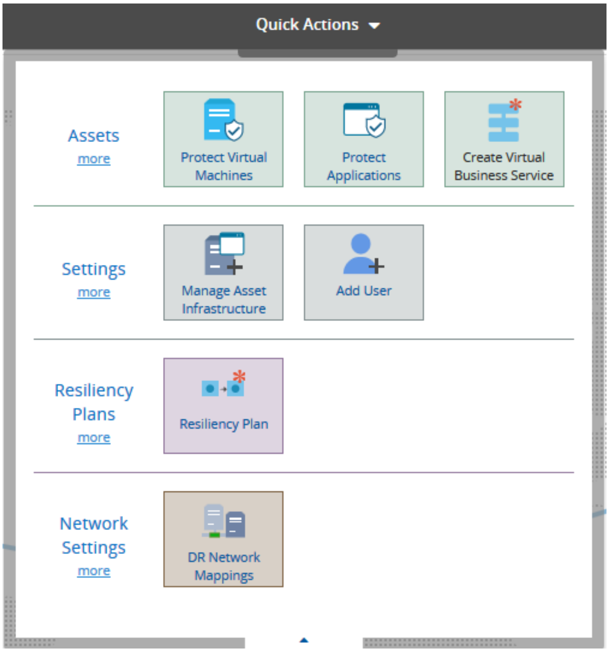

Filtering and searching for objects in the web console

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 general settings. You access the Settings page from the menu bar.

⚙ Settings

Table C-4 Settings page options

Type of setting	Description	More information
Infrastructure	Add and manage Infrastructure Management Servers (IMS) and add the asset infrastructure to the IMS Manage data centers and their network settings Manage Resiliency Managers and resiliency domains	See “Managing data centers” on page 48. See “Removing a Resiliency Manager from a resiliency domain” on page 39.
Updates	View and deploy product updates	See “About applying updates to Resiliency Platform” on page 130.
User Management	Configure authentication domains, users and user groups, and user personas (roles)	See “Managing user authentication and permissions” on page 101.
Licenses	View and manage product licenses	See “Managing licenses” on page 99.
Alerts and Notifications	Configure email and SNMP for alerts and notifications, configure rules for notification	See “Managing settings for alerts and notifications and general product settings” on page 124.
General	Configure general product settings such as purge intervals and telemetry	See “Managing settings for alerts and notifications and general product settings” on page 124.

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 161.

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 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.

Virtual Machines by Type and Platform

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.

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.

Virtual Machines and Applications by Recovery Readiness	Displays the percentage of virtual machines and applications that are unprotected or unmanaged. Use the drop-down list to filter your results.

You can use the Assets icon in the navigation pane to display more detailed information on resiliency groups.

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
🔔	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

Table C-5 Web console icons *(continued)*

Icon	Description	Location
	Assets Opens the Assets page for configuring and viewing resiliency groups and performing start/stop operations or disaster recovery	Navigation pane
	Recovery Automation Opens configuration page for disaster recovery settings	Navigation pane
⋮	Vertical ellipsis Displays list of actions for selected object	To the right of a selected object in a list

Appendix D

Commands used in storage array monitoring and discovery

This appendix includes the following topics:

- EMC Symmetrix storage enclosure commands
- NetApp storage enclosure commands

EMC Symmetrix storage enclosure commands

The following commands are used to discover the storage objects of EMC Symmetrix enclosure. The used SYMCLI version should be 7.0, or later.

- Command to find the version of the installed Solutions Enabler:

  ```
  symcli -out XML
  ```

- Command to find the number of last entry in the audit log of the enclosure:

  ```
  symaudit -sid arrayid show -out XML
  ```

- Command to find the number in the audit log of the enclosure based on the specified filter parameters.

  ```
  symaudit -sid arrayid list -action_code action_codes
  -function_class function_classes -record_num record_no -out XML
  ```

- Command to get the information about the Symmetrix configuration:

  ```
  symcfg -sid arrayid list -out XML
  ```

- Command to get the detailed information about the Symmetrix configuration:

  ```
  symcfg -sid arrayid list -v -out XML
  ```

- Command to get a brief information about the physical disks:

  ```
  symdisk -sid arrayid list -out XML
  ```

- Command to get the detailed information about the physical disks:

  ```
  symdisk -sid arrayid list -v -out XML
  ```

- Command to get the list of all Symmetrix devices:

  ```
  symdev -sid arrayid list -all -out XML
  ```

- Command to get the detailed information of all Symmetrix meta-head devices:

  ```
  symdev -sid arrayid list -meta -v -out XML
  ```

- Command to get the detailed information of Symmetrix devices specified in dev_list:

  ```
  symdev -sid arrayid list -v -devs dev_list -out XML
  ```

- Command to get the list of all Symmetrix devices mapped to front-end directors:

  ```
  symdev -sid arrayid -SA all list -out XML
  ```

- Command to get the list of all RDF groups:

  ```
  symcfg -sid arrayid list -rdfg ALL -out XML
  ```

- Command to get the details of a given RDF group:

  ```
  symrdf -sid arrayid list -rdfg rdf_group_no -out XML
  ```

- Command to get the list of all BCV sessions created on Symmetrix:

  ```
  symmir -sid arrayid list -out XML
  ```

- Command to get the list of all TimeFinder/Clone sessions created on Symmetrix:

  ```
  symclone -sid arrayid list -v -out XML
  ```

- Command to get the list of all TimeFinder/Snap sessions created on Symmetrix:

  ```
  symsnap -sid arrayid list -v -out XML
  ```

- Command to get brief information about thin pools in given array ID:

  ```
  symcfg -sid arrayid list -pool -thin -detail -mb -out XML
  ```

- Command to get the detailed information about the thin pools in given array ID.

  ```
  symcfg -sid arrayid list -pool -thin -detail -v -mb -out XML
  ```

- Command to get the policy association of a given storage group:

  ```
  symfast -sid arrayid show -association -sg sg_name -out XML
  ```

- Command to get the information on FAST policies:

  ```
  symfast -sid arrayid list -fp -v -out XML
  ```

- Command to get the detailed information on FAST tiers:

  ```
  symtier -sid arrayid list -v -out XML
  ```

- Command to get the list of all directors:

```
symcfg -sid arrayid list -DIR ALL -out XML
```

- Command to get the detailed information of all front-end directors:

  ```
  symcfg -sid arrayid list -SA ALL -v -out XML
  ```

- Command to get the detailed information of all Fibre front-end directors:

  ```
  symcfg -sid arrayid list -FA ALL -v -out XML
  ```

- Command to get the detailed information of all FICON directors:

  ```
  symcfg -sid arrayid list -EF ALL -v -out XML
  ```

- Command to list the records within the device masking VCMDB:

  ```
  symmaskdb -sid arrayid list database -out XML
  ```

- Command to list the devices assigned in the device masking VCMDB (applies WWN):

  ```
  symmaskdb -sid arrayid list devs -wwn hba_port -out XML
  ```

- Command to list the devices assigned by records in the device masking VCMDB:

  ```
  symmaskdb -sid arrayid list devs -out XML
  ```

- Command to list the device information by the initiator group:

  ```
  symaccess -sid arrayid list devinfo -out XML
  ```

See “Adding storage enclosures for discovery by the IMS” on page 85.

NetApp storage enclosure commands

The following commands are used to discover various storage objects of NetApp enclosure. The NetApp Data ONTAP version must be 1.4, or later.

- `system-get-info`, `system-get-version`: To discover enclosure details.
- `vfiler-list-info`: To discover Vfilers.
- `disk-list-info`: To discover physical storage information.
- `fcp-adapter-list-info`, `iscsi-adapter-list-info`: To discover adapters.
- `aggr-list-info`: To discover aggregates.
- `volume-list-info`: To discover storage array volumes (FlexVol and traditional).
- `lun-list-info, lun-map-list-info`, `lun-get-serial-number`: To discover LUN details and mapping information.
- `snapmirror-get-status`, `snapshot-list-info`: To discover snapshot details.
- `qtree-list`: To discover array file systems (Qtrees).

- `cifs-share-list-iter-start`, `cifs-share-list-iter-next`: To discover CIFS shares.
- `nfs-exportfs-list-rules`: To discover NFS Shares.
- `quota-report-iter-start`, `quota-report-iter-next`: To discover Quotas.

The following command is used to discover metering and performance statistics for NetApp enclosures:

- `perf-object-get-instances`

See “Adding storage enclosures for discovery by the IMS” on page 85.

Glossary

activity A task or an operation performed on a resiliency group.

add-on An additional software package that can be installed on hosts by the Infrastructure Management Server (IMS) for specialized uses.

asset infrastructure The data center assets that can be added to the Infrastructure Management Server (IMS) for IMS discovery and monitoring. For example, virtual machines or virtualization servers.

assets In Veritas Resiliency Platform, the virtual machines or applications that have been discovered by the Infrastructure Management Server (IMS) and that can be grouped into resiliency groups.

CLISH Command Line Interface SHell. Provides the command line menu on the Veritas Resiliency Platform virtual appliance for use after the initial bootstrap configuration..

data center A location that contains asset infrastructure to be managed by Veritas Resiliency Platform.

For the 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.

host Physical servers, virtual machines, or Hyper-V servers that are added to the Infrastructure Management Server (IMS) as hosts.

Adding the assets as hosts installs the host package that is used by the IMS for discovery and monitoring.

Infrastructure Management Server (IMS) The Veritas Resiliency Platform 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.

migrate A planned activity involving graceful shutdown of virtual machines at the production data center and starting them at the recovery data center. In this process, replication ensures that consistent virtual machine data is made available at the recovery data center.

persona A user role that has access to a predefined set of jobs (operations). Used to assign permissions to users and groups for Veritas Resiliency Platform web console operations.

product role The function configured for a Veritas Resiliency Platform virtual appliance.

For example, a virtual appliance can be configured as a Resiliency Manager, Infrastructure Management Server (IMS) or both.

production data center The data center that is normally used for business. See also recovery data center.

recovery data center The data center that is used if a disaster scenario occurs. See also production data center.

rehearsal A zero-downtime test that mimics the configuration, application data, storage, and the failover behavior of the resiliency group.

Rehearsal verifies the ability of the resiliency group to fail over to the recovery data center during a disaster.

resiliency domain The logical scope of a Resiliency Platform deployment. It can extend across multiple data centers.

resiliency group The unit of management and control in Veritas Resiliency Platform. Related assets are organized into a resiliency group and managed and monitored as a single entity.

Resiliency Manager The Veritas Resiliency Platform component that provides resiliency capabilities within a resiliency domain. It is composed of loosely coupled services, a distributed data repository, and a management console.

resiliency plan A collection of tasks or operations, along with the relevant assets, which are performed in a predefined sequence.

resiliency plan template A template defining the execution sequence of a collection of tasks or operations.

takeover An activity initiated by a user when the production data center is down due to a disaster and the virtual machines need to be restored at the recovery data center to provide business continuity.

tier Within a virtual business service (VBS), resiliency groups are arranged as tiers. Tiers represent the logical dependencies between the resiliency groups and determine the relative order in which the resiliency groups start and stop.

virtual appliance An appliance that includes the operating system environment and the software application which are deployed together as a virtual machine.

The Veritas Resiliency Platform virtual appliance is deployed as a virtual machine and then configured with basic settings and a role (for example, Resiliency Manager).

virtual business service (VBS) A multi-tier IT service where each VBS tier hosts one or more resiliency groups. A VBS groups multiple services as a single unit for visualization, automation, and controlled start and stop in the desired order. You can also migrate/takeover the entire VBS.

web console The web-based management console on the Resiliency Manager that is used to configure the settings for the resiliency domain and perform operations.

Index